

STATA MULTILEVEL MIXED-EFFECTS REFERENCE MANUAL RELEASE 15



A Stata Press Publication
StataCorp LLC
College Station, Texas



Copyright © 1985–2017 StataCorp LLC
All rights reserved
Version 15

Published by Stata Press, 4905 Lakeway Drive, College Station, Texas 77845
Typeset in $\text{T}_{\text{E}}\text{X}$

ISBN-10: 1-59718-244-3

ISBN-13: 978-1-59718-244-7

This manual is protected by copyright. All rights are reserved. No part of this manual may be reproduced, stored in a retrieval system, or transcribed, in any form or by any means—electronic, mechanical, photocopy, recording, or otherwise—without the prior written permission of StataCorp LLC unless permitted subject to the terms and conditions of a license granted to you by StataCorp LLC to use the software and documentation. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document.

StataCorp provides this manual “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. StataCorp may make improvements and/or changes in the product(s) and the program(s) described in this manual at any time and without notice.

The software described in this manual is furnished under a license agreement or nondisclosure agreement. The software may be copied only in accordance with the terms of the agreement. It is against the law to copy the software onto DVD, CD, disk, diskette, tape, or any other medium for any purpose other than backup or archival purposes.

The automobile dataset appearing on the accompanying media is Copyright © 1979 by Consumers Union of U.S., Inc., Yonkers, NY 10703-1057 and is reproduced by permission from CONSUMER REPORTS, April 1979.

Stata, **STATA** Stata Press, Mata, **MATA** and NetCourse are registered trademarks of StataCorp LLC.

Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations.

NetCourseNow is a trademark of StataCorp LLC.

Other brand and product names are registered trademarks or trademarks of their respective companies.

For copyright information about the software, type `help copyright` within Stata.

The suggested citation for this software is

StataCorp. 2017. *Stata: Release 15*. Statistical Software. College Station, TX: StataCorp LLC.

Contents

me	Introduction to multilevel mixed-effects models	1
estat df	Calculate degrees of freedom for fixed effects	48
estat group	Summarize the composition of the nested groups	52
estat icc	Estimate intraclass correlations	53
estat recovariance	Display estimated random-effects covariance matrices	56
estat sd	Display variance components as standard deviations and correlations	58
estat wcorrelation	Display within-cluster correlations and standard deviations	60
mecloglog	Multilevel mixed-effects complementary log-log regression	66
mecloglog postestimation	Postestimation tools for mecloglog	79
meglm	Multilevel mixed-effects generalized linear model	85
meglm postestimation	Postestimation tools for meglm	116
meintreg	Multilevel mixed-effects interval regression	134
meintreg postestimation	Postestimation tools for meintreg	145
melogit	Multilevel mixed-effects logistic regression	154
melogit postestimation	Postestimation tools for melogit	178
menbreg	Multilevel mixed-effects negative binomial regression	190
menbreg postestimation	Postestimation tools for menbreg	207
menl	Nonlinear mixed-effects regression	213
menl postestimation	Postestimation tools for menl	295
meologit	Multilevel mixed-effects ordered logistic regression	309
meologit postestimation	Postestimation tools for meologit	323
meoprobit	Multilevel mixed-effects ordered probit regression	330
meoprobit postestimation	Postestimation tools for meoprobit	344
mepoisson	Multilevel mixed-effects Poisson regression	351
mepoisson postestimation	Postestimation tools for mepoisson	368
meprobit	Multilevel mixed-effects probit regression	375
meprobit postestimation	Postestimation tools for meprobit	388
meqrlogit	Multilevel mixed-effects logistic regression (QR decomposition)	394
meqrlogit postestimation	Postestimation tools for meqrlogit	409
meqrpoisson	Multilevel mixed-effects Poisson regression (QR decomposition)	417
meqrpoisson postestimation	Postestimation tools for meqrpoisson	434
mestreg	Multilevel mixed-effects parametric survival models	441

mestreg postestimation Postestimation tools for mestreg 466

metobit Multilevel mixed-effects tobit regression 478

metobit postestimation Postestimation tools for metobit 489

mixed Multilevel mixed-effects linear regression 496

mixed postestimation Postestimation tools for mixed 562

Glossary 584

Subject and author index 590

Cross-referencing the documentation

When reading this manual, you will find references to other Stata manuals. For example,

[U] [26 Overview of Stata estimation commands](#)

[R] [regress](#)

[D] [reshape](#)

The first example is a reference to chapter 26, *Overview of Stata estimation commands*, in the *User's Guide*; the second is a reference to the `regress` entry in the *Base Reference Manual*; and the third is a reference to the `reshape` entry in the *Data Management Reference Manual*.

All the manuals in the Stata Documentation have a shorthand notation:

[GSM]	<i>Getting Started with Stata for Mac</i>
[GSU]	<i>Getting Started with Stata for Unix</i>
[GSW]	<i>Getting Started with Stata for Windows</i>
[U]	<i>Stata User's Guide</i>
[R]	<i>Stata Base Reference Manual</i>
[BAYES]	<i>Stata Bayesian Analysis Reference Manual</i>
[D]	<i>Stata Data Management Reference Manual</i>
[ERM]	<i>Stata Extended Regression Models Reference Manual</i>
[FMM]	<i>Stata Finite Mixture Models Reference Manual</i>
[FN]	<i>Stata Functions Reference Manual</i>
[G]	<i>Stata Graphics Reference Manual</i>
[IRT]	<i>Stata Item Response Theory Reference Manual</i>
[DSGE]	<i>Stata Linearized Dynamic Stochastic General Equilibrium Reference Manual</i>
[XT]	<i>Stata Longitudinal-Data/Panel-Data Reference Manual</i>
[ME]	<i>Stata Multilevel Mixed-Effects Reference Manual</i>
[MI]	<i>Stata Multiple-Imputation Reference Manual</i>
[MV]	<i>Stata Multivariate Statistics Reference Manual</i>
[PSS]	<i>Stata Power and Sample-Size Reference Manual</i>
[P]	<i>Stata Programming Reference Manual</i>
[SP]	<i>Stata Spatial Autoregressive Models Reference Manual</i>
[SEM]	<i>Stata Structural Equation Modeling Reference Manual</i>
[SVY]	<i>Stata Survey Data Reference Manual</i>
[ST]	<i>Stata Survival Analysis Reference Manual</i>
[TS]	<i>Stata Time-Series Reference Manual</i>
[TE]	<i>Stata Treatment-Effects Reference Manual: Potential Outcomes/Counterfactual Outcomes</i>
[I]	<i>Stata Glossary and Index</i>
[M]	<i>Mata Reference Manual</i>

Description	Quick start	Syntax	Remarks and examples
Acknowledgments	References	Also see	

Description

Mixed-effects models are characterized as containing both fixed effects and random effects. The fixed effects are analogous to standard regression coefficients and are estimated directly. The random effects are not directly estimated (although they may be obtained postestimation) but are summarized according to their estimated variances and covariances. Random effects may take the form of either random intercepts or random coefficients, and the grouping structure of the data may consist of multiple levels of nested groups. As such, mixed-effects models are also known in the literature as multilevel models and hierarchical models. Mixed-effects commands fit mixed-effects models for a variety of distributions of the response conditional on normally distributed random effects.

Mixed-effects linear regression

<code>mixed</code>	Multilevel mixed-effects linear regression

Mixed-effects generalized linear model

<code>meglm</code>	Multilevel mixed-effects generalized linear model

Mixed-effects censored regression

<code>metobit</code>	Multilevel mixed-effects tobit regression
<code>meintreg</code>	Multilevel mixed-effects interval regression

Mixed-effects binary regression

<code>melogit</code>	Multilevel mixed-effects logistic regression
<code>meqrlogit</code>	Multilevel mixed-effects logistic regression (QR decomposition)
<code>meprobit</code>	Multilevel mixed-effects probit regression
<code>mecloglog</code>	Multilevel mixed-effects complementary log-log regression

Mixed-effects ordinal regression

<code>meologit</code>	Multilevel mixed-effects ordered logistic regression
<code>meoprobit</code>	Multilevel mixed-effects ordered probit regression

Mixed-effects count-data regression

<code>mepoisson</code>	Multilevel mixed-effects Poisson regression
<code>meqrpoisson</code>	Multilevel mixed-effects Poisson regression (QR decomposition)
<code>menbreg</code>	Multilevel mixed-effects negative binomial regression

Mixed-effects multinomial regression

Although there is no `memlogit` command, multilevel mixed-effects multinomial logistic models can be fit using `gsem`; see [SEM] [example 41g](#).

Mixed-effects survival model

`mestreg` Multilevel mixed-effects parametric survival models

Nonlinear mixed-effects regression

`menl` Nonlinear mixed-effects regression

Postestimation tools specific to mixed-effects commands

<code>estat df</code>	Calculate and display degrees of freedom for fixed effects
<code>estat group</code>	Summarize the composition of the nested groups
<code>estat icc</code>	Estimate intraclass correlations
<code>estat recovariance</code>	Display the estimated random-effects covariance matrices
<code>estat sd</code>	Display variance components as standard deviations and correlations
<code>estat wcorrelation</code>	Display model-implied within-cluster correlations and standard deviations

Quick start*Linear mixed-effects models*

Linear model of *y* on *x* with random intercepts by *id*

```
mixed y x || id:
```

Three-level linear model of *y* on *x* with random intercepts by *doctor* and *patient*

```
mixed y x || doctor: || patient:
```

Linear model of *y* on *x* with random intercepts and coefficients on *x* by *id*

```
mixed y x || id: x
```

Same model with covariance between the random slope and intercept

```
mixed y x || id: x, covariance(unstructured)
```

Linear model of *y* on *x* with crossed random effects for *id* and *week*

```
mixed y x || _all: R.id || _all: R.week
```

Same model specified to be more computationally efficient

```
mixed y x || _all: R.id || week:
```

Full factorial repeated-measures ANOVA of *y* on *a* and *b* with random effects by *field*

```
mixed y a##b || field:
```

Generalized linear mixed-effects models

Logistic model of y on x with random intercepts by id , reporting odds ratios

```
melogit y x || id: , or
```

Same model specified as a GLM

```
meglm y x || id:, family(bernoulli) link(logit)
```

Three-level ordered probit model of y on x with random intercepts by $doctor$ and $patient$

```
meoprobit y x || doctor: || patient:
```

Nonlinear mixed-effects models

Nonlinear mixed-effects regression of y on x_1 and x_2 with parameters $\{b_0\}$, $\{b_1\}$, $\{b_2\}$, and $\{b_3\}$ and random intercepts U_0 by id

```
menl y = ({b0}+{b1}*x1+{U0[id]})/(1+exp(-(x2-{b2})/{b3}))
```

As above, but using the more efficient specification of the linear combination

```
menl y = ({lc: x1 U0[id]})/(1+exp(-(x2-{b2})/{b3}))
```

As above, but using `define()` to specify the linear combination

```
menl y = {lc:}/(1+exp(-(x2-{b2})/{b3})), define(lc: x1 U0[id])
```

Include a random slope on continuous variable x_1 in the `define()` option, and allow correlation between random slopes U_1 and intercepts U_0

```
menl y = {lc:}/(1+exp(-(x2-{b2})/{b3})),          ///
      define(lc: x1 U0[id] c.x1#U1[id]) covariance(U0 U1, unstructured)
```

Specify a heteroskedastic within-subject error structure that varies as a power of predicted mean values `_yhat`

```
menl y = {lc:}/(1+exp(-(x2-{b2})/{b3})),          ///
      define(lc: x1 U0[id] c.x1#U1[id])            ///
      covariance(U0 U1, unstructured) resvariance(power _yhat)
```

Three-level nonlinear regression of y on x_1 with random intercepts W_0 and slopes W_1 on continuous x_1 by lev_2 and with random intercepts S_0 and slopes S_1 on x_1 by lev_3 , with lev_2 nested within lev_3 , using unstructured covariance for W_0 and W_1 and exchangeable covariance for S_0 and S_1

```
menl y = {phi1:}+{b1}*cos({b2}*x1),                ///
      define(phi1: x1 W0[lev3] S0[lev3>lev2]        ///
        c.x1#(W1[lev3] S1[lev3>lev2]))              ///
      covariance(W0 W1, unstructured) covariance(S0 S1, exchangeable)
```

Syntax

Linear mixed-effects models

```
mixed depvar fe_equation [|| re_equation] [|| re_equation ...] [, options]
```

where the syntax of the fixed-effects equation, *fe_equation*, is

```
[indepvars] [if] [in] [weight] [, fe_options]
```

and the syntax of a random-effects equation, *re_equation*, is the same as below for a generalized linear mixed-effects model.

Generalized linear mixed-effects models

```
mecmd depvar fe_equation [|| re_equation] [|| re_equation ...] [, options]
```

where the syntax of the fixed-effects equation, *fe_equation*, is

```
[indepvars] [if] [in] [, fe_options]
```

and the syntax of a random-effects equation, *re_equation*, is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [, re_options]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

Nonlinear mixed-effects models

```
menl depvar = <menexpr> [if] [in] [, options]
```

<*menexpr*> defines a nonlinear regression function as a substitutable expression that contains model parameters and random effects specified in braces {}, as in `exp({b}+{U[id]})`; see [Random-effects substitutable expressions](#) in [ME] **menl** for details.

Remarks and examples

Remarks are presented under the following headings:

- [Introduction](#)
- [Using mixed-effects commands](#)
- [Mixed-effects models](#)
 - [Linear mixed-effects models](#)
 - [Generalized linear mixed-effects models](#)
 - [Survival mixed-effects models](#)
 - [Nonlinear mixed-effects models](#)
 - [Alternative mixed-effects model specification](#)
 - [Likelihood calculation](#)
 - [Computation time and the Laplacian approximation](#)
 - [Diagnosing convergence problems](#)
 - [Distribution theory for likelihood-ratio test](#)
- [Examples](#)
 - [Two-level models](#)
 - [Covariance structures](#)
 - [Three-level models](#)
 - [Crossed-effects models](#)
 - [Nonlinear models](#)

Introduction

Multilevel models have been used extensively in diverse fields, from the health and social sciences to econometrics. Mixed-effects models for binary outcomes have been used, for example, to analyze the effectiveness of toenail infection treatments ([Lesaffre and Spiessens 2001](#)) and to model union membership of young males ([Vella and Verbeek 1998](#)). Ordered outcomes have been studied by, for example, [Tutz and Hennevogel \(1996\)](#), who analyzed data on wine bitterness, and [De Boeck and Wilson \(2004\)](#), who studied verbal aggressiveness. For applications of mixed-effects models for count responses, see, for example, the study on police stops in New York City ([Gelman and Hill 2007](#)) and the analysis of the number of patents ([Hall, Griliches, and Hausman 1986](#)). [Rabe-Hesketh and Skrondal \(2012\)](#) provide more examples of linear and generalized linear mixed-effects models. Nonlinear mixed-effects (NLME) models are popular in, for example, population pharmacokinetics, bioassays, and studies of biological and agricultural growth processes.

For a comprehensive treatment of mixed-effects models, see, for example, [Searle, Casella, and McCulloch \(1992\)](#); [Verbeke and Molenberghs \(2000\)](#); [Raudenbush and Bryk \(2002\)](#); [Hedeker and Gibbons \(2006\)](#); [McCulloch, Searle, and Neuhaus \(2008\)](#); and [Rabe-Hesketh and Skrondal \(2012\)](#). For NLME models, see, for example, [Davidian and Giltinan \(1995\)](#); [Vonesh and Chinchilli \(1997\)](#); [Demidenko \(2013\)](#); [Pinheiro and Bates \(2000\)](#); and [Davidian and Giltinan \(2003\)](#).

Using mixed-effects commands

Below we summarize general capabilities of the mixed-effects commands. We let *mecmd* stand for any mixed-effects command, such as *mixed*, *melogit*, or *meprobit*, except *menl*. *menl* models the mean function nonlinearly and thus has a different syntax; see [\[ME\] menl](#).

1. Fit a two-level random-intercept model with *levelvar* defining the second level:

```
. mecmd depvar [indepvars] ... || levelvar: , ...
```

2. Fit a two-level random-coefficients model containing the random-effects covariates *revars* at the level *levelvar*:

```
. mecmod depvar [indepvars] ... || levelvar: revars, ...
```

This model assumes an independent covariance structure between the random effects; that is, all covariances are assumed to be 0. There is no statistical justification, however, for imposing any particular covariance structure between random effects at the onset of the analysis. In practice, models with an unstructured random-effects covariance matrix, which allows for distinct variances and covariances between all random-effects covariates (*revars*) at the same level, must be explored first; see [Other covariance structures](#) and [example 3](#) in [ME] [melogit](#) for details.

Stata's commands use the default independent covariance structure for computational feasibility. Numerical methods for fitting mixed-effects models are computationally intensive—computation time increases significantly as the number of parameters increases; see [Computation time and the Laplacian approximation](#) for details. The unstructured covariance is the most general and contains many parameters, which may result in an unreasonable computation time even for relatively simple random-effects models. Whenever feasible, however, you should start your statistical analysis by fitting mixed-effects models with an unstructured covariance between random effects, as we show next.

3. Specify the unstructured covariance between the random effects in the above:

```
. mecmod depvar [indepvars] ... || levelvar: revars, covariance(unstructured) ...
```

4. Fit a three-level nested model with *levelvar1* defining the third level and *levelvar2* defining the second level:

```
. mecmod depvar [indepvars] ... || levelvar1: || levelvar2:, ...
```

5. Fit the above three-level nested model as a two-level model with exchangeable covariance structure at the second level (*mixed*, *meqrlogit*, and *meqrpoisson* only):

```
. mecmod depvar [indepvars] ... || levelvar1: R.levelvar2, cov(exchangeable) ...
```

See [example 11](#) in [ME] [mixed](#) for details about this equivalent specification. This specification may be useful for a more efficient fitting of random-effects models with a mixture of crossed and nested effects.

6. Fit higher-level nested models:

```
. mecmod depvar [indepvars] ... || levelvar1: || levelvar2: || levelvar3: || ...
```

7. Fit a two-way crossed-effects model with the `_all:` notation for each random-effects equation:

```
. mecmod depvar [indepvars] ... || _all: R.factor1 || _all: R.factor2 ...
```

When you use the `_all:` notation for each random-effects equation, the total dimension of the random-effects design equals $r_1 + r_2$, where r_1 and r_2 are the numbers of levels in *factor1* and *factor2*, respectively. This specification may be infeasible for some mixed-effects models; see item 8 below for a more efficient specification of this model.

8. Fit a two-way crossed-effects model with the `_all:` notation for the first random-effects equation only:

```
. mecmod depvar [indepvars] ... || _all: R.factor1 || factor2:, ...
```

Compared with the specification in item 7, this specification requires only $r_1 + 1$ parameters and is thus more efficient; see [Crossed-effects models](#) for details.

9. Fit a two-way full-factorial random-effects model:

```
. mecmd depvar [indepvars] ... || _all: R.factor1 || factor2: || factor1: ...
```

10. Fit a two-level mixed-effects model with a blocked-diagonal covariance structure between *revars1* and *revars2*:

```
. mecmd depvar [indepvars] ... || levelvar: revars1, noconstant ///
|| levelvar: revars2, noconstant ...
```

11. Fit a linear mixed-effects model where the correlation between the residual errors follows an autoregressive process of order 1:

```
. mixed depvar [indepvars] ... || levelvar:, residuals(ar 1, t(time)) ...
```

More residual error structures are available; see [ME] [mixed](#) for details.

12. Fit a two-level linear mixed-effects model accounting for sampling weights *expr1* at the first (residual) level and for sampling weights *expr2* at the level of *levelvar*:

```
. mixed depvar [indepvars] [pweight=expr1] ... || levelvar:, pweight(expr2) ...
```

Mixed-effects commands—with the exception of `mixed`, `meqrlogit`, and `meqrpoisson`—allow constraints on both fixed-effects and random-effects parameters. We provide several examples below of imposing constraints on variance components.

13. Fit a mixed-effects model with the variance of the random intercept on *levelvar* constrained to be 16:

```
. constraint 1 _b[var(_cons[levelvar]):_cons]=16
. mecmd depvar [indepvars] ... || levelvar:, constraints(1) ...
```

14. Fit a mixed-effects model with the variance of the random intercept on *levelvar* and the variance of the random slope on *revar* to be equal:

```
. constraint 1 _b[var(revar[levelvar]):_cons] = _b[var(_cons[levelvar]):_cons]
. mecmd depvar [indepvars] ... || levelvar: revar, constraints(1) ...
```

Note that the constraints above are equivalent to imposing an identity covariance structure for the random-effects equation:

```
. mecmd depvar [indepvars] ... || levelvar: revar, cov(identity) ...
```

15. Assuming four random slopes *revars*, fit a mixed-effects model with the variance components at the level of *levelvar* constrained to have a banded structure:

```
. mat p = (1,.,.,. \ 2,1,.,. \ 3,2,1,. \ 4,3,2,1)
. mecmd depvar [indepvars] ... || levelvar: revars, noconstant ///
covariance(pattern(p)) ...
```

16. Assuming four random slopes *revars*, fit a mixed-effects model with the variance components at the level of *levelvar* constrained to the specified numbers, and with all the covariances constrained to be 0:

```
. mat f = diag((1,2,3,4))
. mecmd depvar [indepvars] ... || levelvar: revars, noconstant ///
covariance(fixed(f)) ...
```

The variance components in models in items 15 and 16 can also be constrained by using the `constraints()` option, but using `covariance(pattern())` or `covariance(fixed())` is more convenient.

Mixed-effects models

Linear mixed-effects models

Linear mixed-effects (LME) models for continuous responses are a generalization of linear regression allowing for the inclusion of random deviations (effects) other than those associated with the overall error term. In matrix notation,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon} \quad (1)$$

where $\mathbf{y}$ is the $n \times 1$ vector of responses, $\mathbf{X}$ is an $n \times p$ design/covariate matrix for the fixed effects $\boldsymbol{\beta}$, and $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. The $n \times 1$ vector of errors $\boldsymbol{\epsilon}$ is assumed to be multivariate normal with mean 0 and variance matrix $\sigma_\epsilon^2 \mathbf{R}$.

The fixed portion of (1), $\mathbf{X}\boldsymbol{\beta}$, is analogous to the linear predictor from a standard OLS regression model with $\boldsymbol{\beta}$ being the regression coefficients to be estimated. For the random portion of (1), $\mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}$, we assume that $\mathbf{u}$ has variance–covariance matrix $\mathbf{G}$ and that $\mathbf{u}$ is orthogonal to $\boldsymbol{\epsilon}$ so that

$$\text{Var} \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\epsilon} \end{bmatrix} = \begin{bmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \sigma_\epsilon^2 \mathbf{R} \end{bmatrix}$$

The random effects $\mathbf{u}$ are not directly estimated (although they may be predicted) but instead are characterized by the elements of $\mathbf{G}$, known as variance components, that are estimated along with the error-covariance parameters that include the overall error variance σ_ϵ^2 and the parameters that are contained within $\mathbf{R}$.

The general forms of the design matrices $\mathbf{X}$ and $\mathbf{Z}$ allow estimation for a broad class of linear models: blocked designs, split-plot designs, growth curves, multilevel or hierarchical designs, etc. They also allow a flexible method of modeling within-cluster correlation. Subjects within the same cluster can be correlated as a result of a shared random intercept, or through a shared random slope on age (for example), or both. The general specification of $\mathbf{G}$ also provides additional flexibility: the random intercept and random slope could themselves be modeled as independent, or correlated, or independent with equal variances, and so forth. The general structure of $\mathbf{R}$ also allows for within-cluster errors to be heteroskedastic and correlated and allows flexibility in exactly how these characteristics can be modeled.

In clustered-data situations, it is convenient not to consider all n observations at once but instead to organize the mixed model as a series of M independent groups (or clusters)

$$\mathbf{y}_j = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j + \boldsymbol{\epsilon}_j \quad (2)$$

for $j = 1, \dots, M$, with cluster j consisting of n_j observations. The response $\mathbf{y}_j$ comprises the rows of $\mathbf{y}$ corresponding with the j th cluster, with $\mathbf{X}_j$ and $\boldsymbol{\epsilon}_j$ defined analogously. The random effects $\mathbf{u}_j$ can now be thought of as M realizations of a $q \times 1$ vector that is normally distributed with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The matrix $\mathbf{Z}_j$ is the $n_j \times q$ design matrix for the j th cluster random effects. Relating this to (1),

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{Z}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Z}_M \end{bmatrix}; \quad \mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_M \end{bmatrix}; \quad \mathbf{G} = \mathbf{I}_M \otimes \boldsymbol{\Sigma}; \quad \mathbf{R} = \mathbf{I}_M \otimes \boldsymbol{\Lambda}$$

where $\boldsymbol{\Lambda}$ denotes the variance matrix of the level-1 errors and $\otimes$ is the Kronecker product.

The mixed-model formulation (2) is from Laird and Ware (1982) and offers two key advantages. First, it makes specifications of random-effects terms easier. If the clusters are schools, you can simply specify a random effect at the school level, as opposed to thinking of what a school-level random effect would mean when all the data are considered as a whole (if it helps, think Kronecker products). Second, representing a mixed-model with (2) generalizes easily to more than one set of random effects. For example, if classes are nested within schools, then (2) can be generalized to allow random effects at both the school and the class-within-school levels.

By our convention on counting and ordering model levels, (2) is a two-level model, with extensions to three, four, or any number of levels. The observation y_{ij} is for individual i within cluster j , and the individuals compose the first level, whereas the clusters compose the second level of the model. In a hypothetical three-level model with classes nested within schools, the observations within classes (the students, presumably) would constitute the first level, the classes would constitute the second level, and the schools would constitute the third level. This differs from certain citations in the classical ANOVA literature and texts such as Pinheiro and Bates (2000) but is the standard in the vast literature on hierarchical models, for example, Skrondal and Rabe-Hesketh (2004).

In Stata, you can use `mixed` to fit linear mixed-effects models; see [ME] `mixed` for a detailed discussion and examples. Various predictions, statistics, and diagnostic measures are available after fitting an LME model with `mixed`. For the most part, calculation centers around obtaining estimates of random effects; see [ME] `mixed postestimation` for a detailed discussion and examples.

Generalized linear mixed-effects models

Generalized linear mixed-effects (GLME) models, also known as generalized linear mixed models (GLMMs), are extensions of generalized linear models allowing for the inclusion of random deviations (effects). In matrix notation,

$$g\{E(\mathbf{y}|\mathbf{X}, \mathbf{u})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim F \quad (3)$$

where $\mathbf{y}$ is the $n \times 1$ vector of responses from the distributional family F , $\mathbf{X}$ is an $n \times p$ design/covariate matrix for the fixed effects $\boldsymbol{\beta}$, and $\mathbf{Z}$ is an $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. The $\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}$ part is called the linear predictor and is often denoted as $\boldsymbol{\eta}$. $g(\cdot)$ is called the link function and is assumed to be invertible such that

$$E(\mathbf{y}|\mathbf{u}) = g^{-1}(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}) = H(\boldsymbol{\eta}) = \boldsymbol{\mu}$$

For notational convenience here and throughout this manual entry, we suppress the dependence of $\mathbf{y}$ on $\mathbf{X}$. Substituting various definitions for $g(\cdot)$ and F results in a wide array of models. For instance, if $g(\cdot)$ is the logit function and $\mathbf{y}$ is distributed as Bernoulli, we have

$$\text{logit}\{E(\mathbf{y}|\mathbf{u})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim \text{Bernoulli}$$

or mixed-effects logistic regression. If $g(\cdot)$ is the natural log function and $\mathbf{y}$ is distributed as Poisson, we have

$$\ln\{E(\mathbf{y}|\mathbf{u})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim \text{Poisson}$$

or mixed-effects Poisson regression.

For the random portion of (3), $\mathbf{Z}\mathbf{u}$, we assume that $\mathbf{u}$ has variance–covariance matrix $\mathbf{G}$ such that

$$\text{Var}(\mathbf{u}) = \mathbf{G}$$

The random effects $\mathbf{u}$ are not directly estimated (although they may be predicted) but instead are characterized by the elements of $\mathbf{G}$, known as variance components.

Analogously to (2), in clustered-data situations, we can write

$$E(\mathbf{y}_j|\mathbf{u}_j) = g^{-1}(\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) \quad \mathbf{y}_j \sim F$$

with all the elements defined as before. In terms of the whole dataset, we now have

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{Z}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Z}_M \end{bmatrix}; \quad \mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_M \end{bmatrix}; \quad \mathbf{G} = \mathbf{I}_M \otimes \boldsymbol{\Sigma}$$

In Stata, you can use `meglm` to fit mixed-effects models for nonlinear responses. Some combinations of families and links are so common that we implemented them as separate commands in terms of `meglm`.

Command	<code>meglm</code> equivalent
<code>melogit</code>	<code>family(bernoulli) link(logit)</code>
<code>meprobit</code>	<code>family(bernoulli) link(probit)</code>
<code>mecloglog</code>	<code>family(bernoulli) link(cloglog)</code>
<code>meologit</code>	<code>family(ordinal) link(logit)</code>
<code>meoprobit</code>	<code>family(ordinal) link(probit)</code>
<code>mepoisson</code>	<code>family(poisson) link(log)</code>
<code>menbreg</code>	<code>family(nbinomial) link(log)</code>

When no family–link combination is specified, `meglm` defaults to a Gaussian family with an identity link. Thus `meglm` can be used to fit linear mixed-effects models; however, for those models we recommend using the more specialized `mixed`, which, in addition to `meglm` capabilities, allows for modeling of the structure of the within-cluster errors; see [ME] [mixed](#) for details.

Various predictions, statistics, and diagnostic measures are available after fitting a GLME model with `meglm` and other `me` commands. For the most part, calculation centers around obtaining estimates of random effects; see [ME] [meglm postestimation](#) for a detailed discussion and examples.

Survival mixed-effects models

Parametric survival mixed-effects models use a trivariate response variable (t_0, t, d) , where each response corresponds to a period under observation $(t_0, t]$ and results in either failure ($d = 1$) or right-censoring ($d = 0$) at time t . See [ST] [streg](#) for background information on parametric survival models. Two often-used models for adjusting survivor functions for the effects of covariates are the accelerated failure-time (AFT) model and the multiplicative or proportional hazards (PH) model.

In the AFT parameterization, the natural logarithm of the survival time, $\log t$, is expressed as a linear function of the covariates. When we incorporate random effects, this yields the model

$$\log(\mathbf{t}_j) = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j + \mathbf{v}_j$$

where $\log(\cdot)$ is an elementwise function, and $\mathbf{v}_j$ is a vector of observation-level errors. The distributional form of the error term determines the regression model.

In the PH model, the covariates have a multiplicative effect on the hazard function

$$h(\mathbf{t}_j) = h_0(\mathbf{t}_j) \exp(\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j)$$

where all the functions are elementwise, and $h_0(\cdot)$ is a baseline hazard function. The functional form of $h_0(\cdot)$ determines the regression model.

In Stata, you can use `mestreg` to fit multilevel mixed-effects parametric survival models for the following distributions and parameterizations.

Distribution	Parameterization
exponential	PH, AFT
loglogistic	AFT
weibull	PH, AFT
lognormal	AFT
gamma	AFT

`mestreg` is suitable only for data that have been set using the `stset` command. By using `stset` on your data, you define the variables `_t0`, `_t`, and `_d`, which serve as the trivariate response. See [ME] `mestreg` for more details about the command. Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects survival model with `mestreg`; see [ME] `mestreg postestimation` for a detailed discussion and examples.

Nonlinear mixed-effects models

NLME models are models containing both fixed effects and random effects where some of, or all, the fixed and random effects enter the model nonlinearly. They can be viewed as a generalization of LME models, in which the conditional mean of the outcome given the random effects is a nonlinear function of the coefficients and random effects. Alternatively, they can be considered as an extension of nonlinear regression models for independent data (see [R] `nl`), in which coefficients may incorporate random effects, allowing them to vary across different levels of hierarchy and thus inducing correlation within observations at the same level.

Using the notation from *Linear mixed-effects models* for LME models for clustered data, we can write an NLME model as

$$\mathbf{y}_j = \boldsymbol{\mu}(\mathbf{A}_j, \boldsymbol{\beta}, \mathbf{u}_j) + \boldsymbol{\epsilon}_j$$

where $\boldsymbol{\mu}(\cdot)$ is a real-valued vector function and $\mathbf{A}_j$ is an $n_j \times l$ matrix of covariates for the j th cluster, which includes both within-subject and between-subject covariates. Do not be surprised to see the $\mathbf{A}_j$ matrix here instead of the more familiar fixed-effects and random-effects design matrices $\mathbf{X}_j$ and $\mathbf{Z}_j$ from previous sections. Because both covariates and parameters can enter the model nonlinearly in NLME, we cannot express the regression function as a function containing the linear term $\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j$ as we can for LME and GLME models. The distributional assumptions on $\mathbf{u}_j$'s and $\boldsymbol{\epsilon}_j$'s are the same as for the LME models.

Parameters of NLME models often have scientifically meaningful interpretations, and research questions are formed based on them. To allow parameters to reflect phenomena of interest, NLME models are often formulated by using a multistage formulation; see *Alternative mixed-effects model specification* below for examples.

We can formulate our previous NLME model as a two-stage hierarchical model:

Stage 1: Individual-level model $y_{ij} = m(\mathbf{x}_{ij}^w, \phi_j) + \epsilon_{ij}, \quad i = 1, \dots, n_j$

Stage 2: Group-level model $\phi_j = \mathbf{d}(\mathbf{x}_j^b, \beta, \mathbf{u}_j), \quad j = 1, \dots, M$

In stage 1, we model the response by using a function $m(\cdot)$, which describes within-subject behavior. This function depends on subject-specific parameters ϕ_j 's, which have a natural physical interpretation, and a vector of within-subject covariates $\mathbf{x}_{ij}^w$. In stage 2, we use a known vector-valued function $\mathbf{d}(\cdot)$ to model between-subject behavior, that is, to model ϕ_j 's and to explain how they vary across subjects. The $\mathbf{d}(\cdot)$ function incorporates random effects and, optionally, a vector of between-subject covariates $\mathbf{x}_j^b$. The general idea is to specify a common functional form for each subject in stage 1 and then allow some parameters to vary randomly across subjects in stage 2.

You can use the `men1` command to fit NLME models to continuous outcomes; see [ME] [men1](#). `men1` supports both the single-equation and multistage model formulations. It supports different covariance structures for random effects and can model heteroskedasticity and correlations within lowest-level groups. Various predictions, statistics, and diagnostic measures are available after fitting an NLME model; see [ME] [men1 postestimation](#).

For an introductory example, see [Nonlinear models](#).

Alternative mixed-effects model specification

In this section, we present a hierarchical or multistage formulation of mixed-effects models where each level is described by its own set of equations. This formulation is common for NLME models; see [Nonlinear mixed-effects models](#).

Consider a random-intercept model that we write here in general terms:

$$y_{ij} = \beta_0 + \beta_1 x_{ij} + u_j + \epsilon_{ij} \quad (4)$$

This single-equation specification contains both level-1 and level-2 effects. In the hierarchical form, we specify a separate equation for each level.

$$\begin{aligned} y_{ij} &= \gamma_{0j} + \beta_1 x_{ij} + \epsilon_{ij} \\ \gamma_{0j} &= \beta_{00} + u_{0j} \end{aligned} \quad (5)$$

The equation for the intercept γ_{0j} consists of the overall mean intercept β_{00} and a cluster-specific random intercept u_{0j} . To fit this model by using, for example, `mixed`, we must translate the multiple-equation notation into a single-equation form. We substitute the second equation into the first one and rearrange terms.

$$\begin{aligned} y_{ij} &= \beta_{00} + u_{0j} + \beta_1 x_{ij} + \epsilon_{ij} \\ &= \beta_{00} + \beta_1 x_{ij} + u_{0j} + \epsilon_{ij} \end{aligned} \quad (6)$$

Note that model (6) is the same as model (4) with $\beta_{00} \equiv \beta_0$ and $u_{0j} \equiv u_j$. Thus the syntax for our generic random-intercept model is

```
. mixed y x || id:
```

where `id` is the variable designating the clusters.

We can extend model (5) to include a random slope. We do so by specifying an additional equation for the slope on x_{ij} .

$$\begin{aligned} y_{ij} &= \gamma_{0j} + \gamma_{1j}x_{ij} + \epsilon_{ij} \\ \gamma_{0j} &= \beta_{00} + u_{0j} \\ \gamma_{1j} &= \beta_{10} + u_{1j} \end{aligned} \tag{7}$$

The additional equation for the slope γ_{1j} consists of the overall mean slope β_{10} and a cluster-specific random slope u_{1j} . We substitute the last two equations into the first one to obtain a reduced-form model.

$$\begin{aligned} y_{ij} &= (\beta_{00} + u_{0j}) + (\beta_{10} + u_{1j})x_{ij} + \epsilon_{ij} \\ &= \beta_{00} + \beta_{10}x_{ij} + u_{0j} + u_{1j}x_{ij} + \epsilon_{ij} \end{aligned}$$

The syntax for this model becomes

```
. mixed y x || id: x, covariance(unstructured)
```

where we specified an unstructured covariance structure for the level-2 u terms.

Here we further extend the random-slope random-intercept model (7) by adding a level-2 covariate z_j into the level-2 equations.

$$\begin{aligned} y_{ij} &= \gamma_{0j} + \gamma_{1j}x_{ij} + \epsilon_{ij} \\ \gamma_{0j} &= \beta_{00} + \beta_{01}z_j + u_{0j} \\ \gamma_{1j} &= \beta_{10} + \beta_{11}z_j + u_{1j} \end{aligned}$$

We substitute as before to obtain a single-equation form:

$$\begin{aligned} y_{ij} &= (\beta_{00} + \beta_{01}z_j + u_{0j}) + (\beta_{10} + \beta_{11}z_j + u_{1j})x_{ij} + \epsilon_{ij} \\ &= \beta_{00} + \beta_{01}z_j + \beta_{10}x_{ij} + \beta_{11}z_jx_{ij} + u_{0j} + u_{1j}x_{ij} + \epsilon_{ij} \end{aligned}$$

Now the fixed-effects portion of the equation contains a constant and variables x , z , and their interaction. Assuming both x and z are continuous variables, we can use the following Stata syntax to fit this model:

```
. mixed y x z c.x#c.z || id: x, covariance(unstructured)
```

Although the `menl` command is not as suitable for fitting LME models as `mixed`, it can accommodate a multistage formulation. For example, (5) can be fit in `menl` as

```
. menl y = {gamma0:}+{b1}*x, define(gamma0: {b00}+{U0[id]})
```

and (7) as

```
. menl y = {gamma0:}+{gamma1:}*x, define(gamma0: {b00}+{U0[id]}) ///
      define(gamma1: {b10}+{U1[id]})
```

In the above `menl`'s specifications, `gamma0` and `gamma1` can be specified more efficiently by using linear combinations; see [ME] [menl](#) for details.

We refer you to [Raudenbush and Bryk \(2002\)](#) and [Rabe-Hesketh and Skrondal \(2012\)](#) for a more thorough discussion and further examples of multistage mixed-model formulations, including three-level models.

Likelihood calculation

The key to fitting mixed models lies in estimating the variance components, and for that there exist many methods. Most of the early literature in LME models dealt with estimating variance components in ANOVA models. For simple models with balanced data, estimating variance components amounts to solving a system of equations obtained by setting expected mean-squares expressions equal to their observed counterparts. Much of the work in extending the ANOVA method to unbalanced data for general ANOVA designs is attributed to [Henderson \(1953\)](#).

The ANOVA method, however, has its shortcomings. Among these is a lack of uniqueness in that alternative, unbiased estimates of variance components could be derived using other quadratic forms of the data in place of observed mean squares ([Searle, Casella, and McCulloch 1992](#), 38–39). As a result, ANOVA methods gave way to more modern methods, such as minimum norm quadratic unbiased estimation (MINQUE) and minimum variance quadratic unbiased estimation (MIVQUE); see [Rao \(1973\)](#) for MINQUE and [LaMotte \(1973\)](#) for MIVQUE. Both methods involve finding optimal quadratic forms of the data that are unbiased for the variance components.

Stata uses maximum likelihood (ML) to fit LME and GLME models. The ML estimates are based on the usual application of likelihood theory, given the distributional assumptions of the model. In addition, for linear mixed-effects models, `mixed` offers the method of restricted maximum likelihood (REML). The basic idea behind REML ([Thompson 1962](#)) is that you can form a set of linear contrasts of the response that do not depend on the fixed effects β but instead depend only on the variance components to be estimated. You then apply ML methods by using the distribution of the linear contrasts to form the likelihood; see the [Methods and formulas](#) section of [\[ME\] mixed](#) for a detailed discussion of ML and REML methods in the context of linear mixed-effects models.

Log-likelihood calculations for fitting any mixed-effects model require integrating out the random effects. For LME models, this integral has a closed-form solution; for GLME and NLME models, it does not. In dealing with this difficulty, early estimation methods avoided the integration altogether. Two such popular methods are the closely related penalized quasiliquelihood (PQL) and marginal quasiliquelihood (MQL) ([Breslow and Clayton 1993](#)). Both PQL and MQL use a combination of iterative reweighted least squares (see [\[R\] glm](#)) and standard estimation techniques for fitting LME models. Efficient computational methods for fitting LME models have existed for some time ([Bates and Pinheiro 1998](#); [Littell et al. 2006](#)), and PQL and MQL inherit this computational efficiency. However, both of these methods suffer from two key disadvantages. First, they have been shown to be biased, and this bias can be severe when clusters are small or intracluster correlation is high ([Rodríguez and Goldman 1995](#); [Lin and Breslow 1996](#)). Second, because they are “quasiliquelihood” methods and not true likelihood methods, their use prohibits comparing nested models via likelihood-ratio (LR) tests, blocking the main avenue of inference involving variance components.

The advent of modern computers has brought with it the development of more computationally intensive methods, such as bias-corrected PQL ([Lin and Breslow 1996](#)), Bayesian Markov-Chain Monte Carlo, and simulated maximum likelihood, just to name a few; see [Ng et al. \(2006\)](#) for a discussion of these alternate strategies (and more) for mixed-effects models for binary outcomes.

One widely used modern method is to directly estimate the integral required to calculate the log likelihood by Gauss–Hermite quadrature or some variation thereof. Because the log likelihood itself is estimated, this method has the advantage of permitting LR tests for comparing nested models. Also, if done correctly, quadrature approximations can be quite accurate, thus minimizing bias. Stata commands for fitting GLME models such as `meqglm` support three types of Gauss–Hermite quadratures: mean–variance adaptive Gauss–Hermite quadrature (MVAGH), mode-curvature adaptive Gauss–Hermite quadrature (MCAGH), and nonadaptive Gauss–Hermite quadrature (GHQ); see [Methods and formulas](#) of [\[ME\] meqglm](#) for a detailed discussion of these quadrature methods. A fourth method, the Laplacian approximation, that does not involve numerical integration is also offered; see [Computation time](#)

and the Laplacian approximation below and *Methods and formulas* of [ME] **meglm** for a detailed discussion of the Laplacian approximation method.

In the context of NLME models, the use of an adaptive quadrature to fit these models can be often computationally infeasible. A popular alternative method used to fit NLME models is the linearization method of Lindstrom and Bates (1990), also known as the conditional first-order linearization method. It is based on a first-order Taylor-series approximation of the mean function and essentially linearizes the mean function with respect to fixed and random effects. The linearization method is computationally efficient because it avoids the intractable integration, but the approximation cannot be made arbitrarily accurate. Despite its potential limiting accuracy, the linearization method has proven the most popular in practice (Fitzmaurice et al. 2009, sec 5.4.8). The linearization method of Lindstrom and Bates (1990), with extensions from Pinheiro and Bates (1995), is the method of estimation in **menl**.

Computation time and the Laplacian approximation

Like many programs that fit generalized linear mixed models, **me** commands can be computationally intensive. This is particularly true for large datasets with many lowest-level clusters, models with many random coefficients, models with many estimable parameters (both fixed effects and variance components), or any combination thereof.

Computation time will also depend on hardware and other external factors but in general is (roughly) a function of $p^2\{M + M(N_Q)^{q_t}\}$, where p is the number of estimable parameters, M is the number of lowest-level (smallest) clusters, N_Q is the number of quadrature points, and q_t is the total dimension of the random effects, that is, the total number of random intercepts and coefficients at all levels.

For a given model and a given dataset, the only prevailing factor influencing computation time is $(N_Q)^{q_t}$. However, because this is a power function, this factor can get prohibitively large. For example, using five quadrature points for a model with one random intercept and three random coefficients, we get $(N_Q)^{q_t} = 5^4 = 625$. Even a modest increase to seven quadrature points would increase this factor by almost fourfold ($7^4 = 2,401$), which, depending on M and p , could drastically slow down estimation. When fitting mixed-effects models, you should always assess whether the approximation is adequate by refitting the model with a larger number of quadrature points. If the results are essentially the same, the lower number of quadrature points can be used.

However, we do not deny a tradeoff between speed and accuracy, and in that spirit we give you the option to choose a (possibly) less accurate solution in the interest of getting quicker results. Toward this end is the limiting case of $N_Q = 1$, otherwise known as the Laplacian approximation; see *Methods and formulas* of [ME] **meglm**. The computational benefit is evident—1 raised to any power equals 1—and the Laplacian approximation has been shown to perform well in certain situations (Liu and Pierce 1994; Tierney and Kadane 1986). When using Laplacian approximation, keep the following in mind:

1. Fixed-effects parameters and their standard errors are well approximated by the Laplacian method. Therefore, if your interest lies primarily here, then the Laplacian approximation may be a viable alternative.
2. Estimates of variance components exhibit bias, particularly the variances.
3. The model log likelihood and comparison LR test are in fair agreement with statistics obtained via quadrature methods.

Although this is by no means the rule, we find the above observations to be fairly typical based on our own experience. Pinheiro and Chao (2006) also make observations similar to points 1 and 2 on the basis of their simulation studies: bias due to Laplace (when present) tends to exhibit itself

more in the estimated variance components than in the estimates of the fixed effects as well as at the lower levels in higher-level models.

Item 3 is of particular interest, because it demonstrates that the Laplacian approximation can produce a decent estimate of the model log likelihood. Consequently, you can use the Laplacian approximation during the model building phase of your analysis, during which you are comparing competing models by using LR tests. Once you settle on a parsimonious model that fits well, you can then increase the number of quadrature points and obtain more accurate parameter estimates for further study.

Of course, sometimes the Laplacian approximation will perform either better or worse than observed here. This behavior depends primarily on cluster size and intracluster correlation, but the relative influence of these factors is unclear. The idea behind the Laplacian approximation is to approximate the posterior density of the random effects given the response with a normal distribution; see [Methods and formulas](#) of [ME] **meglm**. Asymptotic theory dictates that this approximation improves with larger clusters. Of course, the key question, as always, is “How large is large enough?” Also, there are data situations where the Laplacian approximation performs well even with small clusters. Therefore, it is difficult to make a definitive call as to when you can expect the Laplacian approximation to yield accurate results across all aspects of the model.

In conclusion, consider our above advice as a rule of thumb based on empirical evidence.

Diagnosing convergence problems

Given the flexibility of mixed-effects models, you will find that some models fail to converge when used with your data. The default gradient-based method used by mixed-effects commands, except `men1`, is the Newton–Raphson algorithm, requiring the calculation of a gradient vector and Hessian (second-derivative) matrix; see [R] **ml**.

A failure to converge can take any one of three forms:

1. repeated nonconcave or backed-up iterations without convergence;
2. a Hessian (second-derivative) calculation that has become asymmetric, unstable, or has missing values; or
3. the message “standard-error calculation has failed” when computing standard errors.

All three situations essentially amount to the same thing: the Hessian calculation has become unstable, most likely because of a ridge in the likelihood function, a subsurface of the likelihood in which all points give the same value of the likelihood and for which there is no unique solution.

Such behavior is usually the result of one of the following two situations:

- A. A model that is not identified given the data, for example, fitting the three-level nested random intercept model

$$y_{jk} = \mathbf{x}_{jk}\boldsymbol{\beta} + u_k^{(3)} + u_{jk}^{(2)} + \epsilon_{jk}$$

without any replicated measurements at the (j, k) level, that is, with only one i per (j, k) combination. This model is unidentified for such data because the random intercepts $u_{jk}^{(2)}$ are confounded with the overall errors ϵ_{jk} .

- B. A model that contains a variance component whose estimate is really close to 0. When this occurs, a ridge is formed by an interval of values near 0, which produce the same likelihood and look equally good to the optimizer.

For LME models, one useful way to diagnose problems of nonconvergence is to rely on the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1977), normally used by `mixed` only as a means of refining starting values; see *Diagnosing convergence problems* of [ME] `mixed` for details.

If your data and model are nearly unidentified, as opposed to fully unidentified, you may be able to obtain convergence with standard errors by changing some of the settings of the gradient-based optimization. Adding the `difficult` option can be particularly helpful if you are seeing many “nonconcave” messages; you may also consider changing the `technique()` or using the `nonrtolerance` option; see [R] `maximize`.

Regardless of how the convergence problem revealed itself, you may try to obtain better starting values; see *Obtaining better starting values* in [ME] `meglm` for details.

Achieving convergence and diagnosing convergence problems can be even more challenging with NLME models. As with other mixed-effects models, complicated variance–covariance structures for random effects and errors can often lead to overparameterized models that fail to converge. In addition, highly nonlinear mean specifications can lead to multiple solutions and thus to potential convergence to a local maximum. `menl` uses the linearization estimation method that alternates between the penalized least-squares estimation of the fixed-effects parameters and the Newton–Raphson estimation of the random-effects parameters of the approximating LME model, which was the result of the linearization of the original NLME model. This alternating method does not provide a joint Hessian matrix for all parameters, so there is no check for the tolerance of the scaled gradient, and thus the convergence cannot be established in its strict sense. The convergence is declared based on the stopping rules described in *Methods and formulas* of [ME] `menl`. Exploring different initial values to investigate convergence is particularly important with NLME models; see *Obtaining initial values* in [ME] `menl`.

Distribution theory for likelihood-ratio test

When determining the asymptotic distribution of an LR test comparing two nested mixed-effects models, issues concerning boundary problems imposed by estimating strictly positive quantities (that is, variances) can complicate the situation. For example, when performing LR tests involving linear mixed-effects models (whether comparing with linear regression within `mixed` or comparing two separate linear mixed-effects models with `lrttest`), you may thus sometimes see a test labeled as `chibar` rather than the usual `chi2`, or you may see a `chi2` test with a note attached stating that the test is conservative or possibly conservative depending on the hypothesis being tested.

At the heart of the issue is the number of variances being restricted to 0 in the reduced model. If there are none, the usual asymptotic theory holds, and the distribution of the test statistic is χ^2 with degrees of freedom equal to the difference in the number of estimated parameters between both models.

When there is only one variance being set to 0 in the reduced model, the asymptotic distribution of the LR test statistic is a 50:50 mixture of a χ_p^2 and a χ_{p+1}^2 distribution, where p is the number of other restricted parameters in the reduced model that are unaffected by boundary conditions. Stata labels such test statistics as `chibar` and adjusts the significance levels accordingly. See Self and Liang (1987) for the appropriate theory or Gutierrez, Carter, and Drukker (2001) for a Stata-specific discussion.

When more than one variance parameter is being set to 0 in the reduced model, however, the situation becomes more complicated. For example, consider a comparison test versus linear regression for a mixed model with two random coefficients and unstructured covariance matrix

$$\Sigma = \begin{bmatrix} \sigma_0^2 & \sigma_{01} \\ \sigma_{01} & \sigma_1^2 \end{bmatrix}$$

Because the random component of the mixed model comprises three parameters $(\sigma_0^2, \sigma_{01}, \sigma_1^2)$, on the surface it would seem that the LR comparison test would be distributed as χ_3^2 . However, two complications need to be considered. First, the variances σ_0^2 and σ_1^2 are restricted to be positive, and second, constraints such as $\sigma_1^2 = 0$ implicitly restrict the covariance σ_{01} to be 0 as well. From a technical standpoint, it is unclear how many parameters must be restricted to reduce the model to linear regression.

Because of these complications, appropriate and sufficiently general distribution theory for the more-than-one-variance case has yet to be developed. Theory (for example, [Stram and Lee \[1994\]](#)) and empirical studies (for example, [McLachlan and Basford \[1988\]](#)) have demonstrated that, whatever the distribution of the LR test statistic, its tail probabilities are bounded above by those of the χ^2 distribution with degrees of freedom equal to the full number of restricted parameters (three in the above example).

The `mixed` and `me` commands use this reference distribution, the χ^2 with full degrees of freedom, to produce a conservative test and place a note in the output labeling the test as such. Because the displayed significance level is an upper bound, rejection of the null hypothesis based on the reported level would imply rejection on the basis of the actual level.

Examples

Two-level models

► Example 1: Growth-curve model

Consider a longitudinal dataset, used by both [Ruppert, Wand, and Carroll \(2003\)](#) and [Diggle et al. \(2002\)](#), consisting of `weight` measurements of 48 pigs on 9 successive `weeks`. Pigs are identified by the variable `id`. Each pig experiences a linear trend in growth, but overall weight measurements vary from pig to pig. Because we are not really interested in these particular 48 pigs per se, we instead treat them as a random sample from a larger population and model the between-pig variability as a random effect, or in the terminology of (2), as a random-intercept term at the pig level. We thus wish to fit the model

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_j + \epsilon_{ij}$$

for $i = 1, \dots, 9$ weeks and $j = 1, \dots, 48$ pigs. The fixed portion of the model, $\beta_0 + \beta_1 \text{week}_{ij}$, simply states that we want one overall regression line representing the population average. The random effect u_j serves to shift this regression line up or down according to each pig. Because the random effects occur at the pig level (`id`), we fit the model by typing

```

. use http://www.stata-press.com/data/r15/pig
(Longitudinal analysis of pig weights)
. mixed weight week || id:
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -1014.9268
Iteration 1:   log likelihood = -1014.9268
Computing standard errors:
Mixed-effects ML regression              Number of obs   =       432
Group variable: id                     Number of groups =        48
Obs per group:
      min =          9
      avg =         9.0
      max =          9
Wald chi2(1) =    25337.49
Prob > chi2   =      0.0000
Log likelihood = -1014.9268

```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0390124	159.18	0.000	6.133433	6.286359
_cons	19.35561	.5974059	32.40	0.000	18.18472	20.52651

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity					
	var(_cons)	14.81751	3.124226	9.801716	22.40002
	var(Residual)	4.383264	.3163348	3.805112	5.04926

LR test vs. linear model: chibar2(01) = 472.65 Prob >= chibar2 = 0.0000

We explain the output in detail in [example 1](#) of [\[ME\] mixed](#). Here we only highlight the most important points.

1. The first estimation table reports the fixed effects. We estimate $\beta_0 = 19.36$ and $\beta_1 = 6.21$.
2. The second estimation table shows the estimated variance components. The first section of the table is labeled `id: Identity`, meaning that these are random effects at the `id` (pig) level and that their variance–covariance matrix is a multiple of the identity matrix; that is, $\Sigma = \sigma_u^2 \mathbf{I}$. The estimate of $\hat{\sigma}_u^2$ is 14.82 with standard error 3.12.
3. The row labeled `var(Residual)` displays the estimated standard deviation of the overall error term; that is, $\hat{\sigma}_\epsilon^2 = 4.38$. This is the variance of the level-one errors or the variance of the residuals.
4. An LR test comparing the model with one-level ordinary linear regression is provided and is highly significant for these data.

We can predict the random intercept u_j and list the predicted random intercept for the first 10 pigs by typing

```
. predict r_int, reffects
. egen byte tag = tag(id)
. list id r_int if id<=10 & tag
```

	id	r_int
1.	1	-1.683105
10.	2	.8987018
19.	3	-1.952043
28.	4	-1.79068
37.	5	-3.189159
46.	6	-3.780823
55.	7	-2.382344
64.	8	-1.952043
73.	9	-6.739143
82.	10	1.16764

In [example 3](#) of [\[ME\] mixed](#), we show how to fit a random-slope model for these data, and in [example 1](#) of [\[ME\] mixed postestimation](#), we show how to plot the estimated regression lines for each of the pigs.



► Example 2: Split-plot design

Here we replicate the example of a split-plot design from [Kuehl \(2000, 477\)](#). The researchers investigate the effects of nitrogen in four different chemical forms and the effects of thatch accumulation on the quality of golf turf. The experimental plots were arranged in a randomized complete block design with two replications. After two years of nitrogen treatment, the second treatment factor, years of thatch accumulation, was added to the experiment. Each of the eight experimental plots was split into three subplots. Within each plot, the subplots were randomly assigned to accumulate thatch for a period of 2, 5, and 8 years.

```
. use http://www.stata-press.com/data/r15/clippings, clear
(Turfgrass experiment)

. describe
Contains data from http://www.stata-press.com/data/r15/clippings.dta
  obs:                24                Turfgrass experiment
  vars:                4                21 Feb 2016 14:57
  size:               168
```

variable name	storage type	display format	value label	variable label
chlorophyll	float	%9.0g		Chlorophyll content (mg/g) of grass clippings
thatch	byte	%9.0g		Years of thatch accumulation
block	byte	%9.0g		Replication
nitrogen	byte	%17.0g	nitrolab	Nitrogen fertilizer

Sorted by:

Nitrogen treatment is stored in the variable `nitrogen`, and the chemicals used are urea, ammonium sulphate, isobutylidene diurea (IBDU), and sulphur-coated urea (urea SC). The length of thatch accumulation is stored in the variable `thatch`. The response is the chlorophyll content of grass clippings, recorded in mg/g and stored in the variable `chlorophyll`. The `block` variable identifies the replication group.

There are two sources of variation in this example corresponding to the whole-plot errors and the subplot errors. The subplot errors are the residual errors. The whole-plot errors represents variation in the chlorophyll content across nitrogen treatments and replications. We create the variable `wpunit` to represent the whole-plot units that correspond to the levels of the nitrogen treatment and block interaction.

```
. egen wpunit = group(nitrogen block)

. mixed chlorophyll ibn.nitrogen##ibn.thatch ibn.block, noomitted noconstant ||
> wpunit:, reml
note: 8.thatch omitted because of collinearity
note: 1.nitrogen#8.thatch omitted because of collinearity
note: 2.nitrogen#8.thatch omitted because of collinearity
note: 3.nitrogen#8.thatch omitted because of collinearity
note: 4.nitrogen#2.thatch omitted because of collinearity
note: 4.nitrogen#5.thatch omitted because of collinearity
note: 4.nitrogen#8.thatch omitted because of collinearity
note: 2.block omitted because of collinearity

Performing EM optimization:

Performing gradient-based optimization:

Iteration 0:   log restricted-likelihood = -13.212401
Iteration 1:   log restricted-likelihood = -13.203149
Iteration 2:   log restricted-likelihood = -13.203125
Iteration 3:   log restricted-likelihood = -13.203125
```

Computing standard errors:
Mixed-effects REML regression
Group variable: wpunit

Number of obs = 24
Number of groups = 8
Obs per group:
 min = 3
 avg = 3.0
 max = 3
Wald chi2(13) = 2438.36
Prob > chi2 = 0.0000

Log restricted-likelihood = -13.203125

chlorophyll1	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
nitrogen						
urea	5.245833	.3986014	13.16	0.000	4.464589	6.027078
ammonium s..	5.945833	.3986014	14.92	0.000	5.164589	6.727078
IBDU	7.945834	.3986014	19.93	0.000	7.164589	8.727078
Urea (SC)	8.595833	.3986014	21.56	0.000	7.814589	9.377078
thatch						
2	-1.1	.4632314	-2.37	0.018	-2.007917	-.1920828
5	.1500006	.4632314	0.32	0.746	-.7579163	1.057917
nitrogen#						
thatch						
urea#2	-.1500005	.6551081	-0.23	0.819	-1.433989	1.133988
urea#5	.0999994	.6551081	0.15	0.879	-1.183989	1.383988
ammonium s.. #						
2	.8999996	.6551081	1.37	0.169	-.3839887	2.183988
ammonium s.. #						
5	-.1000006	.6551081	-0.15	0.879	-1.383989	1.183988
IBDU#2	-.2000005	.6551081	-0.31	0.760	-1.483989	1.083988
IBDU#5	-1.950001	.6551081	-2.98	0.003	-3.233989	-.6660124
block						
1	-.2916666	.2643563	-1.10	0.270	-.8097955	.2264622

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
wpunit: Identity				
var(_cons)	.0682407	.1195933	.0021994	2.117344
var(Residual)	.2145833	.1072917	.080537	.5717376

LR test vs. linear model: chibar2(01) = 0.53

Prob >= chibar2 = 0.2324

We can calculate the cell means for source of nitrogen and years of thatch accumulation by using `margins`.

```
. margins thatch#nitrogen
```


```
Predictive margins                                Number of obs      =          24
```

```
Expression   : Linear prediction, fixed portion, predict()
```

	Delta-method		z	P> z	[95% Conf. Interval]	
	Margin	Std. Err.				
thatch# nitrogen 2#urea 2 #	3.85	.3760479	10.24	0.000	3.11296	4.58704
ammonium s.. 2#IBDU	5.6	.3760479	14.89	0.000	4.86296	6.33704
2#Urea (SC)	6.5	.3760479	17.29	0.000	5.76296	7.23704
5#urea 5 #	7.35	.3760479	19.55	0.000	6.61296	8.087041
5#urea	5.35	.3760479	14.23	0.000	4.61296	6.087041
ammonium s.. 5#IBDU	5.85	.3760479	15.56	0.000	5.11296	6.58704
5#IBDU	6	.3760479	15.96	0.000	5.26296	6.73704
5#Urea (SC)	8.6	.3760479	22.87	0.000	7.86296	9.337041
8#urea 8 #	5.1	.3760479	13.56	0.000	4.36296	5.837041
ammonium s.. 8#IBDU	5.8	.3760479	15.42	0.000	5.06296	6.53704
8#IBDU	7.8	.3760479	20.74	0.000	7.06296	8.537041
8#Urea (SC)	8.45	.3760479	22.47	0.000	7.712959	9.18704

It is easier to see the effect of the treatments if we plot the impact of the four nitrogen and the three thatch treatments. We can use `marginsplot` to plot the means of chlorophyll content versus years of thatch accumulation by nitrogen source.


```
. marginsplot, ytitle(Chlorophyll (mg/g)) title("")
> subtitle("Mean chlorophyll content of grass clippings versus"
> "nitrogen source for years of thatch accumulation") xsize(3) ysize(3.2)
> legend(cols(1) position(5) ring(0) region(lwidth(none)))
> ylabel(0(2)10, angle(0))
Variables that uniquely identify margins: thatch nitrogen
```



We can see an increase in the mean chlorophyll content over the years of thatch accumulation for all but one nitrogen source.

The marginal means can be obtained by using `margins` on one variable at a time.

```
. margins thatch
Predictive margins                                Number of obs      =          24
Expression   : Linear prediction, fixed portion, predict()
```

	Delta-method					[95% Conf. Interval]	
	Margin	Std. Err.	z	P> z			
thatch							
2	5.825	.188024	30.98	0.000	5.45648	6.19352	
5	6.45	.188024	34.30	0.000	6.08148	6.81852	
8	6.7875	.188024	36.10	0.000	6.41898	7.15602	

```
. margins nitrogen
Predictive margins                                Number of obs      =           24
Expression   : Linear prediction, fixed portion, predict()
```

	Delta-method		z	P> z	[95% Conf. Interval]	
	Margin	Std. Err.				
nitrogen						
urea	4.766667	.2643563	18.03	0.000	4.248538	5.284796
ammonium s..	5.75	.2643563	21.75	0.000	5.231871	6.268129
IBDU	6.766667	.2643563	25.60	0.000	6.248538	7.284796
Urea (SC)	8.133333	.2643563	30.77	0.000	7.615205	8.651462

[Marchenko \(2006\)](#) shows more examples of fitting other experimental designs using linear mixed-effects models.

◀

► Example 3: Binomial counts

We use the data taken from [Agresti \(2013, 219\)](#) on graduate school applications to the 23 departments within the College of Liberal Arts and Sciences at the University of Florida during the 1997–1998 academic year. The dataset contains the department ID (`department`), the number of applications (`nappplied`), and the number of students admitted (`nadmitted`) cross-classified by gender (`female`).

```
. use http://www.stata-press.com/data/r15/admissions, clear
(Graduate school admissions data)

. describe

Contains data from http://www.stata-press.com/data/r15/admissions.dta
   obs:                46                Graduate school admissions data
  vars:                  4                25 Feb 2016 09:28
 size:                 460                (_dta has notes)
```

variable name	storage type	display format	value label	variable label
department	long	%8.0g	dept	department id
nadmitted	byte	%8.0g		number of admissions
nappplied	float	%9.0g		number of applications
female	byte	%8.0g		=1 if female, =0 if male

Sorted by:

We wish to investigate whether admission decisions are independent of gender. Given department and gender, the probability of admission follows a binomial model, that is, $\Pr(Y_{ij} = y_{ij}) = \text{Binomial}(n_{ij}, \pi_{ij})$, where $i = \{0, 1\}$ and $j = 1, \dots, 23$. We fit a mixed-effects binomial logistic model with a random intercept at the department level.

```
. melogit nadmitted female || department:, binomial(napplied) or
Fitting fixed-effects model:
Iteration 0:   log likelihood = -302.47786
Iteration 1:   log likelihood = -300.00004
Iteration 2:   log likelihood = -299.99934
Iteration 3:   log likelihood = -299.99934
Refining starting values:
Grid node 0:   log likelihood = -145.08843
Fitting full model:
Iteration 0:   log likelihood = -145.08843
Iteration 1:   log likelihood = -140.8514
Iteration 2:   log likelihood = -140.61709
Iteration 3:   log likelihood = -140.61628
Iteration 4:   log likelihood = -140.61628
Mixed-effects logistic regression           Number of obs       =          46
Binomial variable:      napplied           Number of groups    =          23
Group variable:         department         Obs per group:
                                           min =              2
                                           avg =             2.0
                                           max =              2
Integration method: mvaghermite           Integration pts.    =          7
                                           Wald chi2(1)       =          2.14
Log likelihood = -140.61628               Prob > chi2        =          0.1435
```

nadmitted	Odds Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
female	1.176898	.1310535	1.46	0.144	.9461357	1.463944
_cons	.7907009	.2057191	-0.90	0.367	.4748457	1.316655
department						
var(_cons)	1.345383	.460702			.6876497	2.632234

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline odds (conditional on zero random effects).
LR test vs. logistic model: chibar2(01) = 318.77 Prob >= chibar2 = 0.0000

The odds of being admitted are higher for females than males but without statistical significance. The estimate of $\hat{\sigma}_u^2$ is 1.35 with the standard error of 0.46. An LR test comparing the model with the one-level binomial regression model favors the random-intercept model, indicating that there is a significant variation in the number of admissions between departments.


We can further assess the model fit by performing a residual analysis. For example, here we predict and plot Anscombe residuals.

```

. predict anscores, anscombe
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)

. twoway (scatter anscores department if female, msymbol(S))
> (scatter anscores department if !female, msymbol(T)),
> yline(-2 2) xline(1/23, lwidth(vvthin) lpattern(dash))
> xlabel(1/23) legend(label(1 "females") label(2 "males"))

```



Anscombe residuals are constructed to be approximately normally distributed, thus residuals that are above two in absolute value are usually considered outliers. In the graph above, the residual for female admissions in department 2 is a clear outlier, suggesting a poor fit for that particular observation; see [ME] [meglm postestimation](#) for more information about Anscombe residuals and other model diagnostics tools.

◀

Covariance structures

► Example 4: Growth-curve model with correlated random effects

Here we extend the model from [example 1](#) of [ME] [me](#) to allow for a random slope on `week` and an unstructured covariance structure between the random intercept and the random slope on `week`.

```
. use http://www.stata-press.com/data/r15/pig, clear
(Longitudinal analysis of pig weights)

. mixed weight week || id: week, covariance(unstructured)

Performing EM optimization:

Performing gradient-based optimization:

Iteration 0:    log likelihood = -868.96185
Iteration 1:    log likelihood = -868.96185

Computing standard errors:

Mixed-effects ML regression              Number of obs      =          432
Group variable: id                      Number of groups   =          48
                                         Obs per group:
                                         min =              9
                                         avg =             9.0
                                         max =              9

                                         Wald chi2(1)      =      4649.17
                                         Prob > chi2        =      0.0000

Log likelihood = -868.96185
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0910745	68.18	0.000	6.031393	6.388399
_cons	19.35561	.3996387	48.43	0.000	18.57234	20.13889


Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: Unstructured				
var(week)	.3715251	.0812958	.2419532	.570486
var(_cons)	6.823363	1.566194	4.351297	10.69986
cov(week,_cons)	-.0984378	.2545767	-.5973991	.4005234
var(Residual)	1.596829	.123198	1.372735	1.857505

LR test vs. linear model: $\chi^2(3) = 764.58$ Prob > $\chi^2 = 0.0000$
Note: LR test is conservative and provided only for reference.

The `unstructured` covariance structure allows for correlation between the random effects. Other covariance structures supported by `mixed`, besides the default `independent`, include `identity` and `exchangeable`; see [ME] `mixed` for details. You can also specify multiple random-effects equations at the same level, in which case the covariance types can be combined to form more complex blocked-diagonal covariance structures; see [example 5](#) below.

We can predict the fitted values and plot the estimated regression line for each of the pigs. The fitted values are based on both the fixed and the random effects.

```
. predict wgt_hat, fitted
. twoway connected wgt_hat week if id<=10, connect(L) ytitle("Predicted weight")
```



◀

► Example 5: Blocked-diagonal covariance structures

In this example, we fit a logistic mixed-effects model with a blocked-diagonal covariance structure of random effects.

We use the data from the 1989 Bangladesh fertility survey ([Huq and Cleland 1990](#)), which polled 1,934 Bangladeshi women on their use of contraception. The women sampled were from 60 districts, identified by the variable `district`. Each district contained either urban or rural areas (variable `urban`) or both. The variable `c_use` is the binary response, with a value of 1 indicating contraceptive use. Other covariates include mean-centered `age` and three indicator variables recording number of children. Below we fit a standard logistic regression model amended to have random coefficients on each indicator variable for children and an overall district random intercept.

```
. use http://www.stata-press.com/data/r15/bangladesh, clear
(Bangladesh Fertility Survey, 1989)
```

```
. melogit c_use urban age child* || district: child*, cov(exchangeable)
> || district:, or

Fitting fixed-effects model:
Iteration 0:   log likelihood = -1229.5485
Iteration 1:   log likelihood = -1228.5268
Iteration 2:   log likelihood = -1228.5263
Iteration 3:   log likelihood = -1228.5263

Refining starting values:
Grid node 0:   log likelihood = -1234.3979

Fitting full model:
Iteration 0:   log likelihood = -1234.3979 (not concave)
Iteration 1:   log likelihood = -1208.0052
Iteration 2:   log likelihood = -1206.4497
Iteration 3:   log likelihood = -1206.2417
Iteration 4:   log likelihood = -1206.2397
Iteration 5:   log likelihood = -1206.2397

Mixed-effects logistic regression
Group variable:      district
Number of obs       =      1,934
Number of groups    =        60
Obs per group:
    min =           2
    avg =          32.2
    max =          118

Integration method: mvaghermite
Integration pts.    =           7
Wald chi2(5)       =      100.01
Prob > chi2        =      0.0000
Log likelihood = -1206.2397
( 1)  [/var(child1[district]) - [/var(child3[district]) = 0
( 2)  [/cov(child1[district],child2[district]) -
      [/cov(child2[district],child3[district]) = 0
( 3)  [/cov(child1[district],child3[district]) -
      [/cov(child2[district],child3[district]) = 0
( 4)  [/var(child2[district]) - [/var(child3[district]) = 0
```

c_use	Odds Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	2.105163	.2546604	6.15	0.000	1.660796	2.668426
age	.9735765	.0077461	-3.37	0.001	.9585122	.9888775
child1	2.992596	.502149	6.53	0.000	2.153867	4.157931
child2	3.879345	.7094125	7.41	0.000	2.710815	5.551584
child3	3.774627	.7055812	7.11	0.000	2.616744	5.444863
_cons	.1859471	.0274813	-11.38	0.000	.1391841	.2484214
district						
var(child1)	.0841518	.0880698			.0108201	.654479
var(child2)	.0841518	.0880698			.0108201	.654479
var(child3)	.0841518	.0880698			.0108201	.654479
var(_cons)	.1870273	.0787274			.0819596	.426786
district						
cov(child1, child2)	.0616875	.0844681	0.73	0.465	-.1038669	.2272419
cov(child1, child3)	.0616875	.0844681	0.73	0.465	-.1038669	.2272419
cov(child2, child3)	.0616875	.0844681	0.73	0.465	-.1038669	.2272419

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline odds (conditional on zero random effects).
LR test vs. logistic model: chi2(3) = 44.57 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

The fixed effects can be interpreted just as you would the output from `logit`. Urban women have roughly double the odds of using contraception as compared with their rural counterparts. Having any number of children will increase the odds from three- to fourfold when compared with the base category of no children. Contraceptive use also decreases with age.

Because we specified `cov(exchangeable)`, the estimated variances on each indicator variable for children are constrained to be the same, and the estimated covariances on each indicator variable for children are constrained to be the same. More complex covariance structures with constraints can be specified using `covariance(pattern())` and `covariance(fixed())`; see [example 6](#) below. 

► Example 6: Meta analysis

In this example, we present a mixed-effects model for meta analysis of clinical trials. The term “meta-analysis” refers to a statistical analysis that involves summary data from similar but independent studies.

[Turner et al. \(2000\)](#) performed a study of nine clinical trials examining the effect of taking diuretics during pregnancy on the risk of pre-eclampsia. The summary data consist of the log odds-ratio (variable `or`) estimated from each study, and the corresponding estimated variance (variable `varor`). The square root of the variance is stored in the variable `std` and the trial identifier is stored in the variable `trial`.

```
. use http://www.stata-press.com/data/r15/diuretics
(Meta analysis of clinical trials studying diuretics and pre-eclampsia)
. list
```

	trial	or	varor	std
1.	1	.04	.16	.4
2.	2	-.92	.12	.3464102
3.	3	-1.12	.18	.4242641
4.	4	-1.47	.3	.5477226
5.	5	-1.39	.11	.3316625
6.	6	-.3	.01	.1
7.	7	-.26	.12	.3464102
8.	8	1.09	.69	.8306624
9.	9	.14	.07	.2645751

In a random-effects modeling of summary data, the observed log odds-ratios are treated as a continuous outcome and assumed to be normally distributed, and the true treatment effect varies randomly among the trials. The random-effects model can be written as

$$y_i \sim N(\theta + \nu_i, \sigma_i^2)$$

$$\nu_i \sim N(0, \tau^2)$$

where y_i is the observed treatment effect corresponding to the i th study, $\theta + \nu_i$ is the true treatment effect, σ_i^2 is the variance of the observed treatment effect, and τ is the between-trial variance component. Our aim is to estimate θ , the global mean.

Notice that the responses y_i do not provide enough information to estimate this model, because we cannot estimate the group-level variance component from a dataset that contains one observation per group. However, we already have estimates for the σ_i 's, therefore we can constrain each σ_i to

be equal to its estimated value, which will allow us to estimate θ and τ . We use `meglm` to estimate this model because the `mixed` command does not support constraints.

In `meglm`, one way to constrain a group of individual variances to specific values is by using the fixed covariance structure (an alternative way is to define each constraint individually with the `constraint` command and specify them in the `constraints()` option). The `covariance(fixed())` option requires a Stata matrix defining the constraints, thus we first create matrix `f` with the values of σ_i , stored in variable `varor`, on the main diagonal. We will use this matrix to constrain the variances.

```
. mkmat varor, mat(f)
. matrix f = diag(f)
```

In the random-effects equation part, we need to specify nine random slopes, one for each trial. Because random-effects equations do not support factor variables (see [U] 11.4.3 **Factor variables**), we cannot use the `i.trial` notation. Instead, we `tabulate` the variable `trial` and use the `generate()` option to create nine dummy variables named `tr1`, `tr2`, ..., `tr9`. We can then fit the model. Because the model is computationally demanding, we use Laplacian approximation instead of the default mean-variance adaptive quadrature; see [Computation time and the Laplacian approximation](#) above for details.

```

. qui tabulate trial, gen(tr)
. meglm or || _all: tr1-tr9, nocons cov(fixed(f)) intm(laplace) nocnsreport
Fitting fixed-effects model:
Iteration 0:   log likelihood = -10.643432
Iteration 1:   log likelihood = -10.643432
Refining starting values:
Grid node 0:   log likelihood = -10.205455
Fitting full model:
Iteration 0:   log likelihood = -10.205455
Iteration 1:   log likelihood = -9.4851561 (backed up)
Iteration 2:   log likelihood = -9.4587068
Iteration 3:   log likelihood = -9.4552982
Iteration 4:   log likelihood = -9.4552759
Iteration 5:   log likelihood = -9.4552759

Mixed-effects GLM                                Number of obs      =           9
Family:                                     Gaussian
Link:                                       identity
Group variable:                             _all              Number of groups   =           1
                                           Obs per group:
                                           min =               9
                                           avg =              9.0
                                           max =               9

Integration method:                             laplace

Log likelihood = -9.4552759                      Wald chi2(0)       =           .
                                           Prob > chi2        =           .

```

or	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_cons	-.5166151	.2059448	-2.51	0.012	-.9202594	-.1129707
_all						
var(tr1)	.16	(constrained)				
var(tr2)	.12	(constrained)				
var(tr3)	.18	(constrained)				
var(tr4)	.3	(constrained)				
var(tr5)	.11	(constrained)				
var(tr6)	.01	(constrained)				
var(tr7)	.12	(constrained)				
var(tr8)	.69	(constrained)				
var(tr9)	.07	(constrained)				
var(e.or)	.2377469	.1950926			.0476023	1.187413

We estimate $\hat{\theta} = -0.52$, which agrees with the estimate reported by [Turner et al. \(2000\)](#).

We can fit the above model in a more efficient way. We can consider the trials as nine independent random variables, each with variance unity, and each being multiplied by a different standard error. To accomplish this, we treat `trial` as a random-effects level, use the standard deviations of the log odds-ratios as a random covariate at the `trial` level, and constrain the variance component of `trial` to unity.

```
. constraint 1 _b[/var(std[trial])] = 1
. meglm or || trial: std, nocons constraints(1)
Fitting fixed-effects model:
Iteration 0:   log likelihood = -10.643432
Iteration 1:   log likelihood = -10.643432
Refining starting values:
Grid node 0:   log likelihood = -10.205455
Fitting full model:
Iteration 0:   log likelihood = -10.205455
Iteration 1:   log likelihood = -9.4851164 (backed up)
Iteration 2:   log likelihood = -9.45869
Iteration 3:   log likelihood = -9.4552794
Iteration 4:   log likelihood = -9.4552759
Iteration 5:   log likelihood = -9.4552759
Mixed-effects GLM                               Number of obs   =           9
Family:                               Gaussian
Link:                               identity
Group variable:                               trial
                                           Number of groups  =           9
                                           Obs per group:
                                           min =           1
                                           avg =          1.0
                                           max =           1
Integration method: mvaghermite                 Integration pts. =           7
Log likelihood = -9.4552759                      Wald chi2(0)     =           .
( 1)  [/var(std[trial])] = 1                     Prob > chi2      =           .
```

or	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_cons	-.5166151	.2059448	-2.51	0.012	-.9202594	-.1129708
trial						
var(std)	1 (constrained)					
var(e.or)	.2377469	.1950926			.0476023	1.187413

The results are the same, but this model took a fraction of the time compared with the less efficient specification.



Three-level models

The methods we have discussed so far extend from two-level models to models with three or more levels with nested random effects. By “nested”, we mean that the random effects shared within lower-level subgroups are unique to the upper-level groups. For example, assuming that classroom effects would be nested within schools would be natural, because classrooms are unique to schools. Below we illustrate a three-level mixed-effects ordered probit model.

Example 7: Three-level ordinal response model

In this example, we fit a three-level ordered probit model. The data are from the Television, School, and Family Smoking Prevention and Cessation Project (Flay et al. 1988; Rabe-Hesketh and Skrondal 2012, chap. 11), where schools were randomly assigned into one of four groups defined

Notes:

1. Our model now has two random-effects equations, separated by `||`. The first is a random intercept (constant only) at the `school` level (level three), and the second is a random intercept at the `class` level (level two). The order in which these are specified (from left to right) is significant—`meoprob` assumes that `class` is nested within `school`.
2. The information on groups is now displayed as a table, with one row for each grouping. You can suppress this table with the `nogroup` or the `noheader` option, which will also suppress the rest of the header.
3. The variance-component estimates are now organized and labeled according to level. The variance component for `class` is labeled `school>class` to emphasize that classes are nested within schools. ◀

The above extends to models with more than two levels of nesting in the obvious manner, by adding more random-effects equations, each separated by `||`. The order of nesting goes from left to right as the groups go from biggest (highest level) to smallest (lowest level).

Crossed-effects models

Not all mixed-effects models contain nested levels of random effects.

► Example 8: Crossed random effects

Returning to our longitudinal analysis of pig weights, suppose that we wish to fit

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_i + v_j + \epsilon_{ij} \tag{8}$$

for the $i = 1, \dots, 9$ weeks and $j = 1, \dots, 48$ pigs and

$$u_i \sim N(0, \sigma_u^2); \quad v_j \sim N(0, \sigma_v^2); \quad \epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$$

all independently. That is, we assume an overall population-average growth curve $\beta_0 + \beta_1 \text{week}$ and a random pig-specific shift. In other words, the effect due to `week`, u_i , is systematic to that week and common to all pigs. The rationale behind (8) could be that, assuming that the pigs were measured contemporaneously, we might be concerned that week-specific random factors such as weather and feeding patterns had significant systematic effects on all pigs.

Model (8) is an example of a two-way crossed-effects model, with the pig effects v_j being crossed with the week effects u_i . One way to fit such models is to consider all the data as one big cluster, and treat u_i and v_j as a series of $9 + 48 = 57$ random coefficients on indicator variables for `week` and `pig`. The random effects **u** and the variance components **G** are now represented as

$$\mathbf{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_9 \\ v_1 \\ \vdots \\ v_{48} \end{bmatrix} \sim N(\mathbf{0}, \mathbf{G}); \quad \mathbf{G} = \begin{bmatrix} \sigma_u^2 \mathbf{I}_9 & \mathbf{0} \\ \mathbf{0} & \sigma_v^2 \mathbf{I}_{48} \end{bmatrix}$$

Because $\mathbf{G}$ is block diagonal, it can be represented as repeated-level equations. All we need is an ID variable to identify all the observations as one big group and a way to tell mixed-effects commands to treat `week` and `pig` as factor variables (or equivalently, as two sets of overparameterized indicator variables identifying weeks and pigs, respectively). The mixed-effects commands support the special group designation `_all` for the former and the `R.varname` notation for the latter.

```
. use http://www.stata-press.com/data/r15/pig
(Longitudinal analysis of pig weights)
. mixed weight week || _all: R.id || _all: R.week
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -1013.824
Iteration 1:   log likelihood = -1013.824
Computing standard errors:
Mixed-effects ML regression              Number of obs      =          432
Group variable: _all                     Number of groups   =           1
                                         Obs per group:
                                             min =          432
                                             avg  =         432.0
                                             max  =          432
                                         Wald chi2(1)       =       13258.28
                                         Prob > chi2        =         0.0000

Log likelihood = -1013.824
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0539313	115.14	0.000	6.104192	6.315599
_cons	19.35561	.6333982	30.56	0.000	18.11418	20.59705

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
_all: Identity					
	var(R.id)	14.83623	3.126142	9.816733	22.42231
_all: Identity					
	var(R.week)	.0849874	.0868856	.0114588	.6303302
	var(Residual)	4.297328	.3134404	3.724888	4.957741

LR test vs. linear model: chi2(2) = 474.85 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

We estimate $\hat{\sigma}_u^2 = 0.08$ and $\hat{\sigma}_v^2 = 14.84$.

The `R.varname` notation is equivalent to giving a list of overparameterized (none dropped) indicator variables for use in a random-effects specification. When you use `R.varname`, mixed-effects commands handle the calculations internally rather than creating the indicators in the data. Because the set of indicators is overparameterized, `R.varname` implies `noconstant`.

Note that the column dimension of our random-effects design is 57. Computation time and memory requirements grow (roughly) quadratically with the dimension of the random effects. As a result, fitting such crossed-effects models is feasible only when the total column dimension is small to moderate. For this reason, mixed-effects commands use the Laplacian approximation as the default estimation method for crossed-effects models; see [Computation time and the Laplacian approximation](#) above for more details.

It is often possible to rewrite a mixed-effects model in a way that is more computationally efficient. For example, we can treat pigs as nested within the `_all` group, yielding the equivalent and more efficient (total column dimension 10) way to fit (8):

```
. mixed weight week || _all: R.week || id:
```

The results of both estimations are identical, but the latter specification, organized at the cluster (pig) level with random-effects dimension 1 (a random intercept) is much more computationally efficient. Whereas with the first form we are limited in how many pigs we can analyze, there is no such limitation with the second form.

All the mixed-effects commands—except `mixed`, `meqrlogit`, and `meqrpoisson`—automatically attempt to recast the less efficient model specification into a more efficient one. However, this automatic conversion may not be sufficient for some complicated mixed-effects specifications, especially if both crossed and nested effects are involved. Therefore, we strongly encourage you to always specify the more efficient syntax; see [Rabe-Hesketh and Skrondal \(2012\)](#) and [Marchenko \(2006\)](#) for additional techniques to make calculations more efficient in more complex mixed-effects models.




Nonlinear models

NLME models are popular in population pharmacokinetics, bioassays, studies of biological and agricultural growth processes, and other applications, where the mean function is a nonlinear function of fixed and random effects. [Remarks and examples](#) of **[ME] menl** provide many examples of fitting different NLME models by using `menl`, including a pharmacokinetics model in [example 15](#). Here we consider simple data from [Draper and Smith \(1998\)](#) that contain trunk circumference (in mm) of five different orange trees measured over seven different time points.

Let's plot our data first.

```
. use http://www.stata-press.com/data/r15/orange
(Growth of orange trees (Draper and Smith, 1998))
. twoway scatter circumf age, connect(L) ylabel(#6 175)
```



Consider the following nonlinear growth model for these data,

$$\text{circumf}_{ij} = \frac{\beta_1}{1 + \exp \left\{ -(\text{age}_{ij} - \beta_2) / \beta_3 \right\}} + \epsilon_{ij}$$

where ϵ_{ij} 's are i.i.d. $N(0, \sigma_\epsilon^2)$. In this model, β_1 can be interpreted as the average asymptotic trunk circumference of trees as $\text{age}_{ij} \rightarrow \infty$. We can crudely estimate it as the average of the trunk circumference values at the last observed time point, which for these data is roughly 175 mm. β_2 is the age at which a tree attains half of the average asymptotic trunk circumference β_1 ; that is, if we set $\text{age}_{ij} = \beta_2$, then $E(\text{circumf}_{ij}) = 0.5\beta_1$. β_3 is a scale parameter that represents the number of days it takes for a tree to grow from 50% to about 73% of the average asymptotic trunk circumference. That is, if we set $\text{age} = t_{0.73} = \beta_2 + \beta_3$, then $E(\text{circumf}_{ij}) = \beta_1 / \{1 + \exp(-1)\} = 0.73\beta_1$ and then $\beta_3 = t_{0.73} - \beta_2$.

The above model can be easily fit by using, for example, `nl`; see [R] [nl](#). However, if we study the graph more carefully, we will notice that there is an increasing variability in the trunk circumferences of trees as they approach their limiting age. So it may be more reasonable to allow β_1 to vary between trees,

$$\text{circumf}_{ij} = \frac{\beta_1 + u_{1j}}{1 + \exp \left\{ -(\text{age}_{ij} - \beta_2) / \beta_3 \right\}} + \epsilon_{ij} \quad (9)$$

where u_{1j} 's are i.i.d. $N(0, \sigma_{u_1}^2)$. We use `menl` to fit this model.

The specification of NLME models in `menl` is fairly straightforward. Following the dependent variable and the equality sign (=), we specify the expression for the mean function as a usual Stata expression but with parameters and random effects enclosed in curly braces (`{}`).


```
. menl circumf = ({b1}+{U1[tree]})/(1+exp(-(age-{b2})/{b3}))
Obtaining starting values by EM:
Alternating PNLs/LME algorithm:
Iteration 1:    linearization log likelihood = -131.584579
Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      35
Group variable: tree                      Number of groups   =       5
                                           Obs per group:
                                           min =             7
                                           avg =            7.0
                                           max =             7

Linearization log likelihood = -131.58458
```

circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b1	191.049	16.15403	11.83	0.000	159.3877	222.7103
/b2	722.556	35.15082	20.56	0.000	653.6616	791.4503
/b3	344.1624	27.14739	12.68	0.000	290.9545	397.3703

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Identity	var(U1)	991.1514	639.4636	279.8776	3510.038
	var(Residual)	61.56371	15.89568	37.11466	102.1184

In the above specification, we used {U1[tree]} to include random intercepts at the tree level in our model. U1 is the name or label associated with these random intercepts.

The output of menl is similar to that of mixed—the header information is displayed first, fixed-effects parameter estimates are displayed in the first or the fixed-effects parameter table, and the estimates of variance components are displayed in the second or the random-effects parameter table.

The header information is similar to that of mixed, but unlike mixed, menl in general does not report a model χ^2 statistic in the header because a test of the joint significance of all fixed-effects parameters (except the constant term) may not be relevant in a nonlinear model. menl also reports the so-called linearization log likelihood. menl uses the linearization method of Lindstrom and Bates (1990), with extensions from Pinheiro and Bates (1995), for estimation. This method is based on the approximation of the NLME model by an LME model, in which a first-order Taylor-series approximation is used to linearize the nonlinear mean function with respect to fixed and random effects; see Introduction and Methods and formulas in [ME] menl for details. The linearization log likelihood is the log likelihood of this approximating LME model. We can use this log likelihood for model comparison of different NLME models and to form likelihood-ratio tests, but note that this is not the log likelihood of the corresponding NLME model. Depending on the accuracy of the approximation, the linearization log likelihood may be close to the true NLME log likelihood.

As part of Stata’s standard estimation output, menl reports z tests against zeros for the estimated fixed-effects parameters. Testing a parameter against zero may or may not be of interest, or may not even be appropriate, in a nonlinear model. In our example, {b3} is the denominator of a fraction, so the test of {b3} against zero may not be feasible in this model. Instead, we may be interested in testing {b3} against, for example, 300, which would correspond to testing whether the average trunk circumference of orange trees increases from 50% to 73% of its asymptotic value in 300 days. We can perform this test by using, for instance, the test command; see [R] test. As a side note, setting $\beta_3 = 0$ in (9) results in a simple random-intercept model, in a limiting sense.

From the random-effects table, the variability in limiting growth β_1 between trees, labeled as `var(U1)`, is statistically significant in this model with an estimate of 991 (mm²) and a 95% CI of [280, 3510].

We can rewrite (9) as a two-stage model,

$$\text{circumf}_{ij} = \frac{\phi_{1j}}{1 + \exp\left\{-\left(\text{age}_{ij} - \phi_{2j}\right) / \phi_{3j}\right\}} + \epsilon_{ij} \quad (10)$$

where the stage 2 specification is

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 \\ \beta_3 \end{bmatrix} \quad (11)$$

The model defined by (10) and (11) is the same as that defined by (9) but with a different parameterization.

In `menl`, we can accommodate this two-stage formulation with the `define()` option. For example, we can fit the two-stage model defined by (10) and (11) as follows:

```
. menl circumf = {phi1:}/(1+exp(-(age-{b2})/{b3})), define(phi1: {b1}+{U1[tree]})
Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:    linearization log likelihood = -131.584579
Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      35
Group variable: tree                      Number of groups   =       5
                                           Obs per group:
                                           min =             7
                                           avg =            7.0
                                           max =             7

Linearization log likelihood = -131.58458
           phi1:  {b1}+{U1[tree]}
```

circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b1	191.049	16.15403	11.83	0.000	159.3877	222.7103
/b2	722.556	35.15082	20.56	0.000	653.6616	791.4503
/b3	344.1624	27.14739	12.68	0.000	290.9545	397.3703

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Identity					
	var(U1)	991.1514	639.4636	279.8776	3510.038
	var(Residual)	61.56371	15.89568	37.11466	102.1184

The results are identical to the previous model. Here we defined a substitutable expression `phi1` in the `define()` option as a function of `{b1}` and `{U1[tree]}` and included it in our main expression as `{phi1:}`. Including a colon (`:`) in `{phi1:}` is important to notify `menl` that it is a substitutable expression rather than a simple scalar parameter `{phi1}`.

In general, we can accommodate multistage formulations by using the `define()` option repeatedly.

More conveniently, we can use a linear-combination specification (see [Linear combinations](#) in [ME] **menl**) within the `define()` option to define the linear combination `{b1}+{U1[tree]}`.

```
. menl circumf = {phi1:}/(1+exp(-(age-{b2})/{b3})), define(phi1: U1[tree], xb)
Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:    linearization log likelihood = -131.584579
Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      35
Group variable: tree                      Number of groups   =       5
                                           Obs per group:
                                           min =              7
                                           avg =             7.0
                                           max =              7
```

Linearization log likelihood = -131.58458
phi1: U1[tree], xb

circumf		Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1	_cons	191.049	16.15403	11.83	0.000	159.3877	222.7103
	/b2	722.556	35.15082	20.56	0.000	653.6616	791.4503
	/b3	344.1624	27.14739	12.68	0.000	290.9545	397.3703

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Identity					
	var(U1)	991.1514	639.4636	279.8776	3510.038
	var(Residual)	61.56371	15.89568	37.11466	102.1184

The `{phi1: U1[tree], xb}` specification used in the `define()` option, but without curly braces, creates a linear combination named `phi1` that contains a constant `{phi1: _cons}` and random intercepts `{U1}` at the `tree` level. In the linear-combination specification, the constant is included automatically unless you specify the `noconstant` option such as `{phi1: U1[tree], xb noconstant}`. Also, you do not specify curly braces around random effects within the linear-combination specification. If we had covariates, say, `x1` and `x2`, that we also wanted to include in the linear combination, we would have used `{phi1: x1 x2 U1[tree]}`. Notice that we did not specify the `xb` option in the previous linear combination. When a linear combination contains more than one term, this option is implied. When a linear combination contains only one term, such as in `{phi1: U1[tree], xb}`, the `xb` option must be specified to request that `menl` treat the specification as a linear combination instead of a scalar parameter; see [Random-effects substitutable expressions](#) in [ME] **menl** for details.

Instead of using `define()`, we could have similarly specified the linear combination directly in the main expression:

```
. menl circumf = {phi1: U1[tree], xb}/(1+exp(-(age-{b2})/{b3}))
(output omitted)
```

However, by using the `define()` option, we simplified the look of the main equation.

We can extend the stage 2 specification (11) to allow, for example, β_2 to vary across trees by including random intercepts at the tree level for ϕ_{2j} ,

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 + u_{2j} \\ \beta_3 \end{bmatrix}$$

We can then fit the corresponding model by using `menl` as follows:

```
. menl circumf = {phi1:}/(1+exp(-(age-{phi2:})/{b3})),
> define(phi1: U1[tree], xb) define(phi2: U2[tree], xb)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1:  linearization log likelihood = -131.605387
Iteration 2:  linearization log likelihood = -131.5827
Iteration 3:  linearization log likelihood = -131.580503
Iteration 4:  linearization log likelihood = -131.580265
Iteration 5:  linearization log likelihood = -131.580257
```

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	35
Group variable: tree	Number of groups	=	5
	Obs per group:		
	min	=	7
	avg	=	7.0
	max	=	7

Linearization log likelihood = -131.58026

```
phi1:  U1[tree], xb
phi2:  U2[tree], xb
```

	circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
phi1	_cons	190.5939	16.211	11.76	0.000	158.8209 222.3669
phi2	_cons	719.6027	35.77597	20.11	0.000	649.4831 789.7223
	/b3	342.0794	26.42036	12.95	0.000	290.2965 393.8624

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Independent					
	var(U1)	1012.15	666.2807	278.557	3677.697
	var(U2)	503.2308	2401.323	.043651	5801504
	var(Residual)	59.27073	18.21297	32.45483	108.2434

The large standard error for the estimate of the variance component `var(U2)` suggests that our model is overparameterized—a common problem when fitting NLME models. We could verify this, for instance, by computing information criteria ([R] [estimates stats](#)) or by performing a likelihood-ratio test ([R] [lrtest](#)).

By default, `menl` assumes an independent covariance structure for the random effects such as `U1` and `U2` in our example. We can specify, for example, an unstructured model by using the `covariance()` option. We demonstrate this only for illustration, given that our simpler model that assumed independence between `U1` and `U2` was already overparameterized.

```
. menl circumf = {phi1:}/(1+exp(-(age-{phi2:})/{b3})),
> define(phi1: U1[tree], xb) define(phi2: U2[tree], xb)
> covariance(U1 U2, unstructured)
```

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = -130.90452
Iteration 2: linearization log likelihood = -130.902051
Iteration 3: linearization log likelihood = -130.901772
Iteration 4: linearization log likelihood = -130.901766
```

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	35
Group variable: tree	Number of groups	=	5
	Obs per group:		
	min	=	7
	avg	=	7.0
	max	=	7

Linearization log likelihood = -130.90177

```
phi1: U1[tree], xb
phi2: U2[tree], xb
```

circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	189.8349	17.20035	11.04	0.000	156.1228	223.5469
phi2						
_cons	709.5333	37.24229	19.05	0.000	636.5397	782.5268
/b3	340.4731	25.52176	13.34	0.000	290.4514	390.4948

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Unstructured					
	var(U1)	1180.097	775.0819	325.7264	4275.458
	var(U2)	1469.879	2777.132	36.22882	59636.04
	cov(U1,U2)	1015.504	1124.568	-1188.609	3219.617
	var(Residual)	56.07332	16.20294	31.82681	98.79145

In `menl`, we need to list the names of the random effects in the `covariance()` option for which we want to specify a covariance structure other than the independent one used by default.

In our example, parameters ϕ_{1j} and ϕ_{2j} were modeled as linear functions of random effects and parameters β_1 and β_2 . The relationship does not have to be linear; see [example 15](#) in [\[ME\] menl](#).

This example has a small number of trees or clusters, so REML estimation would have been more appropriate. We could have obtained REML estimates in our examples by specifying the `reml` option with `menl`.

See [\[ME\] menl](#) for more examples of and details about the `menl` command.

Acknowledgments

We are indebted to Sophia Rabe-Hesketh of the University of California, Berkeley, and coauthor of the Stata Press book *Multilevel and Longitudinal Modeling Using Stata*; Anders Skrondal of the University of Oslo and the Norwegian Institute of Public Health, and coauthor of the Stata Press book *Multilevel and Longitudinal Modeling Using Stata*; and Andrew Pickles of the University of Manchester for their extensive body of work in Stata, both previous and ongoing, in this area.

References

- Agresti, A. 2013. *Categorical Data Analysis*. 3rd ed. Hoboken, NJ: Wiley.
- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies.
<http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Breslow, N. E., and D. G. Clayton. 1993. Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* 88: 9–25.
- Davidian, M., and D. M. Giltinan. 1995. *Nonlinear Models for Repeated Measurement Data*. Boca Raton, FL: Chapman & Hall/CRC.
- . 2003. Nonlinear models for repeated measurement data: An overview and update. *Journal of Agricultural, Biological, and Environmental Statistics* 8: 387–419.
- De Boeck, P., and M. Wilson, ed. 2004. *Explanatory Item Response Models: A Generalized Linear and Nonlinear Approach*. New York: Springer.
- Demidenko, E. 2013. *Mixed Models: Theory and Applications with R*. 2nd ed. Hoboken, NJ: Wiley.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39: 1–38.
- Diggle, P. J., P. J. Heagerty, K.-Y. Liang, and S. L. Zeger. 2002. *Analysis of Longitudinal Data*. 2nd ed. Oxford: Oxford University Press.
- Draper, N., and H. Smith. 1998. *Applied Regression Analysis*. 3rd ed. New York: Wiley.
- Fitzmaurice, G. M., M. Davidian, G. Verbeke, and G. Molenberghs, ed. 2009. *Longitudinal Data Analysis*. Boca Raton, FL: Chapman & Hall/CRC.
- Flay, B. R., B. R. Brannon, C. A. Johnson, W. B. Hansen, A. L. Ulene, D. A. Whitney-Saltiel, L. R. Gleason, S. Sussman, M. D. Gavin, K. M. Glowacz, D. F. Sobol, and D. C. Spiegel. 1988. The television, school, and family smoking cessation and prevention project: I. Theoretical basis and program development. *Preventive Medicine* 17: 585–607.
- Gelman, A., and J. Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge: Cambridge University Press.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Hall, B. H., Z. Griliches, and J. A. Hausman. 1986. Patents and R and D: Is there a lag? *International Economic Review* 27: 265–283.
- Hedeker, D., and R. D. Gibbons. 2006. *Longitudinal Data Analysis*. Hoboken, NJ: Wiley.
- Henderson, C. R. 1953. Estimation of variance and covariance components. *Biometrics* 9: 226–252.
- Huq, N. M., and J. Cleland. 1990. *Bangladesh Fertility Survey 1989 (Main Report)*. National Institute of Population Research and Training.
- Kuehl, R. O. 2000. *Design of Experiments: Statistical Principles of Research Design and Analysis*. 2nd ed. Belmont, CA: Duxbury.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- LaMotte, L. R. 1973. Quadratic estimation of variance components. *Biometrics* 29: 311–330.

- Lesaffre, E., and B. Spiessens. 2001. On the effect of the number of quadrature points in a logistic random-effects model: An example. *Journal of the Royal Statistical Society, Series C* 50: 325–335.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Lindstrom, M. J., and D. M. Bates. 1990. Nonlinear mixed effects models for repeated measures data. *Biometrics* 46: 673–687.
- Littell, R. C., G. A. Milliken, W. W. Stroup, R. D. Wolfinger, and O. Schabenberger. 2006. *SAS System for Mixed Models*. 2nd ed. Cary, NC: SAS Institute.
- Liu, Q., and D. A. Pierce. 1994. A note on Gauss–Hermite quadrature. *Biometrika* 81: 624–629.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Ng, E. S.-W., J. R. Carpenter, H. Goldstein, and J. Rasbash. 2006. Estimation in generalised linear mixed models with binary outcomes by simulated maximum likelihood. *Statistical Modelling* 6: 23–42.
- Pinheiro, J. C., and D. M. Bates. 1995. Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics* 4: 12–35.
- . 2000. *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- Pinheiro, J. C., and E. C. Chao. 2006. Efficient Laplacian and adaptive Gaussian quadrature algorithms for multilevel generalized linear mixed models. *Journal of Computational and Graphical Statistics* 15: 58–81.
- Rabe-Hesketh, S., and A. Skrondal. 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Rao, C. R. 1973. *Linear Statistical Inference and Its Applications*. 2nd ed. New York: Wiley.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Rodríguez, G., and N. Goldman. 1995. An assessment of estimation procedures for multilevel models with binary responses. *Journal of the Royal Statistical Society, Series A* 158: 73–89.
- Ruppert, D., M. P. Wand, and R. J. Carroll. 2003. *Semiparametric Regression*. Cambridge: Cambridge University Press.
- Searle, S. R., G. Casella, and C. E. McCulloch. 1992. *Variance Components*. New York: Wiley.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.
- Skrondal, A., and S. Rabe-Hesketh. 2004. [Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models](#). Boca Raton, FL: Chapman & Hall/CRC.
- Stram, D. O., and J. W. Lee. 1994. Variance components testing in the longitudinal mixed effects model. *Biometrics* 50: 1171–1177.
- Thompson, W. A., Jr. 1962. The problem of negative estimates of variance components. *Annals of Mathematical Statistics* 33: 273–289.
- Tierney, L., and J. B. Kadane. 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* 81: 82–86.
- Turner, R. M., R. Z. Omar, M. Yang, H. Goldstein, and S. G. Thompson. 2000. A multilevel model framework for meta-analysis of clinical trials with binary outcomes. *Statistics in Medicine* 19: 3417–3432.
- Tutz, G., and W. Hennevoigl. 1996. Random effects in ordinal regression models. *Computational Statistics & Data Analysis* 22: 537–557.
- Vella, F., and M. Verbeek. 1998. Whose wages do unions raise? A dynamic model of unionism and wage rate determination for young men. *Journal of Applied Econometrics* 13: 163–183.
- Verbeke, G., and G. Molenberghs. 2000. *Linear Mixed Models for Longitudinal Data*. New York: Springer.

Vonesh, E., and V. M. Chinchilli. 1997. *Linear and Nonlinear Models for the Analysis of Repeated Measurements*. New York: Dekker.

Also see

[ME] [Glossary](#)

Title

estat df — Calculate degrees of freedom for fixed effects

Description

Remarks and examples

Menu for estat

Stored results

Syntax

Also see

Options

Description

`estat df` is for use after estimation with `mixed`.

`estat df` calculates and displays the degrees of freedom (DF) for each fixed effect using the specified methods. This allows for a comparison of different DF methods. `estat df` can also be used to continue with postestimation using a different DF method without rerunning the model.

Menu for estat

Statistics > Postestimation

Syntax

```
estat df [ , method\(df\_methods\) post[ (df\_method) ] eim oim]
```

Options

`method(df_methods)` specifies a list of methods to compute DF. The supported methods are `residual`, `repeated`, `anova`, `satterthwaite`, and `kroger`; more than one method may be specified. Methods `satterthwaite` and `kroger` are only available with REML estimation. If option `dfmethod()` was not specified in the most recently fit `mixed` model, then option `method()` is required. See [Small-sample inference for fixed effects](#) under *Remarks and examples* in [ME] `mixed` for more details.

`post` causes `estat df` to behave like a Stata estimation command. When `post` is specified, `estat df` will post the DF for each fixed effect as well as everything related to the DF computation to `e()` for the method specified in `method()`. Thus, after posting, you could continue to use this DF for other postestimation commands. For example, you could use `test`, `small` to perform Wald *F* tests on linear combination of the fixed effects.

`post` may also be specified using the syntax `post(df_method)`. You must use this syntax if you specify multiple `df_methods` in option `method()`. With this syntax, `estat df` computes the DF using the method specified in `post()` and stores the results in `e()`. Only one computation method may be specified using the syntax `post()`.

The `df_method` specified in `post()` must be one of the DF methods specified in option `method()`. If only one method is specified in option `method()`, then one can simply use `post` to make this DF method active for postestimation and for `mixed` replay.

`eim` specifies that the expected information matrix be used in the DF computation. It can be used only when `method()` contains `kroger` or `satterthwaite`. `eim` is the default.

oim specifies that the observed information matrix be used in the DF computation. It can be used only when method() contains kroger or satterthwaite.

Remarks and examples

► Example 1: Changing the degrees of freedom method

To illustrate the use of estat df, we refit the dental veneer data from [example 14](#) of [ME] mixed using the Kenward–Roger method (option dfmethod(kroger)) to compute the DF for fixed effects.

```
. use http://www.stata-press.com/data/r15/veneer
(Dental veneer data)
. mixed gcf followup base_gcf cda age || patient: followup,
> covariance(unstructured) || tooth:, reml nolog dfmethod(kroger)
Mixed-effects REML regression                               Number of obs   =          110
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
patient	12	2	9.2	12
tooth	55	2	2.0	2

```
DF method: Kenward-Roger                                DF:                min =          10.41
                                                         avg =          28.96
                                                         max =          50.71
```

```
F(4, 27.96) = 1.47
Log restricted-likelihood = -420.92761                    Prob > F           = 0.2370
```

gcf	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
followup	.3009815	1.938641	0.16	0.879	-3.96767	4.569633
base_gcf	-.0183127	.1466261	-0.12	0.901	-.3132419	.2766164
cda	-.329303	.5533506	-0.60	0.554	-1.440355	.7817493
age	-.5773932	.2350491	-2.46	0.033	-1.098324	-.056462
_cons	45.73862	13.21824	3.46	0.002	18.53866	72.93858

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
patient: Unstructured					
	var(followup)	41.88772	18.79997	17.38009	100.9535
	var(_cons)	524.9851	253.0205	204.1287	1350.175
	cov(followup,_cons)	-140.4229	66.57623	-270.9099	-9.935907
tooth: Identity					
	var(_cons)	47.45738	16.63034	23.8792	94.3165
	var(Residual)	48.86704	10.50523	32.06479	74.47382

```
LR test vs. linear model: chi2(4) = 91.12                Prob > chi2 = 0.0000
```

Note: LR test is conservative and provided only for reference.

Rather than specifying option `dftable(pvalue)` or `dftable(ci)` at estimation, we can display the covariate-specific DFs during postestimation by typing

```
. estat df
Degrees of freedom
```

	Kenward-Roger
gcf	
followup	10.96355
base_gcf	47.2708
cda	50.70932
age	10.41127
_cons	25.43377

`estat df` can also compare different DF methods using the `method()` option. For example, we can compare the Kenward–Roger method with the Satterthwaite method by typing

```
. estat df, method(kroger satterthwaite)
Degrees of freedom
```

	Kenward-Roger	Satterthwaite
gcf		
followup	10.96355	10.96355
base_gcf	47.2708	47.2708
cda	50.70932	50.70932
age	10.41127	10.41127
_cons	25.43377	25.43377

The two methods produce the same estimates of DFs for single-hypothesis tests, but the results differ for multiple-hypotheses tests; see [example 4](#) of [\[ME\] mixed postestimation](#) for details.

Suppose that we decide to proceed with the Satterthwaite method in subsequent analysis. Rather than retyping our `mixed` command with the `dfmethod(satterthwaite)` option, we can post the Satterthwaite DFs using the `post` option of `estat df`.

```
. estat df, method(satterthwaite) post
Degrees of freedom
```

	Satterthwaite
gcf	
followup	10.96355
base_gcf	47.2708
cda	50.70932
age	10.41127
_cons	25.43377

The returned values associated with `dfmethod(kroger)` from the `mixed` command will be replaced with those of `dfmethod(satterthwaite)`.

Stored results

`estat df` stores the following in `r()`:

Macros	
<code>r(dfmethods)</code>	DF methods
Matrices	
<code>r(df)</code>	parameter-specific DFs for each method specified in <code>method()</code>
<code>r(V_df)</code>	variance–covariance matrix of the estimators when <code>kroger</code> method is specified

If option `post()` is specified, `estat df` also stores the following in `e()`:

Scalars	
<code>e(F)</code>	overall F test statistic for the method specified in <code>post()</code>
<code>e(ddf_m)</code>	model DDF for the method specified in <code>post()</code>
<code>e(df_max)</code>	maximum DF for the method specified in <code>post()</code>
<code>e(df_avg)</code>	average DF for the method specified in <code>post()</code>
<code>e(df_min)</code>	minimum DF for the method specified in <code>post()</code>
Macros	
<code>e(dfmethod)</code>	DF method specified in <code>post()</code>
<code>e(dftitle)</code>	title for DF method
Matrices	
<code>e(df)</code>	parameter-specific DFs for the method specified in <code>post()</code>
<code>e(V_df)</code>	variance–covariance matrix of the estimators when <code>kroger</code> method is posted

Also see

[ME] [mixed](#) — Multilevel mixed-effects linear regression

[U] [20 Estimation and postestimation commands](#)

Title

estat group — Summarize the composition of the nested groups

[Description](#)
[Also see](#)

[Menu for estat](#)

[Syntax](#)

[Remarks and examples](#)

Description

`estat group` reports the number of groups and minimum, average, and maximum group sizes for each level of the model. Model levels are identified by the corresponding group variable in the data. Because groups are treated as nested, the information in this summary may differ from what you would get if you used the `tabulate` command on each group variable individually.

Menu for estat

Statistics > Postestimation

Syntax

`estat group`

Remarks and examples

See [example 3](#) in [\[ME\] mixed postestimation](#) and [example 4](#) in [\[ME\] menl postestimation](#).

Also see

- [\[ME\] mecloglog](#) — Multilevel mixed-effects complementary log-log regression
- [\[ME\] meglm](#) — Multilevel mixed-effects generalized linear model
- [\[ME\] meintreg](#) — Multilevel mixed-effects interval regression
- [\[ME\] melogit](#) — Multilevel mixed-effects logistic regression
- [\[ME\] menbreg](#) — Multilevel mixed-effects negative binomial regression
- [\[ME\] menl](#) — Nonlinear mixed-effects regression
- [\[ME\] meologit](#) — Multilevel mixed-effects ordered logistic regression
- [\[ME\] meoprobit](#) — Multilevel mixed-effects ordered probit regression
- [\[ME\] mepoisson](#) — Multilevel mixed-effects Poisson regression
- [\[ME\] meprobit](#) — Multilevel mixed-effects probit regression
- [\[ME\] meqrlogit](#) — Multilevel mixed-effects logistic regression (QR decomposition)
- [\[ME\] meqrpoisson](#) — Multilevel mixed-effects Poisson regression (QR decomposition)
- [\[ME\] mestreg](#) — Multilevel mixed-effects parametric survival models
- [\[ME\] metobit](#) — Multilevel mixed-effects tobit regression
- [\[ME\] mixed](#) — Multilevel mixed-effects linear regression
- [\[U\] 20 Estimation and postestimation commands](#)

Title

estat icc — Estimate intraclass correlations

Description	Menu for estat	Syntax	Options
Remarks and examples	Stored results	Methods and formulas	Also see

Description

estat icc is for use after estimation with `mixed`, `meintreg`, `metobit`, `melogit`, `meqrlogit`, and `meprobit`.

estat icc displays the intraclass correlation for pairs of responses at each nested level of the model. Intraclass correlations are available for random-intercept models or for random-coefficient models conditional on random-effects covariates being equal to 0. They are not available for crossed-effects models or with residual error structures other than independent structures.

Menu for estat

Statistics > Postestimation

Syntax

```
estat icc [ , level(#) ]
```

Options

`level(#)` specifies the confidence level, as a percentage, for confidence intervals. The default is `level(95)` or as set by `set level`; see [U] 20.8 Specifying the width of confidence intervals.

Remarks and examples

See, for instance, example 2 in [ME] [mixed postestimation](#) and examples 1 and 4 in [ME] [melogit postestimation](#).

Stored results

estat icc stores the following in `r()`:

Scalars	
<code>r(icc#)</code>	level-# intraclass correlation
<code>r(se#)</code>	standard errors of level-# intraclass correlation
<code>r(level)</code>	confidence level of confidence intervals
Macros	
<code>r(label#)</code>	label for level #
Matrices	
<code>r(ci#)</code>	vector of confidence intervals (lower and upper) for level-# intraclass correlation

For a G -level nested model, # can be any integer between 2 and G .

Methods and formulas

Intraclass correlations

Consider a simple, two-level random-intercept model, stated in terms of a latent linear response, where only $y_{ij} = I(y_{ij}^* > 0)$ is observed for the latent variable,

$$y_{ij}^* = \beta + u_j^{(2)} + \epsilon_{ij}^{(1)}$$

with $i = 1, \dots, n_j$ and level-2 groups $j = 1, \dots, M$. Here β is an unknown fixed intercept, $u_j^{(2)}$ is a level-2 random intercept, and $\epsilon_{ij}^{(1)}$ is a level-1 error term. In a mixed-effects linear and probit regression, errors are assumed to be normally distributed with mean 0 and variance γ . In a mixed-effects logistic regression, errors are assumed to be logistic with mean 0 and variance γ . Random intercepts are assumed to be normally distributed with mean 0 and variance σ_2^2 and to be independent of error terms.

The intraclass correlation for this model is

$$\rho = \text{Corr}(y_{ij}^*, y_{i'j}^*) = \frac{\sigma_2^2}{\gamma + \sigma_2^2}$$

where $\gamma = \sigma_1^2$ for a mixed-effects linear regression, $\gamma = 1$ for a mixed-effects probit regression, and $\gamma = \pi^2/3$ for a mixed-effects logistic regression. The intraclass correlation corresponds to the correlation between the latent responses i and i' from the same group j .

Now consider a three-level nested random-intercept model,

$$y_{ijk}^* = \beta + u_{jk}^{(2)} + u_k^{(3)} + \epsilon_{ijk}^{(1)}$$

for measurements $i = 1, \dots, n_{jk}$ and level-2 groups $j = 1, \dots, M_{1k}$ nested within level-3 groups $k = 1, \dots, M_2$. Here $u_{jk}^{(2)}$ is a level-2 random intercept, $u_k^{(3)}$ is a level-3 random intercept, and $\epsilon_{ijk}^{(1)}$ is a level-1 error term. The random intercepts are assumed to be normally distributed with mean 0 and variances σ_2^2 and σ_3^2 , respectively, and to be mutually independent. The error terms are also independent of the random intercepts.

We can consider two types of intraclass correlations for this model. We will refer to them as level-2 and level-3 intraclass correlations. The level-3 intraclass correlation is

$$\rho^{(3)} = \text{Corr}(y_{ijk}^*, y_{i'j'k}^*) = \frac{\sigma_3^2}{\gamma + \sigma_2^2 + \sigma_3^2}$$

This is the correlation between latent responses i and i' from the same level-3 group k and from different level-2 groups j and j' .

The level-2 intraclass correlation is

$$\rho^{(2)} = \text{Corr}(y_{ijk}^*, y_{i'jk}^*) = \frac{\sigma_2^2 + \sigma_3^2}{\gamma + \sigma_2^2 + \sigma_3^2}$$

This is the correlation between latent responses i and i' from the same level-3 group k and level-2 group j . (Note that level-1 intraclass correlation is undefined.)

More generally, for a G -level nested random-intercept model, the g -level intraclass correlation is defined as

$$\rho^{(g)} = \frac{\sum_{l=g}^G \sigma_l^2}{\gamma + \sum_{l=2}^G \sigma_l^2}$$

The above formulas also apply in the presence of fixed-effects covariates $\mathbf{X}$ in a random-effects model. In this case, intraclass correlations are conditional on fixed-effects covariates and are referred to as residual intraclass correlations. `estat icc` also uses the same formulas to compute intraclass correlations for random-coefficient models, assuming 0 baseline values for the random-effects covariates, and labels them as conditional intraclass correlations.

Intraclass correlations will always fall in $[0,1]$ because variance components are nonnegative. To accommodate the range of an intraclass correlation, we use the logit transformation to obtain confidence intervals. We use the delta method to estimate the standard errors of the intraclass correlations.

Let $\hat{\rho}^{(g)}$ be a point estimate of the intraclass correlation and $\widehat{\text{SE}}(\hat{\rho}^{(g)})$ be its standard error. The $(1 - \alpha) \times 100\%$ confidence interval for $\text{logit}(\rho^{(g)})$ is

$$\text{logit}(\hat{\rho}^{(g)}) \pm z_{\alpha/2} \frac{\widehat{\text{SE}}(\hat{\rho}^{(g)})}{\hat{\rho}^{(g)}(1 - \hat{\rho}^{(g)})}$$

where $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal distribution and $\text{logit}(x) = \ln\{x/(1-x)\}$. Let k_u be the upper endpoint of this interval, and let k_l be the lower. The $(1 - \alpha) \times 100\%$ confidence interval for $\rho^{(g)}$ is then given by

$$\left(\frac{1}{1 + e^{-k_l}}, \frac{1}{1 + e^{-k_u}} \right)$$

Also see

[ME] [meintreg](#) — Multilevel mixed-effects interval regression

[ME] [melogit](#) — Multilevel mixed-effects logistic regression

[ME] [meprobit](#) — Multilevel mixed-effects probit regression

[ME] [meqrlogit](#) — Multilevel mixed-effects logistic regression (QR decomposition)

[ME] [metobit](#) — Multilevel mixed-effects tobit regression

[ME] [mixed](#) — Multilevel mixed-effects linear regression

[U] [20 Estimation and postestimation commands](#)

Title

estat recovariance — Display estimated random-effects covariance matrices

Description

Remarks and examples

Menu for estat

Stored results

Syntax

Also see

Options

Description

`estat recovariance` is for use after estimation with `menl`, `mixed`, `meqrlogit`, and `meqrpoisson`.

`estat recovariance` displays the estimated variance–covariance matrix of the random effects for each level in the model.

Menu for estat

Statistics > Postestimation

Syntax

`estat recovariance` [`, relevel(levelvar) correlation matlist_options`]

Options

`relevel(levelvar)` specifies the level in the model for which the random-effects covariance matrix is to be displayed. By default, the covariance matrices for all levels in the model are displayed. *levelvar* is the name of the model level and is either the name of the variable describing the grouping at that level or is `_all`, a special designation for a group comprising all the estimation data. The `_all` designation is not supported with `menl`.

`correlation` displays the covariance matrix as a correlation matrix.

matlist_options are style and formatting options that control how the matrix (or matrices) is displayed; see [P] [matlist](#) for a list of options that are available.

Remarks and examples

For `menl`, the rows and columns of the matrix are labeled with [full random-effects names](#) as they are defined in the model.

For other commands, the rows and columns of the matrix are labeled as `_cons` for the random intercepts; for random coefficients, the label is the name of the associated variable in the data.

See [example 1](#) in [ME] [mixed postestimation](#).

Stored results

estat recovariance stores the following in `r()`:

Scalars

`r(relevels)` number of levels

Matrices

`r(Cov#)` level-# random-effects covariance matrix

`r(Corr#)` level-# random-effects correlation matrix (if option `correlation` was specified)

For a G -level nested model, # can be any integer between 2 and G .

Also see

[ME] [menl](#) — Nonlinear mixed-effects regression

[ME] [meqrlogit](#) — Multilevel mixed-effects logistic regression (QR decomposition)

[ME] [meqrpoisson](#) — Multilevel mixed-effects Poisson regression (QR decomposition)

[ME] [mixed](#) — Multilevel mixed-effects linear regression

[U] [20 Estimation and postestimation commands](#)

Title

estat sd — Display variance components as standard deviations and correlations

[Description](#)

[Remarks and examples](#)

[Menu for estat](#)

[Stored results](#)

[Syntax](#)

[Also see](#)

[Options](#)

Description

`estat sd` displays the random-effects and within-group error parameter estimates as standard deviations and correlations.

Menu for estat

Statistics > Postestimation

Syntax

```
estat sd [ , variance verbose post coeflegend ]
```

Options

`variance` specifies that `estat sd` display the random-effects and within-group error parameter estimates as variances and covariances. If the `post` option is specified, the estimated variances and covariances and their respective standard errors are posted to `e()`. `variance` is allowed only after `menl`.

`verbose` specifies that the full estimation table be displayed. By default, only the random-effects and within-group error parameters are displayed. This option is implied when `post` is specified.

`post` causes `estat sd` to behave like a Stata estimation (e-class) command. `estat sd` posts the vector of calculated standard deviation and correlation parameters along with the corresponding variance–covariance matrix to `e()`, so that you can treat the estimated parameters just as you would results from any other estimation command. For example, you could use `test` to perform simultaneous tests of hypotheses on the parameters, or you could use `lincom` to create linear combinations.

The following option is not shown in the dialog box:

`coeflegend` specifies that the legend of the coefficients and how to specify them in an expression be displayed rather than displaying the statistics for the coefficients. This option is not supported by `mixed`.

Remarks and examples

See [example 1](#) in [\[ME\] mixed postestimation](#) and [example 16](#) in [\[ME\] menl](#).

Stored results

`estat sd` stores the following in `r()`:

Matrices

<code>r(b)</code>	coefficient vector
<code>r(V)</code>	variance–covariance matrix of the estimators

If `post` is specified, `estat sd` stores the following in `e()`:

Matrices

<code>e(b)</code>	coefficient vector
<code>e(V)</code>	variance–covariance matrix of the estimators

Also see

- [ME] [mecloglog](#) — Multilevel mixed-effects complementary log-log regression
- [ME] [meglm](#) — Multilevel mixed-effects generalized linear model
- [ME] [meintreg](#) — Multilevel mixed-effects interval regression
- [ME] [melogit](#) — Multilevel mixed-effects logistic regression
- [ME] [menbreg](#) — Multilevel mixed-effects negative binomial regression
- [ME] [menl](#) — Nonlinear mixed-effects regression
- [ME] [meologit](#) — Multilevel mixed-effects ordered logistic regression
- [ME] [meoprobit](#) — Multilevel mixed-effects ordered probit regression
- [ME] [mepoisson](#) — Multilevel mixed-effects Poisson regression
- [ME] [meprobit](#) — Multilevel mixed-effects probit regression
- [ME] [meqrllogit](#) — Multilevel mixed-effects logistic regression (QR decomposition)
- [ME] [meqrpoisson](#) — Multilevel mixed-effects Poisson regression (QR decomposition)
- [ME] [mestreg](#) — Multilevel mixed-effects parametric survival models
- [ME] [metobit](#) — Multilevel mixed-effects tobit regression
- [ME] [mixed](#) — Multilevel mixed-effects linear regression
- [U] [20 Estimation and postestimation commands](#)

[Description](#)
[Remarks and examples](#)
[Also see](#)

[Menu for estat](#)
[Stored results](#)

[Syntax](#)
[Methods and formulas](#)

[Options](#)
[Reference](#)

Description

estat wcorrelation is for use after estimation with menl and mixed.

estat wcorrelation displays the overall correlation matrix for a given cluster calculated on the basis of the design of the random effects and their assumed covariance and the correlation structure of the residuals. This allows for a comparison of different multilevel models in terms of the ultimate within-cluster correlation matrix that each model implies.

Menu for estat

Statistics > Postestimation

Syntax

estat wcorrelation [, *options*]

<i>options</i>	Description
at(<i>at_spec</i>)	specify the cluster for which you want the correlation matrix; default is the first two-level cluster encountered in the data
all	display correlation matrix for all the data
<u>c</u> ovariance	display the covariance matrix instead of the correlation matrix
list	list the data corresponding to the correlation matrix
nosort	list the rows and columns of the correlation matrix in the order they were originally present in the data
<u>i</u> terate(#)	maximum number of iterations to compute random effects; default is <code>iterate(50)</code> ; only for use after <code>menl</code>
<u>t</u> olerance(#)	convergence tolerance when computing random effects; default is <code>tolerance(1e-4)</code> ; only for use after <code>menl</code>
format(<i>%fmt</i>)	set the display format; default is <code>format(%6.3f)</code>
<i>matlist_options</i>	style and formatting options that control how matrices are displayed

Options

`at(at_spec)` specifies the cluster of observations for which you want the within-cluster correlation matrix. *at_spec* is

```
relevel_var = value [, relevel_var = value ...]
```

For example, if you specify

```
. estat wcorrelation, at(school = 33)
```

you get the within-cluster correlation matrix for those observations in school 33. If you specify

```
. estat wcorrelation, at(school = 33 classroom = 4)
```

you get the correlation matrix for classroom 4 in school 33.

If `at()` is not specified, then you get the correlations for the first level-two cluster encountered in the data. This is usually what you want.

`all` specifies that you want the correlation matrix for all the data. This is not recommended unless you have a relatively small dataset or you enjoy seeing large $n \times n$ matrices. However, this can prove useful in some cases.

`covariance` specifies that the within-cluster covariance matrix be displayed instead of the default correlations and standard deviations.

`list` lists the model data for those observations depicted in the displayed correlation matrix. With linear mixed-effects models, this option is also useful if you have many random-effects design variables and you wish to see the represented values of these design variables.

`nosort` lists the rows and columns of the correlation matrix in the order that they were originally present in the data. Normally, `estat wcorrelation` will first sort the data according to level variables, by-group variables, and time variables to produce correlation matrices whose rows and columns follow a natural ordering. `nosort` suppresses this.

`iterate(#)` specifies the maximum number of iterations when computing estimates of the random effects. The default is `iterate(50)`. This option is only for use after `menl`.

`tolerance(#)` specifies a convergence tolerance when computing estimates of the random effects. The default is `tolerance(1e-4)`. This option is only for use after `menl`.

`format(%fnt)` sets the display format for the standard-deviation vector and correlation matrix. The default is `format(%6.3f)`.

matlist_options are style and formatting options that control how the matrix (or matrices) is displayed; see [P] [matlist](#) for a list of options that are available.

Remarks and examples

► Example 1: Displaying within-cluster correlations for different clusters

Here we fit a model where different clusters have different within-cluster correlations, and we show how to display them for different clusters. We use the Asian children weight data from [example 6](#) of [\[ME\] mixed](#).

```
. use http://www.stata-press.com/data/r15/childweight
(Weight data on Asian children)
. mixed weight age || id: age, covariance(unstructured)
```

Performing EM optimization:

Performing gradient-based optimization:

```
Iteration 0:  log likelihood = -344.37065
Iteration 1:  log likelihood = -342.83887
Iteration 2:  log likelihood = -342.71863
Iteration 3:  log likelihood = -342.71777
Iteration 4:  log likelihood = -342.71777
```

Computing standard errors:

Mixed-effects ML regression	Number of obs	=	198
Group variable: id	Number of groups	=	68
	Obs per group:		
	min	=	1
	avg	=	2.9
	max	=	5
	Wald chi2(1)	=	755.27
Log likelihood = -342.71777	Prob > chi2	=	0.0000

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
age	3.459671	.1258877	27.48	0.000	3.212936	3.706406
_cons	5.110496	.1494781	34.19	0.000	4.817524	5.403468

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: Unstructured				
var(age)	.202392	.1242868	.0607406	.6743838
var(_cons)	.0970272	.1107998	.0103483	.9097447
cov(age,_cons)	.140134	.0566901	.0290234	.2512445
var(Residual)	1.357922	.1650502	1.070076	1.723198

LR test vs. linear model: chi2(3) = 27.38 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

We use `estat wcorrelation` to display the within-cluster correlations for the first cluster.

```
. estat wcorrelation, list
Standard deviations and correlations for id = 45:
Standard deviations:
      obs |      1      2      3      4      5
-----|-----
      sd |  1.224  1.314  1.448  1.506  1.771
Correlations:
      obs |      1      2      3      4      5
-----|-----
       1 |  1.000
       2 |  0.141  1.000
       3 |  0.181  0.274  1.000
       4 |  0.193  0.293  0.376  1.000
       5 |  0.230  0.348  0.447  0.477  1.000
```

Data:

	id	weight	age
1.	45	5.171	.136893
2.	45	10.86	.657084
3.	45	13.15	1.21834
4.	45	13.2	1.42916
5.	45	15.88	2.27242

We specified the `list` option to display the data associated with the cluster. The next cluster in the dataset has ID 258. To display the within-cluster correlations for this cluster, we specify the `at()` option.

```
. estat wcorrelation, at(id=258) list
Standard deviations and correlations for id = 258:
Standard deviations:
      obs |      1      2      3      4
-----|-----
      sd |  1.231  1.320  1.424  1.782
Correlations:
      obs |      1      2      3      4
-----|-----
       1 |  1.000
       2 |  0.152  1.000
       3 |  0.186  0.270  1.000
       4 |  0.244  0.356  0.435  1.000
```

Data:

	id	weight	age
1.	258	5.3	.19165
2.	258	9.74	.687201
3.	258	9.98	1.12799
4.	258	11.34	2.30527

The within-cluster correlations for this model depend on age. The values for `age` in the two clusters are different, as are the corresponding within-cluster correlations.

See [example 1](#) of [\[ME\] mixed postestimation](#) for a model fit where each cluster had the same model-implied within-cluster correlations.

Stored results

`estat wcorrelation` stores the following in `r()`:

Matrices

<code>r(sd)</code>	standard deviations
<code>r(Corr)</code>	within-cluster correlation matrix
<code>r(Cov)</code>	within-cluster variance–covariance matrix
<code>r(G)</code>	variance–covariance matrix of random effects
<code>r(Z)</code>	model-based design matrix
<code>r(R)</code>	variance–covariance matrix of level-one errors

Results `r(G)`, `r(Z)`, and `r(R)` are available only after `mixed`.

Methods and formulas

Methods and formulas are presented under the following headings:

Linear mixed-effects model

Nonlinear mixed-effects model

Linear mixed-effects model

A two-level linear mixed model of the form

$$\mathbf{y}_j = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j + \boldsymbol{\epsilon}_j$$

implies the marginal model

$$\mathbf{y}_j = \mathbf{X}_j\boldsymbol{\beta} + \boldsymbol{\epsilon}_j^*$$

where $\boldsymbol{\epsilon}_j^* \sim N(\mathbf{0}, \mathbf{V}_j)$, $\mathbf{V}_j = \mathbf{Z}_j\mathbf{G}\mathbf{Z}_j' + \mathbf{R}$. In a marginal model, the random part is described in terms of the marginal or total residuals $\boldsymbol{\epsilon}_j^*$, and $\mathbf{V}_j$ is the covariance structure of these residuals.

`estat wcorrelation` calculates the marginal covariance matrix $\tilde{\mathbf{V}}_j$ for cluster j and by default displays the results in terms of standard deviations and correlations. This allows for a comparison of different multilevel models in terms of the ultimate within-cluster correlation matrix that each model implies.

Calculation of the marginal covariance matrix extends naturally to higher-level models; see, for example, chapter 4.8 in [West, Welch, and Galecki \(2015\)](#).

Nonlinear mixed-effects model

For nonlinear mixed-effects models, there is no closed-form expression for the marginal covariance matrix $\text{Cov}(\mathbf{y}_j)$. This is because it is expressed in terms of a q -dimensional integral (q is the number of random effects in the model), which, in general, is analytically intractable. Under the linear mixed-effects approximation, the marginal covariance matrix is estimated by $\hat{\mathbf{V}}_j = \hat{\mathbf{Z}}_j\hat{\boldsymbol{\Sigma}}\hat{\mathbf{Z}}_j' + \hat{\sigma}^2\hat{\mathbf{A}}_j$, where $\hat{\mathbf{Z}}_j$, $\hat{\boldsymbol{\Sigma}}$, and $\hat{\mathbf{A}}_j$ are defined in [Methods and formulas](#) of [\[ME\] menl](#).

`estat wcorrelation` calculates the estimated marginal covariance matrix $\widehat{\mathbf{V}}_j$ for cluster j and by default displays the results in terms of standard deviations and correlations.

Under the linear mixed-effects approximation, estimation of the marginal covariance matrix extends naturally to higher-level models; see, for example, chapter 4.8 in [West, Welch, and Galecki \(2015\)](#).

Reference

West, B. T., K. B. Welch, and A. T. Galecki. 2015. *Linear Mixed Models: A Practical Guide Using Statistical Software*. 2nd ed. Boca Raton, FL: Chapman & Hall/CRC.

Also see

[ME] [menl](#) — Nonlinear mixed-effects regression

[ME] [mixed](#) — Multilevel mixed-effects linear regression

[U] [20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
Reference	Also see		

Description

mecloglog fits mixed-effects models for binary or binomial responses. The conditional distribution of the response given the random effects is assumed to be Bernoulli, with probability of success determined by the inverse complementary log-log function.

Quick start

Two-level complementary log-log model of y on x with random intercepts by lev2

```
mecloglog y x || lev2:
```

Add binary variable a and random coefficients for a

```
mecloglog y x a || lev2: a
```

As above, but allow the random effects to be correlated

```
mecloglog y x a || lev2: a, covariance(unstructured)
```

Three-level random-intercept model of y on x with lev2 nested within lev3

```
mecloglog y x || lev3: || lev2:
```

Crossed-effects model of y on x with two-way crossed random effects by factors a and b

```
mecloglog y x || _all:R.a || b:
```

Menu

Statistics > Multilevel mixed-effects models > Complementary log-log regression

Syntax

```
mecloglog depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1
<code>asis</code>	retain perfect predictor variables

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>binomial</u> (<i>varname</i> #)	set binomial trials if data are in binomial form
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>eform</u>	report exponentiated coefficients
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<i>maximize_options</i>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<code>mvaghermite</code>	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<code>mcaghermite</code>	mode-curvature adaptive Gauss–Hermite quadrature
<code>ghermite</code>	nonadaptive Gauss–Hermite quadrature
<code>laplace</code>	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: mecloglog.

vce() and weights are not allowed with the *svy* prefix; see [SVY] svy.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

`noconstant` suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

`offset(varname)` specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

`asis` forces retention of perfect predictor variables and their associated, perfectly predicted observations and may produce instabilities in maximization; see [R] probit.

`covariance(vartype)` specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: `independent`, `exchangeable`, `identity`, `unstructured`, `fixed(matname)`, or `pattern(matname)`.

`covariance(independent)` covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is `covariance(independent)` unless a crossed random-effects model is fit, in which case the default is `covariance(identity)`.

`covariance(exchangeable)` structure specifies one common variance for all random effects and one common pairwise covariance.

`covariance(identity)` is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

`covariance(unstructured)` allows for all variances and covariances to be distinct. If an equation consists of *p* random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

`covariance(fixed(matname))` and `covariance(pattern(matname))` covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted.

In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of `matname`. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if `matname[i, j] = matname[k, l]`.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`binomial(varname | #)` specifies that the data are in binomial form; that is, `depvar` records the number of successes from a series of binomial trials. This number of trials is given either as `varname`, which allows this number to vary over the observations, or as the constant `#`. If `binomial()` is not specified (the default), `depvar` is treated as Bernoulli, with any nonzero, nonmissing values indicating positive responses.

`constraints(constraints)`, `collinear`; see [R] [estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] [estimation options](#).

`eform` reports exponentiated coefficients and corresponding standard errors and confidence intervals.

This option may be specified either at estimation or upon replay.

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `noci`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrntolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `mecloglog` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `mecloglog` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[ (gridspec) ]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects complementary log-log (cloglog) regression is cloglog regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

mecloglog allows for many levels of random effects. However, for simplicity, we here consider the two-level model, where for a series of M independent clusters, and conditional on a set of fixed effects $\mathbf{x}_{ij}$ and a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = 1 | \mathbf{x}_{ij}, \mathbf{u}_j) = H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are the binary-valued y_{ij} , and we follow the standard Stata convention of treating $y_{ij} = 1$ if *depvar* _{ij} $\neq 0$ and treating $y_{ij} = 0$ otherwise. The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard cloglog regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$, so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Finally, because this is cloglog regression, $H(\cdot)$ is the inverse of the complementary log-log function that maps the linear predictor to the probability of a success ($y_{ij} = 1$) with $H(v) = 1 - \exp\{-\exp(v)\}$.

Model (1) may also be stated in terms of a latent linear response, where only $y_{ij} = I(y_{ij}^* > 0)$ is observed for the latent

$$y_{ij}^* = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

The errors ϵ_{ij} are independent and identically extreme-value (Gumbel) distributed with the mean equal to Euler's constant and variance $\sigma_\epsilon^2 = \pi^2/6$, independently of $\mathbf{u}_j$. This nonsymmetric error distribution is an alternative to the symmetric error distribution underlying logistic and probit analysis and is usually used when the positive (or negative) outcome is rare.

Below we present two short examples of mixed-effects cloglog regression; refer to [ME] **me** and [ME] **meglm** for examples of other random-effects models. A two-level cloglog model can also be fit using **xtcloglog** with the **re** option; see [XT] **xtcloglog**. In the absence of random effects, mixed-effects cloglog regression reduces to standard cloglog regression; see [R] **cloglog**.

► Example 1: Two-level random-intercept model

In example 1 of [XT] **xtcloglog**, we analyze unionization of women in the United States over the period 1970–1988. The women are identified by the variable **idcode**. Here we refit that model with **mecloglog**. Because the original example used 12 integration points by default, we request 12 integration points as well.

```
. use http://www.stata-press.com/data/r15/union
(NLS Women 14-24 in 1968)
```

```
. mecloglog union age grade not_smsa south##c.year || idcode:, intpoints(12)
Fitting fixed-effects model:
Iteration 0:   log likelihood = -14237.139
Iteration 1:   log likelihood = -13546.159
Iteration 2:   log likelihood = -13540.611
Iteration 3:   log likelihood = -13540.607
Iteration 4:   log likelihood = -13540.607
Refining starting values:
Grid node 0:   log likelihood = -11104.448
Fitting full model:
Iteration 0:   log likelihood = -11104.448
Iteration 1:   log likelihood = -10617.891
Iteration 2:   log likelihood = -10537.919
Iteration 3:   log likelihood = -10535.946
Iteration 4:   log likelihood = -10535.941
Iteration 5:   log likelihood = -10535.941
Mixed-effects cloglog regression
Group variable:      idcode
Number of obs       =      26,200
Number of groups    =       4,434
Obs per group:
    min =           1
    avg =          5.9
    max =          12
Integration method: mvaghermite
Integration pts.    =          12
Wald chi2(6)       =      248.12
Prob > chi2        =      0.0000
Log likelihood = -10535.941
```

union	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
age	.0128542	.0119441	1.08	0.282	-.0105559	.0362642
grade	.0699965	.0138551	5.05	0.000	.0428409	.097152
not_smsa	-.1982009	.0649258	-3.05	0.002	-.3254531	-.0709488
1.south	-2.049901	.4892644	-4.19	0.000	-3.008842	-1.090961
year	-.0006158	.0123999	-0.05	0.960	-.0249191	.0236875
south#c.year						
1	.0164457	.0060685	2.71	0.007	.0045516	.0283399
_cons	-3.277375	.6610552	-4.96	0.000	-4.57302	-1.981731
idcode						
var(_cons)	3.489803	.1630921			3.184351	3.824555

```
LR test vs. cloglog model: chibar2(01) = 6009.33      Prob >= chibar2 = 0.0000
```

The estimates are practically the same. `xtcloglog` reports the estimated variance component as a standard deviation, $\hat{\sigma}_u = 1.86$. `mecloglog` reports $\hat{\sigma}_u^2 = 3.49$, the square root of which is 1.87. We find that age and education each have a positive effect on union membership, although the former is not statistically significant. Women who live outside of metropolitan areas are less likely to unionize.

The estimated variance of the random intercept at the individual level, $\hat{\sigma}^2$, is 3.49 with standard error 0.16. The reported likelihood-ratio test shows that there is enough variability between women to favor a mixed-effects cloglog regression over an ordinary cloglog regression; see [Distribution theory for likelihood-ratio test](#) in [ME] **me** for a discussion of likelihood-ratio testing of variance components.

► Example 2: Three-level random-intercept model

Rabe-Hesketh, Touloupoulou, and Murray (2001) analyzed data from a study that measured the cognitive ability of patients with schizophrenia compared with their relatives and control subjects. Cognitive ability was measured as the successful completion of the “Tower of London”, a computerized task, measured at three levels of difficulty. For all but one of the 226 subjects, there were three measurements (one for each difficulty level). Because patients’ relatives were also tested, a family identifier, `family`, was also recorded.

We fit a cloglog model with response `dtlm`, the indicator of cognitive function, and with covariates `difficulty` and a set of indicator variables for `group`, with the controls (`group==1`) being the base category. We also allow for random effects due to families and due to subjects within families. The first is a random intercept (constant only) at the `family` level, and the second is a random intercept at the `subject` level. The order in which these are specified (from left to right) is significant—`mecloglog` assumes that `subject` is nested within `family`. The equations are separated by `||`.

```

. use http://www.stata-press.com/data/r15/towerlondon
(Tower of London data)
. mecloglog dtlm difficulty i.group || family: || subject:
Fitting fixed-effects model:
Iteration 0:   log likelihood = -337.21921
Iteration 1:   log likelihood = -313.79023
Iteration 2:   log likelihood = -313.56906
Iteration 3:   log likelihood = -313.56888
Iteration 4:   log likelihood = -313.56888
Refining starting values:
Grid node 0:   log likelihood = -314.57061
Fitting full model:
Iteration 0:   log likelihood = -314.57061   (not concave)
Iteration 1:   log likelihood = -308.82101
Iteration 2:   log likelihood = -305.71841
Iteration 3:   log likelihood = -305.26804
Iteration 4:   log likelihood = -305.26516
Iteration 5:   log likelihood = -305.26516
Mixed-effects cloglog regression                               Number of obs   =           677

```

Group Variable	No. of Groups	Observations per Group Minimum	Average	Maximum
family	118	2	5.7	27
subject	226	2	3.0	3

```

Integration method: mvaghermite                               Integration pts. =           7
Wald chi2(3) = 83.32
Log likelihood = -305.26516                                   Prob > chi2    = 0.0000

```

dtlm	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
difficulty	-1.342844	.1501508	-8.94	0.000	-1.637135 -1.048554
group					
2	-.1331007	.269389	-0.49	0.621	-.6610935 .3948922
3	-.7714314	.3097099	-2.49	0.013	-1.378452 -.164411
_cons	-1.6718	.2290325	-7.30	0.000	-2.120695 -1.222905
family					
var(_cons)	.2353453	.2924064			.0206122 2.687117
family>					
subject					
var(_cons)	.7737687	.4260653			.2629714 2.276742

```

LR test vs. cloglog model: chi2(2) = 16.61                   Prob > chi2 = 0.0002
Note: LR test is conservative and provided only for reference.

```

After adjusting for the random-effects structure, the probability of successful completion of the Tower of London decreases dramatically as the level of difficulty increases. Also, schizophrenics (group==3) tended not to perform as well as the control subjects.



The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by ||.

Stored results

mecloglog stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>mecloglog</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	<code>cloglog</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>cloglog</code>
<code>e(family)</code>	bernoulli or binomial
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	offset
<code>e(binomial)</code>	binomial number of trials
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vctype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>

e(marginswexp)	weight expression for margins
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved
Matrices	
e(b)	coefficient vector
e(Cns)	constraints matrix
e(ilog)	iteration log (up to 20 iterations)
e(gradient)	gradient vector
e(N_g)	group counts
e(g_min)	group-size minimums
e(g_avg)	group-size averages
e(g_max)	group-size maximums
e(V)	variance-covariance matrix of the estimators
e(V_modelbased)	model-based variance
Functions	
e(sample)	marks estimation sample

Methods and formulas

`mecloglog` is a convenience command for `meglm` with a `cloglog` link and a `bernoulli` or `binomial` family; see [ME] [meglm](#).

Model (1) assumes Bernoulli data, a special case of the binomial. Because binomial data are also supported by `mecloglog` (option `binomial()`), the methods presented below are in terms of the more general binomial mixed-effects model.

For a two-level binomial model, consider the response y_{ij} as the number of successes from a series of r_{ij} Bernoulli trials (replications). For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$f(\mathbf{y}_j | \mathbf{u}_j) = \prod_{i=1}^{n_j} \left[\binom{r_{ij}}{y_{ij}} \{H(\eta_{ij})\}^{y_{ij}} \{1 - H(\eta_{ij})\}^{r_{ij} - y_{ij}} \right] \\ = \exp \left(\sum_{i=1}^{n_j} \left[y_{ij} \log \{H(\eta_{ij})\} - (r_{ij} - y_{ij}) \exp(\eta_{ij}) + \log \binom{r_{ij}}{y_{ij}} \right] \right)$$

for $\eta_{ij} = \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j + \text{offset}_{ij}$ and $H(v) = 1 - \exp\{-\exp(v)\}$.

Defining $\mathbf{r}_j = (r_{j1}, \dots, r_{jn_j})'$ and

$$c(\mathbf{y}_j, \mathbf{r}_j) = \sum_{i=1}^{n_j} \log \binom{r_{ij}}{y_{ij}}$$

where $c(\mathbf{y}_j, \mathbf{r}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j | \mathbf{u}_j) = \exp [\mathbf{y}_j' \log \{H(\boldsymbol{\eta}_j)\} - (\mathbf{r}_j - \mathbf{y}_j)' \exp(\boldsymbol{\eta}_j) + c(\mathbf{y}_j, \mathbf{r}_j)]$$

where $\boldsymbol{\eta}_j$ is formed by stacking the row vectors $\boldsymbol{\eta}_{ij}$. We extend the definitions of the functions $H(\cdot)$, $\log(\cdot)$, and $\exp(\cdot)$ to be vector functions where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ , the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned}\mathcal{L}_j(\beta, \Sigma) &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int f(\mathbf{y}_j | \mathbf{u}_j) \exp(-\mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp\{c(\mathbf{y}_j, \mathbf{r}_j)\} (2\pi)^{-q/2} |\Sigma|^{-1/2} \int \exp\{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j\end{aligned}\tag{2}$$

where

$$h(\beta, \Sigma, \mathbf{u}_j) = \mathbf{y}_j' \log\{H(\boldsymbol{\eta}_j)\} - (\mathbf{r}_j - \mathbf{y}_j)' \exp(\boldsymbol{\eta}_j) - \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see *Methods and formulas* in [ME] **meglm** for details.

mecloglog supports multilevel weights and survey data; see *Methods and formulas* in [ME] **meglm** for details.

Reference

Rabe-Hesketh, S., T. Touloupoulou, and R. M. Murray. 2001. Multilevel modeling of cognitive function in schizophrenic patients and their first degree relatives. *Multivariate Behavioral Research* 36: 279–298.

Also see

- [ME] **mecloglog postestimation** — Postestimation tools for **mecloglog**
- [ME] **melogit** — Multilevel mixed-effects logistic regression
- [ME] **meprobbit** — Multilevel mixed-effects probit regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: mecloglog** — Bayesian multilevel complementary log-log regression
- [SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtcloglog** — Random-effects and population-averaged cloglog models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation command is of special interest after `mecloglog`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
density	predicted density function
distribution	predicted distribution function
pearson	Pearson residuals
deviance	deviance residuals
anscombe	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
Integration	
<i>int_options</i>	integration options
pearson, deviance, anscombe may not be combined with <code>marginal</code> .	
<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only
<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<i>int_options</i>	integration options
<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`mu`, the default, calculates the predicted mean, that is, the probability of a positive outcome.

`eta`, `xb`, `stdp`, `density`, `distribution`, `pearson`, `deviance`, `anscombe`, `scores`, `conditional()`, `marginal`, and `nooffset`; see [ME] [meglm postestimation](#).

`reffects`, `ebmeans`, `ebmodes`, and `reses()`; see [ME] [meglm postestimation](#).

Integration

`intpoints()`, `iterate()`, and `tolerance()`; see [ME] [meglm postestimation](#).

margins

Description for margins

`margins` estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
<u>density</u>	not allowed with margins
<u>distribution</u>	not allowed with margins
<u>pearson</u>	not allowed with margins
<u>deviance</u>	not allowed with margins
<u>anscombe</u>	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options `conditional(ebmeans)` and `conditional(ebmodes)` are not allowed with `margins`.
Option `marginal` is assumed where applicable if `conditional(fixedonly)` is not specified.

Statistics not allowed with `margins` are functions of stochastic quantities other than $e(b)$.
For the full syntax, see [\[R\] margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects complementary log-log model with `mecloglog`. Here we show a short example of predicted probabilities and predicted random effects; refer to [\[ME\] meglm postestimation](#) for additional examples.

► Example 1: Obtaining predicted probabilities and random effects

In [example 2](#) of [\[ME\] mecloglog](#), we analyzed the cognitive ability (`dt1m`) of patients with schizophrenia compared with their relatives and control subjects, by using a three-level complementary log-log model with random effects at the family and subject levels. Cognitive ability was measured as the successful completion of the “Tower of London”, a computerized task, measured at three levels of difficulty.

Here we list the data for family 16:

```
. list family subject dtlm pr re1 re2 if family==16, sepby(subject)
```

	family	subject	dtlm	pr	re1	re2
208.	16	5	1	.486453	.4184933	.2760492
209.	16	5	0	.1597047	.4184933	.2760492
210.	16	5	0	.0444156	.4184933	.2760492
211.	16	34	1	.9659582	.4184933	1.261488
212.	16	34	1	.5862808	.4184933	1.261488
213.	16	34	1	.205816	.4184933	1.261488
214.	16	35	0	.5571261	.4184933	-.1616545
215.	16	35	1	.1915688	.4184933	-.1616545
216.	16	35	0	.0540124	.4184933	-.1616545

We can see that the predicted random effects (**re1**) at the family level are the same for all members of the family. Similarly, the predicted random effects (**re2**) at the individual level are constant within each individual.



Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [\[ME\] meglm postestimation](#).

Also see

- [\[ME\] mecloglog](#) — Multilevel mixed-effects complementary log-log regression
- [\[ME\] meglm postestimation](#) — Postestimation tools for meglm
- [\[U\] 20 Estimation and postestimation commands](#)

[Description](#)
[Options](#)
[References](#)

[Quick start](#)
[Remarks and examples](#)
[Also see](#)

[Menu](#)
[Stored results](#)

[Syntax](#)
[Methods and formulas](#)

Description

`meglm` fits multilevel mixed-effects generalized linear models. `meglm` allows a variety of distributions for the response conditional on normally distributed random effects.

Quick start

Without weights

Random-effects probit regression of `y` on `x1` with random intercepts by `lev2`

```
meglm y x1 || lev2:, family(binomial) link(probit)
```

As above, but fit a logit model and report odds ratios

```
meglm y x1 || lev2:, family(binomial) or
```

Two-level gamma model of `y` with fixed and random coefficients on `x1`

```
meglm y x1 || lev2: x1, family(gamma)
```

Nested three-level random-intercept Poisson model reporting incidence-rate ratios

```
meglm y x1 || lev3: || lev2:, family(poisson) irr
```

Two-level linear regression of `y` on `x1` and `x2` with random intercepts by `lev2`, random coefficients on `x2`, and robust standard errors

```
meglm y x1 x2 || lev2: x2, vce(robust)
```

With weights

Two-level linear regression of `y` on `x` with random intercepts by `psu` for two-stage sampling with PSU-level and observation-level sampling weights `wvar2` and `wvar1`, respectively

```
meglm y x [pweight=wvar1] || psu:, pweight(wvar2)
```

Add secondary sampling stage with units identified by `ssu` having weights `wvar2` and PSU-level weights `wvar3` for a three-level random-intercept model

```
meglm y x [pw=wvar1] || psu:, pw(wvar3) || ssu:, pw(wvar2)
```

Same as above, but `svyset` data first

```
svyset psu, weight(wvar3) || ssu, weight(wvar2) || _n, weight(wvar1)  
svy: meglm y x || psu: || ssu:
```

Menu

Statistics > Multilevel mixed-effects models > Generalized linear model (GLM)

Syntax

```
meglm depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R. varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress the constant term from the fixed-effects equation
<code>exposure(<i>varname</i>_e)</code>	include $\ln(\textit{varname}_e)$ in model with coefficient constrained to 1
<code>offset(<i>varname</i>_o)</code>	include <i>varname</i> _o in model with coefficient constrained to 1
<code>asis</code>	retain perfect predictor variables

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>family</u> (<i>family</i>)	distribution of <i>depvar</i> ; default is <code>family(gaussian)</code>
<u>link</u> (<i>link</i>)	link function; default varies per family
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>eform</u>	report exponentiated fixed-effects coefficients
<u>irr</u>	report fixed-effects coefficients as incidence-rate ratios
<u>or</u>	report fixed-effects coefficients as odds ratios
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics

<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the R. notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the R. notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed(matname)</u>	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern(matname)</u>	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>family</i>	Description
<u>gaussian</u>	Gaussian (normal); the default
<u>bernoulli</u>	Bernoulli
<u>binomial</u> [# <i>varname</i>]	binomial; default number of binomial trials is 1
<u>gamma</u>	gamma
<u>nbinomial</u> [mean <u>constant</u>]	negative binomial; default dispersion is mean
<u>ordinal</u>	ordinal
<u>poisson</u>	Poisson

<i>link</i>	Description
<u>identity</u>	identity
<u>log</u>	log
<u>logit</u>	logit
<u>probit</u>	probit
<u>cloglog</u>	complementary log-log

<i>intmethod</i>	Description
<u>mvaghermite</u>	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mcaghermite</u>	mode-curvature adaptive Gauss–Hermite quadrature
<u>ghermite</u>	nonadaptive Gauss–Hermite quadrature
<u>laplace</u>	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] *bayes*: *meglm*. *vce()* and *weights* are not allowed with the *svy* prefix; see [SVY] *svy*.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified. Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box. See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

exposure(*varname_e*) specifies a variable that reflects the amount of exposure over which the *depvar* events were observed for each observation; $\ln(\text{varname}_e)$ is included in the fixed-effects portion of the model with the coefficient constrained to be 1.

offset(*varname_o*) specifies that *varname_o* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

asis forces retention of perfect predictor variables and their associated, perfectly predicted observations and may produce instabilities in maximization; see [R] **probit**.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: **independent**, **exchangeable**, **identity**, **unstructured**, **fixed**(*matname*), or **pattern**(*matname*).

covariance(**independent**) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is **covariance**(**independent**) unless a crossed random-effects model is fit, in which case the default is **covariance**(**identity**).

covariance(**exchangeable**) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(**identity**) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(**unstructured**) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p+1)/2$ unique parameters.

covariance(**fixed**(*matname*)) and **covariance**(**pattern**(*matname*)) covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a **fixed**(*matname*) covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a **pattern**(*matname*) covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if *matname*[i, j] = *matname*[k, l].

fweight(*varname*) specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, [**fw**=*fwivar*]. *varname* can be any valid Stata variable name, and you can specify **fweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable *wt1* would hold the first-level (the observation-level) frequency weights, and *wt2* would hold the second-level (the school-level) frequency weights.

iweight(*varname*) specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, [**iw**=*iwivar*]. *varname* can be any valid Stata variable name, and you can specify **iweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`family(family)` specifies the distribution of *depvar*; `family(gaussian)` is the default.

`link(link)` specifies the link function; the default is the canonical link for the `family()` specified except for the gamma and negative binomial families.

If you specify both `family()` and `link()`, not all combinations make sense. You may choose from the following combinations:

	identity	log	logit	probit	cloglog
Gaussian	D	x			
Bernoulli			D	x	x
binomial			D	x	x
gamma		D			
negative binomial		D			
ordinal			D	x	x
Poisson		D			

D denotes the default.

`constraints(constraints)`, `collinear`; see [R] estimation options.

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] *vce_option*. If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] estimation options.

`eform` reports exponentiated fixed-effects coefficients and corresponding standard errors and confidence intervals. This option may be specified either at estimation or upon replay.

`irr` reports estimated fixed-effects coefficients transformed to incidence-rate ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated or stored. `irr` may be specified either at estimation or upon replay. This option is allowed for count models only.

`or` reports estimated fixed-effects coefficients transformed to odds ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated. `or` may be specified at estimation or upon replay. This option is allowed for logistic models only.

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrntolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meglm` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with **meglm** but are not shown in the dialog box:

startvalues(svmethod) specifies how starting values are to be computed. Starting values specified in **from()** override the computed starting values.

startvalues(zero) specifies that starting values be set to 0.

startvalues(constantonly) builds on **startvalues(zero)** by fitting a constant-only model to obtain estimates of the intercept and auxiliary parameters, and it substitutes 1 for the variances of random effects.

startvalues(fixedonly[, iterate(#)]) builds on **startvalues(constantonly)** by fitting a full fixed-effects model to obtain estimates of coefficients along with intercept and auxiliary parameters, and it continues to use 1 for the variances of random effects. This is the default behavior. **iterate(#)** limits the number of iterations for fitting the fixed-effects model.

startvalues(iv[, iterate(#)]) builds on **startvalues(fixedonly)** by using instrumental-variable methods with generalized residuals to obtain variances of random effects. **iterate(#)** limits the number of iterations for fitting the instrumental-variable model.

startvalues(iterate(#)) limits the number of iterations for fitting the default model (fixed effects).

startgrid[(gridspec)] performs a grid search on variance components of random effects to improve starting values. No grid search is performed by default unless the starting values are found to be not feasible, in which case **meglm** runs **startgrid()** to perform a “minimal” search involving q^3 likelihood evaluations, where q is the number of random effects. Sometimes this resolves the problem. Usually, however, there is no problem and **startgrid()** is not run by default. There can be benefits from running **startgrid()** to get better starting values even when starting values are feasible.

startgrid() is a brute-force approach that tries various values for variances and covariances and chooses the ones that work best. You may already be using a default form of **startgrid()** without knowing it. If you see **meglm** displaying Grid node 1, Grid node 2, ... following Grid node 0 in the iteration log, that is **meglm** doing a default search because the original starting values were not feasible. The default form tries 0.1, 1, and 10 for all variances of all random effects.

startgrid(numlist) specifies values to try for variances of random effects.

startgrid(covspec) specifies the particular variances and covariances in which grid searches are to be performed. *covspec* is *name[level]* for variances and *name1[level]*name2[level]* for covariances. For example, the variance of the random intercept at level *id* is specified as *_cons[id]*, and the variance of the random slope on variable *week* at the same level is specified as *week[id]*. The residual variance for the linear mixed-effects model is specified as *e.depvar*, where *depvar* is the name of the dependent variable. The covariance between the random slope and the random intercept above is specified as *_cons[id]*week[id]*.

startgrid(numlist covspec) combines the two syntaxes. You may also specify **startgrid()** multiple times so that you can search the different ranges for different variances and covariances.

noestimate specifies that the model is not to be fit. Instead, starting values are to be shown (as modified by the above options if modifications were made), and they are to be shown using the **coeflegend** style of output.

dnumerical specifies that during optimization, the gradient vector and Hessian matrix be computed using numerical techniques instead of analytical formulas. By default, analytical formulas for computing the gradient and Hessian are used for all integration methods except **intmethod(laplace)**.

coeflegend; see [\[R\] estimation options](#).

Remarks and examples

For a general introduction to `me` commands, see [ME] [me](#). For additional examples of mixed-effects models for binary and binomial outcomes, see [ME] [melogit](#), [ME] [meprobit](#), and [ME] [mecloglog](#). For additional examples of mixed-effects models for ordinal responses, see [ME] [meologit](#) and [ME] [meoprobit](#). For additional examples of mixed-effects models for multinomial outcomes, see [SEM] [example 41g](#). For additional examples of mixed-effects models for count outcomes, see [ME] [mepoisson](#) and [ME] [menbreg](#). For additional examples of mixed-effects parametric survival models, see [ME] [mestreg](#). For additional examples of mixed-effects models for censored outcomes, see [ME] [metobit](#) and [ME] [meintreg](#).

Remarks are presented under the following headings:

- [Introduction](#)
- [Two-level models for continuous responses](#)
- [Two-level models for nonlinear responses](#)
- [Three-level models for nonlinear responses](#)
- [Crossed-effects models](#)
- [Obtaining better starting values](#)
- [Survey data](#)
- [Video example](#)

Introduction

`meglm` fits multilevel mixed-effects generalized linear models of the form

$$g\{E(\mathbf{y}|\mathbf{X}, \mathbf{u})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim F \quad (1)$$

where $\mathbf{y}$ is the $n \times 1$ vector of responses from the distributional family F , $\mathbf{X}$ is an $n \times p$ design/covariate matrix for the fixed effects $\boldsymbol{\beta}$, and $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. The $\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}$ part is called the linear predictor, and it is often denoted as $\boldsymbol{\eta}$. The linear predictor also contains the offset or exposure variable when `offset()` or `exposure()` is specified. $g(\cdot)$ is called the link function and is assumed to be invertible such that

$$E(\mathbf{y}|\mathbf{X}, \mathbf{u}) = g^{-1}(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}) = H(\boldsymbol{\eta}) = \boldsymbol{\mu}$$

For notational convenience here and throughout this manual entry, we suppress the dependence of $\mathbf{y}$ on $\mathbf{X}$. Substituting various definitions for $g(\cdot)$ and F results in a wide array of models. For instance, if $\mathbf{y}$ is distributed as Gaussian (normal) and $g(\cdot)$ is the identity function, we have

$$E(\mathbf{y}) = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim \text{normal}$$

or mixed-effects linear regression. If $g(\cdot)$ is the logit function and $\mathbf{y}$ is distributed as Bernoulli, we have

$$\text{logit}\{E(\mathbf{y})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim \text{Bernoulli}$$

or mixed-effects logistic regression. If $g(\cdot)$ is the natural log function and $\mathbf{y}$ is distributed as Poisson, we have

$$\ln\{E(\mathbf{y})\} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \quad \mathbf{y} \sim \text{Poisson}$$

or mixed-effects Poisson regression. In fact, some combinations of families and links are so common that we implemented them as separate commands in terms of `meglm`.

Command	meglm equivalent
melogit	family(bernoulli) link(logit)
meprobit	family(bernoulli) link(probit)
mecloglog	family(bernoulli) link(cloglog)
meologit	family(ordinal) link(logit)
meoprobit	family(ordinal) link(probit)
mepoisson	family(poisson) link(log)
menbreg	family(nbinomial) link(log)

When no family–link combination is specified, `meglm` defaults to a Gaussian family with an identity link. Thus `meglm` can be used to fit linear mixed-effects models; however, for those models we recommend using the more specialized `mixed`, which, in addition to `meglm` capabilities, allows for modeling of the structure of the residual errors; see [ME] [mixed](#) for details.

The random effects $\mathbf{u}$ are assumed to be distributed as multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The random effects are not directly estimated (although they may be predicted), but instead are characterized by the variance components, the elements of $\mathbf{G} = \text{Var}(\mathbf{u})$.

The general forms of the design matrices $\mathbf{X}$ and $\mathbf{Z}$ allow estimation for a broad class of generalized mixed-effects models: blocked designs, split-plot designs, growth curves, multilevel or hierarchical designs, etc. They also allow a flexible method of modeling within-cluster correlation. Subjects within the same cluster can be correlated as a result of a shared random intercept, or through a shared random slope on a covariate, or both. The general specification of variance components also provides additional flexibility—the random intercept and random slope could themselves be modeled as independent, or correlated, or independent with equal variances, and so forth.

Comprehensive treatments of mixed models are provided by, for example, [Searle, Casella, and McCulloch \(1992\)](#); [Verbeke and Molenberghs \(2000\)](#); [Raudenbush and Bryk \(2002\)](#); [Hedeker and Gibbons \(2006\)](#); [McCulloch, Searle, and Neuhaus \(2008\)](#); and [Rabe-Hesketh and Skrondal \(2012\)](#).

The key to fitting mixed models lies in estimating the variance components, and for that there exist many methods; see, for example, [Breslow and Clayton \(1993\)](#); [Lin and Breslow \(1996\)](#); [Bates and Pinheiro \(1998\)](#); and [Ng et al. \(2006\)](#). `meglm` uses maximum likelihood (ML) to estimate model parameters. The ML estimates are based on the usual application of likelihood theory, given the distributional assumptions of the model.

Returning to (1): in clustered-data situations, it is convenient not to consider all n observations at once but instead to organize the mixed model as a series of M independent groups (or clusters)

$$g\{E(\mathbf{y}_j)\} = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j \tag{2}$$

for $j = 1, \dots, M$, with cluster j consisting of n_j observations. The response $\mathbf{y}_j$ comprises the rows of $\mathbf{y}$ corresponding with the j th cluster, with $\mathbf{X}_j$ defined analogously. The random effects $\mathbf{u}_j$ can now be thought of as M realizations of a $q \times 1$ vector that is normally distributed with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The matrix $\mathbf{Z}_i$ is the $n_j \times q$ design matrix for the j th cluster random effects. Relating this to (1), note that

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{Z}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Z}_M \end{bmatrix}; \quad \mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_M \end{bmatrix}; \quad \mathbf{G} = \mathbf{I}_M \otimes \Sigma$$

where $\mathbf{I}_M$ is the $M \times M$ identity matrix and $\otimes$ is the Kronecker product.

The mixed-model formula (2) is from [Laird and Ware \(1982\)](#) and offers two key advantages. First, it makes specifications of random-effects terms easier. If the clusters are schools, you can simply specify a random effect at the school level, as opposed to thinking of what a school-level random effect would mean when all the data are considered as a whole (if it helps, think Kronecker products). Second, representing a mixed-model with (2) generalizes easily to more than one set of random effects. For example, if classes are nested within schools, then (2) can be generalized to allow random effects at both the school and the class-within-school levels.

Two-level models for continuous responses

We begin with a simple application of (2).

► Example 1: Two-level linear mixed model

Consider a longitudinal dataset, used by both [Ruppert, Wand, and Carroll \(2003\)](#) and [Diggle et al. \(2002\)](#), consisting of `weight` measurements of 48 pigs on 9 successive `weeks`. Pigs are identified by the variable `id`. Each pig experiences a linear trend in growth but overall weight measurements vary from pig to pig. Because we are not really interested in these particular 48 pigs per se, we instead treat them as a random sample from a larger population and model the between-pig variability as a random effect, or in the terminology of (2), as a random-intercept term at the pig level. We thus wish to fit the model

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_j + \epsilon_{ij}$$

for $i = 1, \dots, 9$ weeks and $j = 1, \dots, 48$ pigs. The fixed portion of the model, $\beta_0 + \beta_1 \text{week}_{ij}$, simply states that we want one overall regression line representing the population average. The random effect u_j serves to shift this regression line up or down according to each pig. Because the random effects occur at the pig level (`id`), we fit the model by typing

```
. use http://www.stata-press.com/data/r15/pig
(Longitudinal analysis of pig weights)
. meglm weight week || id:
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1251.2506
Iteration 1:   log likelihood = -1251.2506
Refining starting values:
Grid node 0:   log likelihood = -1150.6253
Fitting full model:
Iteration 0:   log likelihood = -1150.6253   (not concave)
Iteration 1:   log likelihood = -1036.1793
Iteration 2:   log likelihood = -1017.912
Iteration 3:   log likelihood = -1014.9537
Iteration 4:   log likelihood = -1014.9268
Iteration 5:   log likelihood = -1014.9268
```


Mixed-effects GLM		Number of obs	=	432
Family:	Gaussian			
Link:	identity			
Group variable:	id	Number of groups	=	48
		Obs per group:		
		min	=	9
		avg	=	9.0
		max	=	9
Integration method:	mvaghermite	Integration pts.	=	7
Log likelihood = -1014.9268		Wald chi2(1)	=	25337.48
		Prob > chi2	=	0.0000

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0390124	159.18	0.000	6.133433	6.286359
_cons	19.35561	.5974047	32.40	0.000	18.18472	20.52651
id						
var(_cons)	14.81745	3.124202			9.801687	22.39989
var(e.weight)	4.383264	.3163349			3.805112	5.049261

LR test vs. linear model: chibar2(01) = 472.65 Prob >= chibar2 = 0.0000

At this point, a guided tour of the model specification and output is in order:

1. By typing `weight week`, we specified the response, `weight`, and the fixed portion of the model in the same way that we would if we were using `regress` or any other estimation command. Our fixed effects are a coefficient on `week` and a constant term.
2. When we added `|| id:`, we specified random effects at the level identified by the group variable `id`, that is, the pig level (level two). Because we wanted only a random intercept, that is all we had to type.
3. The estimation log displays a set of iterations from optimizing the log likelihood. By default, these are Newton–Raphson iterations, but other methods are available by specifying the appropriate *maximize_options*; see [\[R\] maximize](#).
4. The header describes the model, presents a summary of the random-effects group, reports the integration method used to fit the model, and reports a Wald test against the null hypothesis that all the coefficients on the independent variables in the mean equation are 0. Here the null hypothesis is rejected at all conventional levels. You can suppress the group information with the `nogroup` or the `noheader` option, which will suppress the rest of the header as well.
5. The estimation table reports the fixed effects, followed by the random effects, followed by the overall error term.
 - a. For the fixed-effects part, we estimate $\beta_0 = 19.36$ and $\beta_1 = 6.21$.
 - b. The random-effects equation is labeled `id`, meaning that these are random effects at the `id` (pig) level. We have only one random effect at this level, the random intercept. The variance of the level-two errors, σ_u^2 , is estimated as 14.82 with standard error 3.12.
 - c. The row labeled `var(e.weight)` displays the estimated variance of the overall error term: $\sigma_\epsilon^2 = 4.38$. This is the variance of the level-one errors, that is, the residuals.
6. Finally, a likelihood-ratio test comparing the model with ordinary linear regression is provided and is highly significant for these data. See [Distribution theory for likelihood-ratio test](#) in [\[ME\] me](#) for a discussion of likelihood-ratio testing of variance components.

See [Remarks and examples](#) in `[ME] mixed` for further analysis of these data including a random-slope model and a model with an unstructured covariance structure.

Two-level models for nonlinear responses

By specifying different family–link combinations, we can fit a variety of mixed-effects models for nonlinear responses. Here we replicate one of the models from [example 2](#) of `melogit`.

➤ Example 2: Two-level logistic regression model

[Ng et al. \(2006\)](#) analyzed a subsample of data from the 1989 Bangladesh fertility survey ([Huq and Cleland 1990](#)), which polled 1,934 Bangladeshi women on their use of contraception. The women sampled were from 60 districts, identified by the variable `district`. Each district contained either urban or rural areas (variable `urban`) or both. The variable `c_use` is the binary response, with a value of 1 indicating contraceptive use. Other covariates include mean-centered `age` and three indicator variables recording number of children.

We fit a standard logistic regression model, amended to have a random intercept for each district and a random slope on the indicator variable `urban`. We fit the model by typing

Three-level models for nonlinear responses

Two-level models extend naturally to models with three or more levels with nested random effects. Here we replicate the model from [example 2](#) of [\[ME\] meologit](#).

► Example 3: Three-level ordered logistic regression model

We use the data from the Television, School, and Family Smoking Prevention and Cessation Project ([Flay et al. 1988](#); [Rabe-Hesketh and Skrondal 2012](#), chap. 11), where schools were randomly assigned into one of four groups defined by two treatment variables. Students within each school are nested in classes, and classes are nested in schools. The dependent variable is the tobacco and health knowledge (THK) scale score collapsed into four ordered categories. We regress the outcome on the treatment variables, social resistance classroom curriculum and TV intervention, and their interaction and control for the pretreatment score.

```
. use http://www.stata-press.com/data/r15/tvsfpors
. meglm thk prethk cc##tv || school: || class:, family(ordinal) link(logit)

Fitting fixed-effects model:
Iteration 0:   log likelihood =  -2212.775
Iteration 1:   log likelihood =  -2125.509
Iteration 2:   log likelihood =  -2125.1034
Iteration 3:   log likelihood =  -2125.1032

Refining starting values:
Grid node 0:   log likelihood =  -2152.1514

Fitting full model:
Iteration 0:   log likelihood =  -2152.1514   (not concave)
Iteration 1:   log likelihood =  -2125.9213   (not concave)
Iteration 2:   log likelihood =  -2120.1861
Iteration 3:   log likelihood =  -2115.6177
Iteration 4:   log likelihood =  -2114.5896
Iteration 5:   log likelihood =  -2114.5881
Iteration 6:   log likelihood =  -2114.5881

Mixed-effects GLM                Number of obs   =       1,600
Family:                          ordinal
Link:                            logit
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
school	28	18	57.1	137
class	135	1	11.9	28

Integration method: mvaghermite				Integration pts. = 7	
Log likelihood = -2114.5881				Wald chi2(4) = 124.39	
				Prob > chi2 = 0.0000	
thk	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
prethk	.4085273	.039616	10.31	0.000	.3308814 .4861731
1.cc	.8844369	.2099124	4.21	0.000	.4730161 1.295858
1.tv	.236448	.2049065	1.15	0.249	-.1651614 .6380575
cc#tv					
1 1	-.3717699	.2958887	-1.26	0.209	-.951701 .2081612
/cut1	-.0959459	.1688988			-.4269815 .2350896
/cut2	1.177478	.1704946			.8433151 1.511642
/cut3	2.383672	.1786736			2.033478 2.733865
school					
var(_cons)	.0448735	.0425387			.0069997 .2876749
school>class					
var(_cons)	.1482157	.0637521			.063792 .3443674
LR test vs. ologit model: chi2(2) = 21.03				Prob > chi2 = 0.0000	
Note: LR test is conservative and provided only for reference.					

Notes:

1. Our model now has two random-effects equations, separated by ||. The first is a random intercept (constant only) at the `school` level (level three), and the second is a random intercept at the `class` level (level two). The order in which these are specified (from left to right) is significant—`meglm` assumes that `class` is nested within `school`.
2. The information on groups is now displayed as a table, with one row for each grouping. You can suppress this table with the `nogroup` or the `noheader` option, which will suppress the rest of the header, as well.
3. The variance-component estimates are now organized and labeled according to level. The variance component for `class` is labeled `school>class` to emphasize that classes are nested within schools.

We refer you to [example 2](#) of [\[ME\] meologit](#) and [example 1](#) of [\[ME\] meologit postestimation](#) for a substantive interpretation of the results.



The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by ||. The order of nesting goes from left to right as the groups go from biggest (highest level) to smallest (lowest level).

Crossed-effects models

Not all mixed models contain nested levels of random effects. In this section, we consider a crossed-effects model, that is, a mixed-effects model in which the levels of random effects are not nested; see [\[ME\] me](#) for more information on crossed-effects models.

► Example 4: Crossed-effects logistic regression model

We use the salamander cross-breeding data from [Karim and Zeger \(1992\)](#) as analyzed in [Rabe-Hesketh and Skrondal \(2012, chap. 16.10\)](#). The salamanders come from two populations—whiteside and roughbutt—and are labeled whiteside males (`wsm`), whiteside females (`wsf`), roughbutt males (`rbm`), and roughbutt females (`rbf`). Male identifiers are recorded in the variable `male`, and female identifiers are recorded in the variable `female`. The salamanders were divided into groups such that each group contained 60 male–female pairs, with each salamander having three potential partners from the same population and three potential partners from the other population. The outcome (`y`) is coded 1 if there was a successful mating and is coded 0 otherwise; see the references for a detailed description of the mating experiment.

We fit a crossed-effects logistic regression for successful mating, where each male has the same value of his random intercept across all females, and each female has the same value of her random intercept across all males.

To fit a crossed-effects model in Stata, we use the `_all: R.varname` syntax. We treat the entire dataset as one super cluster, denoted `_all`, and we nest each gender within the super cluster by using the `R.varname` notation. `R.male` requests a random intercept for each level of `male` and imposes an identity covariance structure on the random effects; that is, the variances of the random intercepts are restricted to be equal for all male salamanders. `R.female` accomplishes the same for the female salamanders. In Stata, we type

```
. use http://www.stata-press.com/data/r15/salamander
. meglm y wsm##wsf || _all: R.male || _all: R.female, family(bernoulli)
> link(logit) or
note: crossed random-effects model specified; option intmethod(laplace)
implied
Fitting fixed-effects model:
Iteration 0:   log likelihood = -223.13998
Iteration 1:   log likelihood = -222.78752
Iteration 2:   log likelihood = -222.78735
Iteration 3:   log likelihood = -222.78735
Refining starting values:
Grid node 0:   log likelihood = -211.58149
Fitting full model:
Iteration 0:   log likelihood = -211.58149
Iteration 1:   log likelihood = -209.33737
Iteration 2:   log likelihood = -209.27869
Iteration 3:   log likelihood = -209.27674
Iteration 4:   log likelihood = -209.27659
Iteration 5:   log likelihood = -209.27659
Mixed-effects GLM                                Number of obs      =       360
Family:                               Bernoulli
Link:                                   logit
Group variable:                          _all          Number of groups   =         1
                                           Obs per group:
                                           min =             360
                                           avg =            360.0
                                           max =             360

Integration method:      laplace

Log likelihood = -209.27659                      Wald chi2(3)       =       40.62
                                           Prob > chi2        =       0.0000
```

y	Odds Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
1.wsm	.4955894	.2135158	-1.63	0.103	.2130086	1.153047
1.wsf	.0547934	.0296028	-5.38	0.000	.0190045	.1579796
wsm##wsf 1 1	36.17588	21.76004	5.97	0.000	11.12796	117.6041
_cons	2.740692	1.039892	2.66	0.008	1.302842	5.765392
_all>male var(_cons)	1.041024	.505147			.4021828	2.694624
_all>female var(_cons)	1.174373	.5503834			.4686844	2.942601

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline odds (conditional on zero random effects).
LR test vs. logistic model: chi2(2) = 27.02 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

Because we specified a crossed-effects model, meglm defaulted to the method of Laplacian approximation to calculate the likelihood; see [Computation time and the Laplacian approximation](#) in [ME] **me** for a discussion of computational complexity of mixed-effects models, and see [Methods and formulas](#) below for the formulas used by the Laplacian approximation method.

The estimates of the random intercepts suggest that the heterogeneity among the female salamanders, 1.17, is larger than the heterogeneity among the male salamanders, 1.04.

Setting both random intercepts to 0, the odds of successful mating for a roughbutt male–female pair are given by the estimate of `_cons`, 2.74. [Rabe-Hesketh and Skrondal \(2012, chap. 16.10\)](#) show how to calculate the odds ratios for the other three salamander pairings.

◀

The `R.varname` notation is equivalent to giving a list of overparameterized (none dropped) indicator variables for use in a random-effects specification. When you specify `R.varname`, `meglm` handles the calculations internally rather than creating the indicators in the data. Because the set of indicators is overparameterized, `R.varname` implies `noconstant`. You can include factor variables in the fixed-effects specification by using standard methods; see [\[U\] 11.4.3 Factor variables](#). However, random-effects equations support only the `R.varname` factor specification. For more complex factor specifications (such as interactions) in random-effects equations, use `generate` to form the variables manually.

□ Technical note

We fit the salamander model by using

```
. meglm y wsm##wsf || _all: R.male || _all: R.female ...
```

as a direct way to demonstrate the `R.` notation. However, we can technically treat female salamanders as nested within the `_all` group, yielding the equivalent way to fit the model:

```
. meglm y wsm##wsf || _all: R.male || female: ...
```

We leave it to you to verify that both produce identical results. As we note in [example 8](#) of [\[ME\] me](#), the latter specification, organized at the cluster (female) level with random-effects dimension one (a random intercept) is, in general, much more computationally efficient.

□

Obtaining better starting values

Given the flexibility of mixed-effects models, you will find that some models “fail to converge” when used with your data; see [Diagnosing convergence problems](#) in [\[ME\] me](#) for details. What we say below applies regardless of how the convergence problem revealed itself. You might have seen the error message “initial values not feasible” or some other error message, or you might have an infinite iteration log.

`meglm` provides two options to help you obtain better starting values: `startvalues()` and `startgrid()`.

`startvalues(svmethod)` allows you to specify one of four starting-value calculation methods: `zero`, `constantonly`, `fixedonly`, or `iv`. By default, `meglm` uses `startvalues(fixedonly)`. Evidently, that did not work for you. Try the other methods, starting with `startvalues(iv)`:

```
. meglm ..., ... startvalues(iv)
```

If that does not solve the problem, proceed through the others.

By the way, if you have starting values for some parameters but not others—perhaps you fit a simplified model to get them—you can combine the options `startvalues()` and `from()`:

```
. meglm ..., ...                               // simplified model
. matrix b = e(b)
. meglm ..., ... from(b) startvalues(iv)        // full model
```


The other special option `meglm` provides is `startgrid()`, which can be used with or without `startvalues()`. `startgrid()` is a brute-force approach that tries various values for variances and covariances and chooses the ones that work best.

1. You may already be using a default form of `startgrid()` without knowing it. If you see `meglm` displaying Grid node 1, Grid node 2, ... following Grid node 0 in the iteration log, that is `meglm` doing a default search because the original starting values were not feasible. The default form tries 0.1, 1, and 10 for all variances of all random effects and, if applicable, for the residual variance.
2. `startgrid(numlist)` specifies values to try for variances of random effects.
3. `startgrid(covspec)` specifies the particular variances and covariances in which grid searches are to be performed. Variances and covariances are specified in the usual way. `startgrid(_cons[id] x[id] _cons[id]*x[id])` specifies that 0.1, 1, and 10 be tried for each member of the list.
4. `startgrid(numlist covspec)` combines the two syntaxes. You can specify `startgrid()` multiple times so that you can search the different ranges for different variances and covariances.

Our advice to you is the following:

1. If you receive an iteration log and it does not contain Grid node 1, Grid node 2, ..., then specify `startgrid(.1 1 10)`. Do that whether the iteration log was infinite or ended with some other error. In this case, we know that `meglm` did not run `startgrid()` on its own because it did not report Grid node 1, Grid node 2, etc. Your problem is poor starting values, not infeasible ones.

A synonym for `startgrid(.1 1 10)` is just `startgrid` without parentheses.

Be careful, however, if you have many random effects. Specifying `startgrid()` could run a long time because it runs all possible combinations. If you have 10 random effects, that means $10^3 = 1,000$ likelihood evaluations.

If you have many random effects, rerun your difficult `meglm` command including option `iterate(#)` and look at the results. Identify the problematic variances and search across them only. Do not just look for variances going to 0. Variances getting really big can be a problem, too, and even reasonable values can be a problem. Use your knowledge and intuition about the model.

Perhaps you will try to fit your model by specifying `startgrid(.1 1 10 _cons[id] x[id] _cons[id]*x[id])`.

Values 0.1, 1, and 10 are the default. Equivalent to specifying `startgrid(.1 1 10 _cons[id] x[id] _cons[id]*x[id])` is `startgrid(_cons[id] x[id] _cons[id]*x[id])`.

Look at covariances as well as variances. If you expect a covariance to be negative but it is positive, then try negative starting values for the covariance by specifying `startgrid(-.1 -1 -10 _cons[id]*x[id])`.

Remember that you can specify `startgrid()` multiple times. Thus you might specify both `startgrid(_cons[id] x[id])` and `startgrid(-.1 -1 -10 _cons[id]*x[id])`.

2. If you receive the message “initial values not feasible”, you know that `meglm` already tried the default `startgrid()`.

The default `startgrid()` only tried the values 0.1, 1, and 10, and only tried them on the variances of random effects. You may need to try different values or try the same values on covariances or variances of errors of observed endogenous variables.

We suggest you first rerun the model causing difficulty and include the `noestimate` option. If, looking at the results, you have an idea of which variance or covariance is a problem, or if you have few variances and covariances, we would recommend running `startgrid()` first. On the other hand, if you have no idea as to which variance or covariance is the problem and you have many of them, you will be better off if you first simplify the model. After doing that, if your simplified model does not include all the variances and covariances, you can specify a combination of `from()` and `startgrid()`.

Survey data

Multilevel modeling of survey data is a little different from standard modeling in that weighted sampling can take place at multiple levels in the model, resulting in multiple sampling weights. Most survey datasets, regardless of the design, contain one overall inclusion weight for each observation in the data. This weight reflects the inverse of the probability of ultimate selection, and by “ultimate” we mean that it factors in all levels of clustered sampling, corrections for noninclusion and oversampling, poststratification, etc.

For simplicity, in what follows, assume a simple two-stage sampling design where groups are randomly sampled and then individuals within groups are sampled. Also assume that no additional weight corrections are performed; that is, sampling weights are simply the inverse of the probability of selection. The sampling weight for observation i in cluster j in our two-level sample is then $w_{ij} = 1/\pi_{ij}$, where π_{ij} is the probability that observation i, j is selected. If you were performing a standard analysis such as OLS regression with `regress`, you would simply use a variable holding w_{ij} as your `pweight` variable, and the fact that it came from two levels of sampling would not concern you. Perhaps you would type `vce(cluster groupvar)` where `groupvar` identifies the top-level groups to get standard errors that control for correlation within these groups, but you would still use only one weight variable.

Now take these same data and fit a two-level model with `meglm`. As seen in (5) in *Methods and formulas* later in this entry, it is not sufficient to use the single sampling weight w_{ij} , because weights enter the log likelihood at both the group level and the individual level. Instead, what is required for a two-level model under this sampling design is w_j , the inverse of the probability that group j is selected in the first stage, and $w_{i|j}$, the inverse of the probability that individual i from group j is selected at the second stage conditional on group j already being selected. You cannot use w_{ij} without making any assumptions about w_j .

Given the rules of conditional probability, $w_{ij} = w_j w_{i|j}$. If your dataset has only w_{ij} , then you will need to either assume equal probability sampling at the first stage ($w_j = 1$ for all j) or find some way to recover w_j from other variables in your data; see [Rabe-Hesketh and Skrondal \(2006\)](#) and the references therein for some suggestions on how to do this, but realize that there is little yet known about how well these approximations perform in practice.

What you really need to fit your two-level model are data that contain w_j in addition to either w_{ij} or $w_{i|j}$. If you have w_{ij} —that is, the unconditional inclusion weight for observation i, j —then you need to divide w_{ij} by w_j to obtain $w_{i|j}$.

➤ Example 5: Two-level logistic regression model with weights

Rabe-Hesketh and Skrondal (2006) analyzed data from the 2000 Programme for International Student Assessment (PISA) study on reading proficiency among 15-year-old American students, as performed by the Organisation for Economic Co-operation and Development (OECD). The original study was a three-stage cluster sample, where geographic areas were sampled at the first stage, schools at the second, and students at the third. Our version of the data does not contain the geographic-areas variable, so we treat this as a two-stage sample where schools are sampled at the first stage and students at the second.

```
. use http://www.stata-press.com/data/r15/pisa2000
(Programme for International Student Assessment (PISA) 2000 data)

. describe
Contains data from http://www.stata-press.com/data/r15/pisa2000.dta
  obs:                2,069                Programme for International
                                         Student Assessment (PISA) 2000
                                         data
vars:                  11                12 Jun 2016 10:08
size:                 37,242            (_dta has notes)
```

variable name	storage type	display format	value label	variable label
female	byte	%8.0g		1 if female
isei	byte	%8.0g		International socio-economic index
w_fstuw	float	%9.0g		Student-level weight
wrschbw	float	%9.0g		School-level weight
high_school	byte	%8.0g		1 if highest level by either parent is high school
college	byte	%8.0g		1 if highest level by either parent is college
one_for	byte	%8.0g		1 if one parent foreign born
both_for	byte	%8.0g		1 if both parents are foreign born
test_lang	byte	%8.0g		1 if English (the test language) is spoken at home
pass_read	byte	%8.0g		1 if passed reading proficiency threshold
id_school	int	%8.0g		School ID

Sorted by:

For student i in school j , where the variable `id_school` identifies the schools, the variable `w_fstuw` is a student-level overall inclusion weight (w_{ij} , not $w_{i|j}$) adjusted for noninclusion and nonparticipation of students, and the variable `wrschbw` is the school-level weight w_j adjusted for oversampling of schools with more minority students. The weight adjustments do not interfere with the methods prescribed above, and thus we can treat the weight variables simply as w_{ij} and w_j , respectively.

Rabe-Hesketh and Skrondal (2006) fit a two-level logistic model for passing a reading proficiency threshold. We will do the same using `meglm`, but first we must reproduce the “method 1” adjusted weight variables that were used. The “method 1” adjustment scales the first-level weights so that they sum to the effective sample size of their corresponding second-level cluster.

```
. sort id_school
. generate sqw = w_fstuw * w_fstuw
. by id_school: egen sumw = sum(w_fstuw)
. by id_school: egen sumsqw = sum(sqw)
. generate pstis1 = w_fstuw*sumw/sumsqw
```

The new variable `pst1s1` holds the adjusted first-level weights. [Rabe-Hesketh and Skrondal \(2006\)](#) also included the school mean socioeconomic index as a covariate in their analysis. We reproduce this variable using `egen`.

```
. by id_school: egen mn_isei = mean(isei)
```

Here is the fitted model:

```
. meglm pass_read female isei mn_isei high_school college test_lang one_for
> both_for [pw=pst1s1], family(bernoulli) link(logit)
> || id_school:, pweight(wnrschbw)
(output omitted)
Mixed-effects GLM                                Number of obs      =       2,069
Family:                                Bernoulli
Link:                                  logit
Group variable:                id_school          Number of groups   =        148
                                           Obs per group:
                                           min =             1
                                           avg =            14.0
                                           max =             28
Integration method: mvaghermite                Integration pts.   =         7
                                           Wald chi2(8)      =       88.30
Log pseudolikelihood = -197395.98              Prob > chi2       =       0.0000
                                           (Std. Err. adjusted for 148 clusters in id_school)
```

pass_read	Coef.	Robust Std. Err.	z	P> z	[95% Conf. Interval]	
female	.6221369	.1540088	4.04	0.000	.3202852	.9239887
isei	.018215	.0048057	3.79	0.000	.0087959	.027634
mn_isei	.0682472	.0164337	4.15	0.000	.0360378	.1004566
high_school	.1028108	.477141	0.22	0.829	-.8323683	1.03799
college	.4531688	.5053447	0.90	0.370	-.5372885	1.443626
test_lang	.6251822	.3821182	1.64	0.102	-.1237557	1.37412
one_for	-.1089314	.2739724	-0.40	0.691	-.6459075	.4280447
both_for	-.2804038	.3264681	-0.86	0.390	-.9202696	.359462
_cons	-5.877565	.954525	-6.16	0.000	-7.7484	-4.006731
id_school var(_cons)	.2955769	.1243375			.1295996	.6741201

Notes:

1. We specified the level-one weights using standard Stata weight syntax, that is, `[pw=pst1s1]`.
2. We specified the level-two weights via the `pweight(wnrschbw)` option as part of the random-effects specification for the `id_school` level. As such, it is treated as a school-level weight. Accordingly, `wnrschbw` needs to be constant within schools, and `meglm` did check for that before estimating.
3. As is the case with other estimation commands in Stata, standard errors in the presence of sampling weights are robust.
4. Robust standard errors are clustered at the top level of the model, and this will always be true unless you specify `vce(cluster clustvar)`, where `clustvar` identifies an even higher level of grouping.

➤ Example 6: Two-level logistic regression model with survey weights

meglm also supports the svy prefix (see [SVY] svy) for the linearized variance estimator. Here we refit the model from the previous example using the svy prefix after we svyset (see [SVY] svyset) the survey design variables.

```
. svyset id_school, weight(wnrschbw) || _n, weight(pst1s1)
Note: Stage 1 is sampled with replacement; further stages will be ignored for
      variance estimation.
      pweight: <none>
      VCE: linearized
      Single unit: missing
      Strata 1: <one>
      SU 1: id_school
      FPC 1: <zero>
      Weight 1: wnrschbw
      Strata 2: <one>
      SU 2: <observations>
      FPC 2: <zero>
      Weight 2: pst1s1

. svy: meglm pass_read female isei mn_isei high_school college test_lang
> one_for both_for, family(bernoulli) link(logit) || id_school:
(running meglm on estimation sample)

Survey: Mixed-effects GLM
Number of strata   =          1          Number of obs   =       2,069
Number of PSUs    =        148          Population size = 346,373.74
                                   Design df       =        147
                                   F(      8,    140)    =       10.51
                                   Prob > F          =       0.0000
```

pass_read	Linearized		t	P> t	[95% Conf. Interval]	
	Coef.	Std. Err.				
female	.6221369	.1540088	4.04	0.000	.3177796	.9264943
isei	.018215	.0048057	3.79	0.000	.0087177	.0277122
mn_isei	.0682472	.0164337	4.15	0.000	.0357704	.100724
high_school	.1028108	.477141	0.22	0.830	-.8401311	1.045753
college	.4531688	.5053447	0.90	0.371	-.5455101	1.451848
test_lang	.6251822	.3821182	1.64	0.104	-.1299725	1.380337
one_for	-.1089314	.2739724	-0.40	0.692	-.6503648	.432502
both_for	-.2804038	.3264681	-0.86	0.392	-.925581	.3647734
_cons	-5.877565	.954525	-6.16	0.000	-7.763929	-3.991201
id_school						
var(_cons)	.2955769	.1243375			.1287156	.6787495

Notes:

1. We svyset the design variables: id_school is the PSU variable, wnrschbw contains weights at the PSU level, _n specifies that the students are identified by the individual observations, and pst1s1 contains our adjusted student-level conditional weights.
2. svyset notes the lack of a finite population correction in the first stage and informs us that only the first-stage unit information will be used in the linearized variance estimator. However, the svy prefix will still pass the stage-two weights to meglm.
3. svy produces a different header, giving us an estimate of the population size, the design degrees of freedom, and the number of first-stage sampling units.

Video example

[Tour of multilevel GLMs](#)

Stored results

`meglm` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_cat)</code>	number of categories (with ordinal outcomes)
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(N_clust)</code>	number of clusters
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>gsem</code>
<code>e(cmd2)</code>	<code>meglm</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	name of marginal model
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	link
<code>e(family)</code>	family
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	offset
<code>e(binomial)</code>	binomial number of trials (with binomial models)
<code>e(dispersion)</code>	mean or constant (with negative binomial models)
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program

<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	b V
<code>e(estat_cmd)</code>	program used to implement estat
<code>e(predict)</code>	program used to implement predict
<code>e(marginsnotok)</code>	predictions disallowed by margins
<code>e(marginswtype)</code>	weight type for margins
<code>e(marginswexp)</code>	weight expression for margins
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as asbalanced
<code>e(asobserved)</code>	factor variables <code>fvset</code> as asobserved

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(cat)</code>	category values (with ordinal outcomes)
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

Methods and formulas are presented under the following headings:

[Introduction](#)
[Gauss–Hermite quadrature](#)
[Adaptive Gauss–Hermite quadrature](#)
[Laplacian approximation](#)
[Survey data](#)

Introduction

Without a loss of generality, consider a two-level generalized mixed-effects model

$$E(\mathbf{y}_j | \mathbf{X}_j, \mathbf{u}_j) = g^{-1}(\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j), \quad \mathbf{y} \sim F$$

for $j = 1, \dots, M$ clusters, with the j th cluster consisting of n_j observations, where, for the j th cluster, $\mathbf{y}_j$ is the $n_j \times 1$ response vector, $\mathbf{X}_j$ is the $n_j \times p$ matrix of fixed predictors, $\mathbf{Z}_j$ is the $n_j \times q$ matrix of random predictors, $\mathbf{u}_j$ is the $q \times 1$ vector of random effects, $\boldsymbol{\beta}$ is the $p \times 1$ vector of regression coefficients on the fixed predictors, and we use $\boldsymbol{\Sigma}$ to denote the unknown $q \times q$ variance matrix of the random effects. For simplicity, we consider a model with no auxiliary parameters.

Let $\boldsymbol{\eta}_j$ be the linear predictor, $\boldsymbol{\eta}_j = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j$, that also includes the offset or the exposure variable when `offset()` or `exposure()` is specified. Let y_{ij} and η_{ij} be the i th individual elements of $\mathbf{y}_j$ and $\boldsymbol{\eta}_j$, $i = 1, \dots, n_j$. Let $f(y_{ij} | \eta_{ij})$ be the conditional density function for the response at observation i . Because the observations are assumed to be conditionally independent, we can overload the definition of $f(\cdot)$ with vector inputs to mean

$$\log f(\mathbf{y}_j | \boldsymbol{\eta}_j) = \sum_{j=1}^{n_i} \log f(y_{ij} | \eta_{ij})$$

The random effects $\mathbf{u}_j$ are assumed to be multivariate normal with mean $\mathbf{0}$ and variance Σ . The likelihood function for cluster j is given by

$$\begin{aligned}\mathcal{L}_j(\beta, \Sigma) &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int_{\mathbb{R}^q} f(\mathbf{y}_j | \boldsymbol{\eta}_j) \exp \left(-\frac{1}{2} \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j \right) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int_{\mathbb{R}^q} \exp \left\{ \log f(\mathbf{y}_j | \boldsymbol{\eta}_j) - \frac{1}{2} \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j \right\} d\mathbf{u}_j\end{aligned}\quad (3)$$

where $\mathbb{R}$ denotes the set of values on the real line and $\mathbb{R}^q$ is the analog in q -dimensional space.

The multivariate integral in (3) is generally not tractable, so we must use numerical methods to approximate the integral. We can use a change-of-variables technique to transform this multivariate integral into a set of nested univariate integrals. Each univariate integral can then be evaluated using a form of Gaussian quadrature. `meglm` supports three types of Gauss–Hermite quadratures: mean–variance adaptive Gauss–Hermite quadrature (MVAGH), mode–curvature adaptive Gauss–Hermite quadrature (MCAGH), and nonadaptive Gauss–Hermite quadrature (GHQ). `meglm` also offers the Laplacian-approximation method, which is used as a default method for crossed mixed-effects models. Below we describe the four methods. The methods described below are based on [Skrdal and Rabe-Hesketh \(2004, chap. 6.3\)](#).

Gauss–Hermite quadrature

Let $\mathbf{u}_j = \mathbf{L}\mathbf{v}_j$, where $\mathbf{v}_j$ is a $q \times 1$ random vector whose elements are independently standard normal variables and $\mathbf{L}$ is the Cholesky decomposition of Σ , $\Sigma = \mathbf{L}\mathbf{L}'$. Then $\boldsymbol{\eta}_j = \mathbf{X}_j\beta + \mathbf{Z}_j\mathbf{L}\mathbf{v}_j$, and the likelihood in (3) becomes

$$\begin{aligned}\mathcal{L}_j(\beta, \Sigma) &= (2\pi)^{-q/2} \int_{\mathbb{R}^q} \exp \left\{ \log f(\mathbf{y}_j | \boldsymbol{\eta}_j) - \frac{1}{2} \mathbf{v}_j' \mathbf{v}_j \right\} d\mathbf{v}_j \\ &= (2\pi)^{-q/2} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \exp \left\{ \log f(\mathbf{y}_j | \boldsymbol{\eta}_j) - \frac{1}{2} \sum_{k=1}^q v_{jk}^2 \right\} dv_{j1}, \dots, dv_{jq}\end{aligned}\quad (4)$$

Consider a q -dimensional quadrature grid containing r quadrature points in each dimension. Let $\mathbf{a}_{\mathbf{k}} = (a_{k_1}, \dots, a_{k_q})'$ be a point on this grid, and let $\mathbf{w}_{\mathbf{k}} = (w_{k_1}, \dots, w_{k_q})'$ be the vector of corresponding weights. The GHQ approximation to the likelihood is

$$\begin{aligned}\mathcal{L}_j^{\text{GHQ}}(\beta, \Sigma) &= \sum_{k_1=1}^r \dots \sum_{k_q=1}^r \left[\exp \left\{ \log f(\mathbf{y}_j | \boldsymbol{\eta}_{j\mathbf{k}}) \right\} \prod_{p=1}^q w_{k_p} \right] \\ &= \sum_{k_1=1}^r \dots \sum_{k_q=1}^r \left[\exp \left\{ \sum_{i=1}^{n_j} \log f(y_{ij} | \eta_{ij\mathbf{k}}) \right\} \prod_{p=1}^q w_{k_p} \right]\end{aligned}$$

where

$$\boldsymbol{\eta}_{j\mathbf{k}} = \mathbf{X}_j\beta + \mathbf{Z}_j\mathbf{L}\mathbf{a}_{\mathbf{k}}$$

and $\eta_{ij\mathbf{k}}$ is the i th element of $\boldsymbol{\eta}_{j\mathbf{k}}$.

Adaptive Gauss–Hermite quadrature

This section sets the stage for MVAGH quadrature and MCAGH quadrature.

Let's reconsider the likelihood in (4). We use $\phi(\mathbf{v}_j)$ to denote a multivariate standard normal with mean $\mathbf{0}$ and variance $\mathbf{I}_q$, and we use $\phi(\mathbf{v}_j|\boldsymbol{\mu}_j, \boldsymbol{\Lambda}_j)$ to denote a multivariate normal with mean $\boldsymbol{\mu}_j$ and variance $\boldsymbol{\Lambda}_j$.

For fixed model parameters, the posterior density for $\mathbf{v}_j$ is proportional to

$$\phi(\mathbf{v}_j)f(\mathbf{y}_j|\boldsymbol{\eta}_j)$$

where

$$\boldsymbol{\eta}_j = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{L}\mathbf{v}_j$$

It is reasonable to assume that this posterior density can be approximated by a multivariate normal density with mean vector $\boldsymbol{\mu}_j$ and variance matrix $\boldsymbol{\Lambda}_j$. Instead of using the prior density of $\mathbf{v}_j$ as the weighting distribution in the integral, we can use our approximation for the posterior density,

$$\mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \int_{\mathbb{R}^q} \frac{f(\mathbf{y}_j|\boldsymbol{\eta}_j)\phi(\mathbf{v}_j)}{\phi(\mathbf{v}_j|\boldsymbol{\mu}_j, \boldsymbol{\Lambda}_j)} \phi(\mathbf{v}_j|\boldsymbol{\mu}_j, \boldsymbol{\Lambda}_j) d\mathbf{v}_j$$

Then the MVAGH approximation to the likelihood is

$$\mathcal{L}_j^{\text{MVAGH}}(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \sum_{k_1=1}^r \cdots \sum_{k_q=1}^r \left[\exp \{ \log f(\mathbf{y}_j|\boldsymbol{\eta}_{jk}) \} \prod_{p=1}^q w_{jk_p}^* \right]$$

where

$$\boldsymbol{\eta}_{jk} = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{L}\mathbf{a}_{jk}^*$$

and $\mathbf{a}_{jk}^*$ and $w_{jk_p}^*$ are the abscissas and weights after an orthogonalizing transformation of $\mathbf{a}_{jk}$ and w_{jk_p} , respectively, which eliminates posterior covariances between the random effects.

Estimates of $\boldsymbol{\mu}_j$ and $\boldsymbol{\Lambda}_j$ are computed using one of two different methods. The mean $\boldsymbol{\mu}_j$ and variance $\boldsymbol{\Lambda}_j$ are computed iteratively by updating the posterior moments with the MVAGH approximation, starting with a $\mathbf{0}$ mean vector and identity variance matrix. For the MCAGH approximation, $\boldsymbol{\mu}_j$ and $\boldsymbol{\Lambda}_j$ are computed by optimizing the integrand with respect to $\mathbf{v}_j$, where $\boldsymbol{\mu}_j$ is the optimal value and $\boldsymbol{\Lambda}_j$ is the curvature at $\boldsymbol{\mu}_j$.

Laplacian approximation

Consider the likelihood in (3) and denote the argument in the exponential function by

$$h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) = \log f(\mathbf{y}_j|\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - \frac{1}{2}\mathbf{u}_j'\boldsymbol{\Sigma}^{-1}\mathbf{u}_j$$

The Laplacian approximation is based on a second-order Taylor expansion of $h(\beta, \Sigma, \mathbf{u}_j)$ about the value of $\mathbf{u}_j$ that maximizes it. The first and second partial derivatives with respect to $\mathbf{u}_j$ are

$$h'(\beta, \Sigma, \mathbf{u}_j) = \frac{\partial h(\beta, \Sigma, \mathbf{u}_j)}{\partial \mathbf{u}_j} = \mathbf{Z}_j' \frac{\partial \log f(\mathbf{y}_j | \boldsymbol{\eta}_j)}{\partial \boldsymbol{\eta}_j} - \Sigma^{-1} \mathbf{u}_j$$

$$h''(\beta, \Sigma, \mathbf{u}_j) = \frac{\partial^2 h(\beta, \Sigma, \mathbf{u}_j)}{\partial \mathbf{u}_j \partial \mathbf{u}_j'} = \mathbf{Z}_j' \frac{\partial^2 \log f(\mathbf{y}_j | \boldsymbol{\eta}_j)}{\partial \boldsymbol{\eta}_j \partial \boldsymbol{\eta}_j'} \mathbf{Z}_j - \Sigma^{-1}$$

The maximizer of $h(\beta, \Sigma, \mathbf{u}_j)$ is $\hat{\mathbf{u}}_j$ such that $h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{0}$. The integral in (3) is proportional to the posterior density of $\mathbf{u}_j$ given the data, so $\hat{\mathbf{u}}_j$ is also the posterior mode.

Let

$$\hat{\mathbf{p}}_j = \mathbf{X}_j \beta + \mathbf{Z}_j \hat{\mathbf{u}}_j$$

$$\mathbf{S}_1 = \frac{\partial \log f(\mathbf{y}_j | \hat{\mathbf{p}}_j)}{\partial \hat{\mathbf{p}}_j}$$

$$\mathbf{S}_2 = \frac{\partial \mathbf{S}_1}{\partial \hat{\mathbf{p}}_j'} = \frac{\partial^2 \log f(\mathbf{y}_j | \hat{\mathbf{p}}_j)}{\partial \hat{\mathbf{p}}_j \partial \hat{\mathbf{p}}_j'}$$

$$\mathbf{H}_j = h''(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{Z}_j' \mathbf{S}_2 \mathbf{Z}_j - \Sigma^{-1}$$

then

$$\mathbf{0} = h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{Z}_j' \mathbf{S}_1 - \Sigma^{-1} \hat{\mathbf{u}}_j$$

Given the above, the second-order Taylor approximation takes the form

$$h(\beta, \Sigma, \mathbf{u}_j) \approx h(\beta, \Sigma, \hat{\mathbf{u}}_j) + \frac{1}{2} (\mathbf{u}_j - \hat{\mathbf{u}}_j)' \mathbf{H}_j (\mathbf{u}_j - \hat{\mathbf{u}}_j)$$

because the first-order derivative term is 0. The integral is approximated by

$$\int_{\mathbb{R}^q} \exp\{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j \approx (2\pi)^{q/2} |\mathbf{H}_j|^{-1/2} \exp\{h(\beta, \Sigma, \hat{\mathbf{u}}_j)\}$$

Thus the Laplacian approximated log likelihood is

$$\log \mathcal{L}_j^{\text{Lap}}(\beta, \Sigma) = -\frac{1}{2} \log |\Sigma| - \frac{1}{2} \log |\mathbf{H}_j| + h(\beta, \Sigma, \hat{\mathbf{u}}_j)$$

The log likelihood for the entire dataset is simply the sum of the contributions of the M individual clusters, namely, $\mathcal{L}(\beta, \Sigma) = \sum_{j=1}^M \mathcal{L}_j(\beta, \Sigma)$.

Maximization of $\mathcal{L}(\beta, \Sigma)$ is performed with respect to (β, σ^2) , where σ^2 is a vector comprising the unique elements of Σ . Parameter estimates are stored in $\mathbf{e}(\mathbf{b})$ as $(\hat{\beta}, \hat{\sigma}^2)$, with the corresponding variance–covariance matrix stored in $\mathbf{e}(\mathbf{V})$. In the presence of auxiliary parameters, their estimates and standard errors are included in $\mathbf{e}(\mathbf{b})$ and $\mathbf{e}(\mathbf{V})$, respectively.

Survey data

In the presence of sampling weights, following [Rabe-Hesketh and Skrondal \(2006\)](#), the weighted log pseudolikelihood for a two-level model is given as

$$\mathcal{L}(\beta, \Sigma) = \sum_{j=1}^M w_j \log \int_{-\infty}^{\infty} \exp \left\{ \sum_{i=1}^{n_j} w_{i|j} \log f(y_{ij} | \eta_{ij}) \right\} \phi(\mathbf{v}_{j1}) d\mathbf{v}_{j1} \quad (5)$$

where w_j is the inverse of the probability of selection for the j th cluster; $w_{i|j}$ is the inverse of the conditional probability of selection of individual i , given the selection of cluster j , $f(\cdot)$ is as defined previously; and $\phi(\cdot)$ is the standard multivariate normal density.

Weighted estimation is achieved through the direct application of w_j and $w_{i|j}$ into the likelihood calculations as detailed above to reflect replicated clusters for w_j and replicated observations within clusters for $w_{i|j}$. Because this estimation is based on replicated clusters and observations, frequency weights are handled similarly.

Weights are not allowed with crossed models or the Laplacian approximation.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies. <http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Breslow, N. E., and D. G. Clayton. 1993. Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* 88: 9–25.
- Cameron, A. C., and P. K. Trivedi. 2010. *Microeconometrics Using Stata*. Rev. ed. College Station, TX: Stata Press.
- Canette, I. 2011. Including covariates in crossed-effects models. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2010/12/22/including-covariates-in-crossed-effects-models/>.
- Diggle, P. J., P. J. Heagerty, K.-Y. Liang, and S. L. Zeger. 2002. *Analysis of Longitudinal Data*. 2nd ed. Oxford: Oxford University Press.
- Flay, B. R., B. R. Brannon, C. A. Johnson, W. B. Hansen, A. L. Ulene, D. A. Whitney-Saltiel, L. R. Gleason, S. Sussman, M. D. Gavin, K. M. Glowacz, D. F. Sobol, and D. C. Spiegel. 1988. The television, school, and family smoking cessation and prevention project: I. Theoretical basis and program development. *Preventive Medicine* 17: 585–607.
- Harville, D. A. 1977. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association* 72: 320–338.
- Hedeker, D., and R. D. Gibbons. 2006. *Longitudinal Data Analysis*. Hoboken, NJ: Wiley.
- Hocking, R. R. 1985. *The Analysis of Linear Models*. Monterey, CA: Brooks/Cole.
- Horton, N. J. 2011. [Stata tip 95: Estimation of error covariances in a linear model](#). *Stata Journal* 11: 145–148.
- Huq, N. M., and J. Cleland. 1990. *Bangladesh Fertility Survey 1989 (Main Report)*. National Institute of Population Research and Training.
- Karim, M. R., and S. L. Zeger. 1992. Generalized linear models with random effects; salamander mating revisited. *Biometrics* 48: 631–644.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- LaMotte, L. R. 1973. Quadratic estimation of variance components. *Biometrics* 29: 311–330.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.

- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- Ng, E. S.-W., J. R. Carpenter, H. Goldstein, and J. Rasbash. 2006. Estimation in generalised linear mixed models with binary outcomes by simulated maximum likelihood. *Statistical Modelling* 6: 23–42.
- Nichols, A. 2007. [Causal inference with observational data](#). *Stata Journal* 7: 507–541.
- Pantazis, N., and G. Touloumi. 2010. [Analyzing longitudinal data in the presence of informative drop-out: The `jmrel` command](#). *Stata Journal* 10: 226–251.
- Rabe-Hesketh, S., and A. Skrondal. 2006. Multilevel modelling of complex survey data. *Journal of the Royal Statistical Society, Series A* 169: 805–827.
- . 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Ruppert, D., M. P. Wand, and R. J. Carroll. 2003. *Semiparametric Regression*. Cambridge: Cambridge University Press.
- Schunck, R., and F. Perales. 2017. [Within- and between-cluster effects in generalized linear mixed models: A discussion of approaches and the `xthybrid` command](#). *Stata Journal* 17: 89–115.
- Searle, S. R. 1989. Obituary: Charles Roy Henderson 1911–1989. *Biometrics* 45: 1333–1335.
- Searle, S. R., G. Casella, and C. E. McCulloch. 1992. *Variance Components*. New York: Wiley.
- Skrondal, A., and S. Rabe-Hesketh. 2004. [Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models](#). Boca Raton, FL: Chapman & Hall/CRC.
- Verbeke, G., and G. Molenberghs. 2000. *Linear Mixed Models for Longitudinal Data*. New York: Springer.

Also see

- [ME] [meglm postestimation](#) — Postestimation tools for meglm
- [ME] [mixed](#) — Multilevel mixed-effects linear regression
- [ME] [menl](#) — Nonlinear mixed-effects regression
- [ME] [me](#) — Introduction to multilevel mixed-effects models
- [BAYES] [bayes: meglm](#) — Bayesian multilevel generalized linear model
- [R] [glm](#) — Generalized linear models
- [SEM] [intro 5](#) — Tour of models ([Multilevel mixed-effects models](#))
- [SVY] [svy estimation](#) — Estimation commands for survey data
- [U] [20 Estimation and postestimation commands](#)

Postestimation commands

Remarks and examples

Also see

predict

Methods and formulas

margins

References

Postestimation commands

The following postestimation command is of special interest after `meglm`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and raw, Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
<code>mu</code>	mean response; the default
<code>pr</code>	synonym for <code>mu</code> for ordinal and binary response models
<code>eta</code>	fitted linear predictor
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code>density</code>	predicted density function
<code>distribution</code>	predicted distribution function
<code>residuals</code>	raw residuals; available only with the Gaussian family
<code>pearson</code>	Pearson residuals
<code>deviance</code>	deviance residuals
<code>anscombe</code>	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
<u>outcome</u> (<i>outcome</i>)	outcome category for predicted probabilities for ordinal models
Integration	
<u>int_options</u>	integration options

`pearson`, `deviance`, `anscombe` may not be combined with `marginal`.
For ordinal outcomes, you specify one or *k* new variables in `newvarlist` with `mu` and `pr`, where *k* is the number of outcomes. If you do not specify `outcome()`, these options assume `outcome(#1)`.

<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only

<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<u>int_options</u>	integration options

<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`mu`, the default, calculates the expected value of the outcome.
`pr` calculates predicted probabilities and is a synonym for `mu`. This option is available only for ordinal and binary response models.
`eta` calculates the fitted linear prediction.

xb calculates the linear prediction $\mathbf{x}\beta$ using the estimated fixed effects (coefficients) in the model. This is equivalent to fixing all random effects in the model to their theoretical (prior) mean value of 0.

stdp calculates the standard error of the fixed-effects linear predictor $\mathbf{x}\beta$.

density calculates the density function. This prediction is computed using the current values of the observed variables, including the dependent variable.

distribution calculates the distribution function. This prediction is computed using the current values of the observed variables, including the dependent variable.

residuals calculates raw residuals, that is, responses minus the fitted values. This option is available only for the Gaussian family.

pearson calculates Pearson residuals. Pearson residuals that are large in absolute value may indicate a lack of fit.

deviance calculates deviance residuals. Deviance residuals are recommended by McCullagh and Nelder (1989) as having the best properties for examining the goodness of fit of a GLM. They are approximately normally distributed if the model is correctly specified. They can be plotted against the fitted values or against a covariate to inspect the model fit.

anscombe calculates Anscombe residuals, which are designed to closely follow a normal distribution.

conditional(ctype) and **marginal** specify how random effects are handled in computing *statistic*.

conditional() specifies that *statistic* will be computed conditional on specified or estimated random effects.

conditional(ebmeans), the default, specifies that empirical Bayes means be used as the estimates of the random effects. These estimates are also known as posterior mean estimates of the random effects.

conditional(ebmodes) specifies that empirical Bayes modes be used as the estimates of the random effects. These estimates are also known as posterior mode estimates of the random effects.

conditional(fixedonly) specifies that all random effects be set to zero, equivalent to using only the fixed portion of the model.

marginal specifies that the predicted *statistic* be computed marginally with respect to the random effects, which means that *statistic* is calculated by integrating the prediction function with respect to all the random effects over their entire support.

Although this is not the default, marginal predictions are often very useful in applied analysis. They produce what are commonly called population-averaged estimates. They are also required by **margins**.

nooffset is relevant only if you specified **offset(varname_o)** or **exposure(varname_e)** with **meglm**. It modifies the calculations made by **predict** so that they ignore the offset or the exposure variable; the linear prediction is treated as $\mathbf{X}\beta + \mathbf{Z}\mathbf{u}$ rather than $\mathbf{X}\beta + \mathbf{Z}\mathbf{u} + \text{offset}$, or $\mathbf{X}\beta + \mathbf{Z}\mathbf{u} + \ln(\text{exposure})$, whichever is relevant.

outcome(outcome) specifies the outcome for which the predicted probabilities are to be calculated. **outcome()** should contain either one value of the dependent variable or one of #1, #2, ..., with #1 meaning the first category of the dependent variable, #2 meaning the second category, etc.

reffects calculates estimates of the random effects using empirical Bayes predictions. By default, or if the **ebmeans** option is specified, empirical Bayes means are computed. With the **ebmodes** option, empirical Bayes modes are computed. You must specify q new variables, where q is the number of random-effects terms in the model. However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you.

ebmeans specifies that empirical Bayes means be used to predict the random effects.

ebmodes specifies that empirical Bayes modes be used to predict the random effects.

reses(*stub* | newvarlist*) calculates standard errors of the empirical Bayes estimators and stores the result in *newvarlist*. This option requires the **reffects** option. You must specify q new variables, where q is the number of random-effects terms in the model. However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you. The new variables will have the same storage type as the corresponding random-effects variables.

The **reffects** and **reses**() options often generate multiple new variables at once. When this occurs, the random effects (and standard errors) contained in the generated variables correspond to the order in which the variance components are listed in the output of **meglm**. The generated variables are also labeled to identify their associated random effect.

scores calculates the scores for each coefficient in **e(b)**. This option requires a new variable list of length equal to the number of columns in **e(b)**. Otherwise, use the *stub** syntax to have **predict** generate enumerated variables with prefix *stub*.

Integration

intpoints(#) specifies the number of quadrature points used to compute marginal predictions and the empirical Bayes means; the default is the value from estimation.

iterate(#) specifies the maximum number of iterations when computing statistics involving empirical Bayes estimators; the default is the value from estimation.

tolerance(#) specifies convergence tolerance when computing statistics involving empirical Bayes estimators; the default is the value from estimation.

margins

Description for margins

margins estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [, options]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
pr	synonym for mu for ordinal and binary response models
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
density	not allowed with margins
distribution	not allowed with margins
residuals	not allowed with margins
pearson	not allowed with margins
deviance	not allowed with margins
anscombe	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options conditional(ebmeans) and conditional(ebmodes) are not allowed with margins.
Option marginal is assumed where applicable if conditional(fixedonly) is not specified.

Statistics not allowed with margins are functions of stochastic quantities other than e(b).
For the full syntax, see [R] margins.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects model using meglm. For the most part, calculation centers around obtaining predictions of the random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation.

We used the binary variables, `rural` and `urban`, instead of the factor notation `i.urban` because, although supported in the fixed-effects specification of the model, such notation is not supported in random-effects specifications.

This particular model allows for district random effects that are specific to the rural and urban areas of that district and that can be interpreted as such. We can obtain predictions of posterior means of the random effects and their standard errors by typing

```
. predict re_rural re_urban, reffects reses(se_rural se_urban)
(calculating posterior means of random effects)
(using 7 quadrature points)
```

The order in which we specified the variables to be generated corresponds to the order in which the variance components are listed in `meglm` output. If in doubt, a simple `describe` will show how these newly generated variables are labeled just to be sure.

Having generated estimated random effects and standard errors, we can now list them for the first 10 districts:

```
. by district, sort: generate tag = (_n==1)
. list district re_rural se_rural re_urban se_urban if district <= 10 & tag,
> sep(0)
```

	district	re_rural	se_rural	re_urban	se_urban
1.	1	-.9523374	.316291	-.5619418	.2329456
118.	2	-.0425217	.3819309	-5.01e-18	.4896702
138.	3	8.57e-18	.6231232	.2229486	.4658747
140.	4	-.2703357	.3980832	.574464	.3962131
170.	5	.0691029	.3101591	.0074569	.4650451
209.	6	-.3939819	.2759802	.2622263	.4177785
274.	7	-.1904756	.4043461	-6.86e-18	.4896702
292.	8	.0382993	.3177392	.2250237	.4654329
329.	9	-.3715211	.3919996	.0628076	.453568
352.	10	-.5624707	.4763545	-1.90e-17	.4896702

The estimated standard errors are conditional on the values of the estimated model parameters: β and the components of Σ . Their interpretation is therefore not one of standard sample-to-sample variability but instead one that does not incorporate uncertainty in the estimated model parameters; see [Methods and formulas](#). That stated, conditional standard errors can still be used as a measure of relative precision, provided that you keep this caveat in mind.

You can also obtain predictions of posterior modes and compare them with the posterior means:

```
. predict mod_rural mod_urban, reffects ebmodes
(calculating posterior modes of random effects)

. list district re_rural mod_rural re_urban mod_urban if district <= 10 & tag,
> sep(0)
```

	district	re_rural	mod_rural	re_urban	mod_urban
1.	1	-.9523374	-.9295366	-.5619418	-.5584528
118.	2	-.0425217	-.0306312	-5.01e-18	0
138.	3	8.57e-18	0	.2229486	.2223551
140.	4	-.2703357	-.2529507	.574464	.5644512
170.	5	.0691029	.0789803	.0074569	.0077525
209.	6	-.3939819	-.3803784	.2622263	.2595116
274.	7	-.1904756	-.1737696	-6.86e-18	0
292.	8	.0382993	.0488528	.2250237	.2244676
329.	9	-.3715211	-.3540084	.0628076	.0605462
352.	10	-.5624707	-.535444	-1.90e-17	0

The two sets of predictions are fairly close.

Because not all districts contain both urban and rural areas, some of the posterior modes are 0 and some of the posterior means are practically 0. A closer examination of the data reveals that district 3 has no rural areas, and districts 2, 7, and 10 have no urban areas.

Had we imposed an unstructured covariance structure in our model, the estimated posterior modes and posterior means in the cases in question would not be exactly 0 because of the correlation between urban and rural effects. For instance, if a district has no urban areas, it can still yield a nonzero (albeit small) random-effects estimate for a nonexistent urban area because of the correlation with its rural counterpart; see [example 1](#) of [ME] [meqrlogit postestimation](#) for details.



➤ Example 2: Calculating predicted probabilities

Continuing with the model from [example 1](#), we can obtain predicted probabilities, and unless we specify the `fixedonly` option, these predictions will incorporate the estimated subject-specific random effects $\tilde{\mathbf{u}}_j$.

```
. predict pr, pr
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)
```

The predicted probabilities for observation i in cluster j are obtained by applying the inverse link function to the linear predictor, $\hat{p}_{ij} = g^{-1}(\mathbf{x}_{ij}\hat{\boldsymbol{\beta}} + \mathbf{z}_{ij}\tilde{\mathbf{u}}_j)$; see [Methods and formulas](#) for details. Because we specified the `means` option, the calculation uses posterior means for $\tilde{\mathbf{u}}_j$. You can use the `modes` option to obtain predictions based on the posterior modes for $\tilde{\tilde{\mathbf{u}}}_j$.

```
. predict prm, pr conditional(ebmodes)
(predictions based on fixed effects and posterior modes of random effects)
```

We can list the two sets of predicted probabilities together with the actual outcome for some district, let's say district 38:

```
. list c_use pr prm if district == 38
```

	c_use	pr	prm
1228.	yes	.5783408	.5780864
1229.	no	.5326623	.5324027
1230.	yes	.6411679	.6409279
1231.	yes	.5326623	.5324027
1232.	yes	.5718783	.5716228
1233.	no	.3447686	.344533
1234.	no	.4507973	.4505391
1235.	no	.1940524	.1976133
1236.	no	.2846738	.2893007
1237.	no	.1264883	.1290078
1238.	no	.206763	.2104961
1239.	no	.202459	.2061346
1240.	no	.206763	.2104961
1241.	no	.1179788	.1203522

The two sets of predicted probabilities are fairly close.

For mixed-effects models with many levels or many random effects, the calculation of the posterior means of random effects or any quantities that are based on the posterior means of random effects may take a long time. This is because we must resort to numerical integration to obtain the posterior means. In contrast, the calculation of the posterior modes of random effects is usually orders of magnitude faster because there is no numerical integration involved. For this reason, empirical modes are often used in practice as an approximation to empirical means. Note that for linear mixed-effects models, the two predictors are the same.

We can compare the observed values with the predicted values by constructing a classification table. Defining success as $\hat{y}_{ij} = 1$ if $\hat{p}_{ij} > 0.5$ and defining $\hat{y}_{ij} = 0$ otherwise, we obtain the following table.

```
. generate p_use = pr > .5
. label var p_use "Predicted outcome"
. tab2 c_use p_use, row
-> tabulation of c_use by p_use
```

Key
<i>frequency</i>
<i>row percentage</i>

Use contracept ion	Predicted outcome		Total
	0	1	
no	991	184	1,175
	84.34	15.66	100.00
yes	423	336	759
	55.73	44.27	100.00
Total	1,414	520	1,934
	73.11	26.89	100.00


The model correctly classified 84% of women who did not use contraceptives but only 44% of women who did. In the next example, we will look at some residual diagnostics.



➤ Example 3: A look at residual diagnostics

Continuing our discussion from [example 2](#), here we look at residual diagnostics. `meglm` offers three kinds of predicted residuals for nonlinear responses—Pearson, Anscombe, and deviance. Of the three, Anscombe residuals are designed to be approximately normally distributed; thus we can check for outliers by plotting Anscombe residuals against observation numbers and seeing which residuals are greater than 2 in absolute value.

```
. predict anscombe, anscombe
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)
. generate n = _n
. label var n "observation number"
. twoway (scatter anscombe n if c_use) (scatter anscombe n if !c_use),
> yline(-2 2) legend(off) text(2.5 1400 "contraceptive use")
> text(-.1 500 "no contraceptive use")
```



There seem to be some outliers among residuals that identify women who use contraceptives. We could examine the observations corresponding to the outliers, or we could try fitting a model with perhaps a different covariance structure, which we leave as an exercise.

◀

➤ Example 4: Using predicted random effects for ranking purposes

In [example 3](#) of [\[ME\] meglm](#), we estimated the effects of two treatments on the tobacco and health knowledge (THK) scale score of students in 28 schools. The dependent variable was collapsed into four ordered categories, and we fit a three-level ordinal logistic regression.


```
. use http://www.stata-press.com/data/r15/tvsfpor, clear
. meologit thk prethk i.cc##i.tv || school: || class:
Fitting fixed-effects model:
Iteration 0:  log likelihood = -2212.775
Iteration 1:  log likelihood = -2125.509
Iteration 2:  log likelihood = -2125.1034
Iteration 3:  log likelihood = -2125.1032
Refining starting values:
Grid node 0:  log likelihood = -2152.1514
Fitting full model:
(output omitted)
Mixed-effects ologit regression                                Number of obs      =          1,600
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
school	28	18	57.1	137
class	135	1	11.9	28

```
Integration method: mvaghermite                                Integration pts.   =           7
                                                                Wald chi2(4)      =        124.39
Log likelihood = -2114.5881                                     Prob > chi2       =         0.0000
```

thk	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
prethk	.4085273	.039616	10.31	0.000	.3308814	.4861731
1.cc	.8844369	.2099124	4.21	0.000	.4730161	1.295858
1.tv	.236448	.2049065	1.15	0.249	-.1651614	.6380575
cc#tv						
1 1	-.3717699	.2958887	-1.26	0.209	-.951701	.2081612
/cut1	-.0959459	.1688988			-.4269815	.2350896
/cut2	1.177478	.1704946			.8433151	1.511642
/cut3	2.383672	.1786736			2.033478	2.733865
school						
var(_cons)	.0448735	.0425387			.0069997	.2876749
school>class						
var(_cons)	.1482157	.0637521			.063792	.3443674

```
LR test vs. ologit model: chi2(2) = 21.03                      Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

Not surprisingly, the level of knowledge before the intervention is a good predictor of the level of knowledge after the intervention. The social resistance classroom curriculum is effective in raising the knowledge score, but the TV intervention and the interaction term are not.


We can rank schools by their institutional effectiveness by plotting the random effects at the school level.

```
. predict re_school re_class, reffects reses(se_school se_class)
(calculating posterior means of random effects)
(using 7 quadrature points)
. generate lower = re_school - 1.96*se_school
. generate upper = re_school + 1.96*se_school
. egen tag = tag(school)
```

```

. gsort +re_school -tag
. generate rank = sum(tag)
. generate labpos = re_school + 1.96*se_school + .1
. twoway (rcap lower upper rank) (scatter re_school rank)
> (scatter labpos rank, mlabel(school) msymbol(none) mlabpos(0)),
> xtitle(rank) ytitle(predicted posterior mean) legend(off)
> xscale(range(0 28)) xlabel(1/28) ysize(2)

```



Although there is some variability in the predicted posterior means, we cannot see significant differences among the schools in this example.

◀

Methods and formulas

Continuing the discussion in *Methods and formulas* of [ME] **meglm** and using the definitions and formulas defined there, we begin by considering the prediction of the random effects $\mathbf{u}_j$ for the j th cluster in a two-level model. Prediction of random effects in multilevel generalized linear models involves assigning values to random effects, and there are many methods for doing so; see [Skrondal and Rabe-Hesketh \(2009\)](#) and [Skrondal and Rabe-Hesketh \(2004, chap. 7\)](#) for a comprehensive review. Stata offers two methods of predicting random effects: empirical Bayes means (also known as posterior means) and empirical Bayes modes (also known as posterior modes). Below we provide more details about the two methods.

Let $\hat{\boldsymbol{\theta}}$ denote the estimated model parameters comprising $\hat{\boldsymbol{\beta}}$ and the unique elements of $\hat{\boldsymbol{\Sigma}}$. Empirical Bayes (EB) predictors of the random effects are the means or modes of the empirical posterior distribution with the parameter estimates $\boldsymbol{\theta}$ replaced with their estimates $\hat{\boldsymbol{\theta}}$. The method is called “empirical” because $\hat{\boldsymbol{\theta}}$ is treated as known. EB combines the prior information about the random effects with the likelihood to obtain the conditional posterior distribution of random effects. Using Bayes’s theorem, the empirical conditional posterior distribution of random effects for cluster j is

$$\begin{aligned}
 \omega(\mathbf{u}_j | \mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}}) &= \frac{\Pr(\mathbf{y}_j, \mathbf{u}_j | \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}})}{\Pr(\mathbf{y}_j | \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}})} \\
 &= \frac{f(\mathbf{y}_j | \mathbf{u}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\beta}}) \phi(\mathbf{u}_j; \hat{\boldsymbol{\Sigma}})}{\int f(\mathbf{y}_j | \mathbf{u}_j) \phi(\mathbf{u}_j) d\mathbf{u}_j} \\
 &= \frac{f(\mathbf{y}_j | \mathbf{u}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\beta}}) \phi(\mathbf{u}_j; \hat{\boldsymbol{\Sigma}})}{\mathcal{L}_j(\hat{\boldsymbol{\theta}})}
 \end{aligned}$$

The denominator is just the likelihood contribution of the j th cluster.

EB mean predictions of random effects, $\tilde{\mathbf{u}}$, also known as posterior means, are calculated as

$$\begin{aligned}\tilde{\mathbf{u}} &= \int_{\mathbb{R}^q} \mathbf{u}_j \omega(\mathbf{u}_j | \mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}}) d\mathbf{u}_j \\ &= \frac{\int_{\mathbb{R}^q} \mathbf{u}_j f(\mathbf{y}_j | \mathbf{u}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\beta}}) \phi(\mathbf{u}_j; \hat{\boldsymbol{\Sigma}}) d\mathbf{u}_j}{\int_{\mathbb{R}^q} f(\mathbf{y}_j | \mathbf{u}_j) \phi(\mathbf{u}_j) d\mathbf{u}_j}\end{aligned}$$

where we use the notation $\tilde{\mathbf{u}}$ rather than $\hat{\mathbf{u}}$ to distinguish predicted values from estimates. This multivariate integral is approximated by the mean–variance adaptive Gaussian quadrature; see [Methods and formulas](#) of [ME] **meglm** for details about the quadrature. If you have multiple random effects within a level or random effects across levels, the calculation involves orthogonalizing transformations using the Cholesky transformation because the random effects are no longer independent under the posterior distribution.

In a linear mixed-effects model, the posterior density is multivariate normal, and EB means are also best linear unbiased predictors (BLUPs); see [Skrondal and Rabe-Hesketh \(2004, 227\)](#). In generalized mixed-effects models, the posterior density tends to multivariate normal as cluster size increases.

EB modal predictions can be approximated by solving for the mode $\tilde{\mathbf{u}}_j$ in

$$\frac{\partial}{\partial \mathbf{u}_j} \log \omega(\tilde{\mathbf{u}}_j | \mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}}) = 0$$

Because the denominator in $\omega(\cdot)$ does not depend on $\mathbf{u}$, we can omit it from the calculation to obtain

$$\begin{aligned}& \frac{\partial}{\partial \mathbf{u}_j} \log \left\{ f(\mathbf{y}_j | \mathbf{u}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\beta}}) \phi(\mathbf{u}_j; \hat{\boldsymbol{\Sigma}}) \right\} \\ &= \frac{\partial}{\partial \mathbf{u}_j} \log f(\mathbf{y}_j | \mathbf{u}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\beta}}) + \frac{\partial}{\partial \mathbf{u}_j} \log \phi(\mathbf{u}_j; \hat{\boldsymbol{\Sigma}}) = 0\end{aligned}$$

The calculation of EB modes does not require numerical integration, and for that reason they are often used in place of EB means. As the posterior density gets closer to being multivariate normal, EB modes get closer and closer to EB means.

Just like there are many methods of assigning values to the random effects, there exist many methods of calculating standard errors of the predicted random effects; see [Skrondal and Rabe-Hesketh \(2009\)](#) for a comprehensive review.

Stata uses the posterior standard deviation as the standard error of the posterior means predictor of random effects. The EB posterior covariance matrix of the random effects is given by

$$\text{cov}(\tilde{\mathbf{u}}_j | \mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}}) = \int_{\mathbb{R}^q} (\mathbf{u}_j - \tilde{\mathbf{u}}_j)(\mathbf{u}_j - \tilde{\mathbf{u}}_j)' \omega(\mathbf{u}_j | \mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j; \hat{\boldsymbol{\theta}}) d\mathbf{u}_j$$

The posterior covariance matrix and the integrals are approximated by the mean–variance adaptive Gaussian quadrature; see [Methods and formulas](#) of [ME] **meglm** for details about the quadrature.

Conditional standard errors for the estimated posterior modes are derived from standard theory of maximum likelihood, which dictates that the asymptotic variance matrix of $\tilde{\mathbf{u}}_j$ is the negative inverse of the Hessian, $g''(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \tilde{\mathbf{u}}_j)$.

In what follows, we show formulas using the posterior means estimates of random effects $\tilde{\mathbf{u}}_j$, which are used by default or if the `means` option is specified. If the `modes` option is specified, $\tilde{\mathbf{u}}_j$ are simply replaced with the posterior modes $\hat{\mathbf{u}}_j$ in these formulas.

For any i th observation in the j th cluster in a two-level model, define the linear predictor as

$$\hat{\eta}_{ij} = \mathbf{x}_{ij}\hat{\boldsymbol{\beta}} + \mathbf{z}_{ij}\tilde{\mathbf{u}}_j$$

The linear predictor includes the offset or exposure variable if one was specified during estimation, unless the `nooffset` option is specified. If the `fixedonly` option is specified, $\hat{\eta}$ contains the linear predictor for only the fixed portion of the model, $\hat{\eta}_{ij} = \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}$.

The predicted mean, conditional on the random effects $\tilde{\mathbf{u}}_j$, is

$$\hat{\mu}_{ij} = g^{-1}(\hat{\eta}_{ij})$$

where $g^{-1}(\cdot)$ is the inverse link function for $\mu_{ij} = g^{-1}(\eta_{ij})$ defined as follows for the available links in `link(link)`:

<i>link</i>	Inverse link
<code>identity</code>	η_{ij}
<code>logit</code>	$1/\{1 + \exp(-\eta_{ij})\}$
<code>probit</code>	$\Phi(\eta_{ij})$
<code>log</code>	$\exp(\eta_{ij})$
<code>cloglog</code>	$1 - \exp\{-\exp(\eta_{ij})\}$

By default, random effects and any statistic based on them—`mu`, `fitted`, `pearson`, `deviance`, `anscombe`—are calculated using posterior means of random effects unless option `modes` is specified, in which case the calculations are based on posterior modes of random effects.

Raw residuals are calculated as the difference between the observed and fitted outcomes,

$$\nu_{ij} = y_{ij} - \hat{\mu}_{ij}$$

and are only defined for the Gaussian family.

Let r_{ij} be the number of Bernoulli trials in a binomial model, α be the conditional overdispersion parameter under the mean parameterization of the negative binomial model, and δ be the conditional overdispersion parameter under the constant parameterization of the negative binomial model.

Pearson residuals are raw residuals divided by the square root of the variance function

$$\nu_{ij}^P = \frac{\nu_{ij}}{\{V(\hat{\mu}_{ij})\}^{1/2}}$$

where $V(\hat{\mu}_{ij})$ is the family-specific variance function defined as follows for the available families in `family(family)`:

<i>family</i>	Variance function $V(\widehat{\mu}_{ij})$
bernoulli	$\widehat{\mu}_{ij}(1 - \widehat{\mu}_{ij})$
binomial	$\widehat{\mu}_{ij}(1 - \widehat{\mu}_{ij}/r_{ij})$
gamma	$\widehat{\mu}_{ij}^2$
gaussian	1
nbinomial mean	$\widehat{\mu}_{ij}(1 + \alpha\widehat{\mu}_{ij})$
nbinomial constant	$\widehat{\mu}_{ij}(1 + \delta)$
ordinal	not defined
poisson	$\widehat{\mu}_{ij}$

Deviance residuals are calculated as

$$\nu_{ij}^D = \text{sign}(\nu_{ij})\sqrt{\widehat{d}_{ij}^2}$$

where the squared deviance residual $\widehat{d}_{ij}^2$ is defined as follows:

<i>family</i>	Squared deviance $\widehat{d}_{ij}^2$
bernoulli	$-2\log(1 - \widehat{\mu}_{ij})$ if $y_{ij} = 0$ $-2\log(\widehat{\mu}_{ij})$ if $y_{ij} = 1$
binomial	$2r_{ij}\log\left(\frac{r_{ij}}{r_{ij} - \widehat{\mu}_{ij}}\right)$ if $y_{ij} = 0$ $2y_{ij}\log\left(\frac{y_{ij}}{\widehat{\mu}_{ij}}\right) + 2(r_{ij} - y_{ij})\log\left(\frac{r_{ij} - y_{ij}}{r_{ij} - \widehat{\mu}_{ij}}\right)$ if $0 < y_{ij} < r_{ij}$ $2r_{ij}\log\left(\frac{r_{ij}}{\widehat{\mu}_{ij}}\right)$ if $y_{ij} = r_{ij}$
gamma	$-2\left\{\log\left(\frac{y_{ij}}{\widehat{\mu}_{ij}}\right) - \frac{\widehat{\nu}_{ij}}{\widehat{\mu}_{ij}}\right\}$
gaussian	$\widehat{\nu}_{ij}^2$
nbinomial mean	$2\log(1 + \alpha\widehat{\mu}_{ij})\alpha$ if $y_{ij} = 0$ $2y_{ij}\log\left(\frac{y_{ij}}{\widehat{\mu}_{ij}}\right) - \frac{2}{\alpha}(1 + \alpha y_{ij})\log\left(\frac{1 + \alpha y_{ij}}{1 + \alpha\widehat{\mu}_{ij}}\right)$ otherwise
nbinomial constant	not defined
ordinal	not defined
poisson	$2\widehat{\mu}_{ij}$ if $y_{ij} = 0$ $2y_{ij}\log\left(\frac{y_{ij}}{\widehat{\mu}_{ij}}\right) - 2\widehat{\nu}_{ij}$ otherwise

Anscombe residuals, denoted ν_{ij}^A , are calculated as follows:

<i>family</i>	Anscombe residual ν_{ij}^A
bernoulli	$\frac{3 \left\{ y_{ij}^{2/3} \mathcal{H}(y_{ij}) - \hat{\mu}_{ij}^{2/3} \mathcal{H}(\hat{\mu}_{ij}) \right\}}{2 (\hat{\mu}_{ij} - \hat{\mu}_{ij}^2)^{1/6}}$
binomial	$\frac{3 \left\{ y_{ij}^{2/3} \mathcal{H}(y_{ij}/r_{ij}) - \hat{\mu}_{ij}^{2/3} \mathcal{H}(\hat{\mu}_{ij}/r_{ij}) \right\}}{2 (\hat{\mu}_{ij} - \hat{\mu}_{ij}^2/r_{ij})^{1/6}}$
gamma	$\frac{3(y_{ij}^{1/3} - \hat{\mu}_{ij}^{1/3})}{\hat{\mu}_{ij}^{1/3}}$
gaussian	ν_{ij}
nbinomial mean	$\frac{\mathcal{H}(-\alpha y_{ij}) - \mathcal{H}(-\alpha \hat{\mu}_{ij}) + 1.5(y_{ij}^{2/3} - \hat{\mu}_{ij}^{2/3})}{(\hat{\mu}_{ij} + \alpha \hat{\mu}_{ij}^2)^{1/6}}$
nbinomial constant	not defined
ordinal	not defined
poisson	$\frac{3(y_{ij}^{2/3} - \hat{\mu}_{ij}^{2/3})}{2\hat{\mu}_{ij}^{1/6}}$

where $\mathcal{H}(t)$ is a specific univariate case of the Hypergeometric2F1 function (Wolfram 2003, 780), defined here as $\mathcal{H}(t) = {}_2F_1(2/3, 1/3, 5/3, t)$.

For a discussion of the general properties of the various residuals, see Hardin and Hilbe (2012, chap. 4).

References

- Hardin, J. W., and J. M. Hilbe. 2012. *Generalized Linear Models and Extensions*. 3rd ed. College Station, TX: Stata Press.
- McCullagh, P., and J. A. Nelder. 1989. *Generalized Linear Models*. 2nd ed. London: Chapman & Hall/CRC.
- Skrondal, A., and S. Rabe-Hesketh. 2004. *Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models*. Boca Raton, FL: Chapman & Hall/CRC.
- . 2009. Prediction in multilevel generalized linear models. *Journal of the Royal Statistical Society, Series A* 172: 659–687.
- Wolfram, S. 2003. *The Mathematica Book*. 5th ed. Champaign, IL: Wolfram Media.

Also see

[ME] **meglm** — Multilevel mixed-effects generalized linear model

[U] **20 Estimation and postestimation commands**

Title

meintreg — Multilevel mixed-effects interval regression

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`meintreg` fits mixed-effects models for continuous responses where the dependent variable may be measured as point data, interval-censored data, left-censored data, or right-censored data. Thus, it is a generalization of the models fit by `metobit`. The dependent variable must be specified using two variables that indicate how it was measured.

Quick start

Two-level interval regression on `x` with random intercepts by `lev2` of the interval-measured dependent variable with lower endpoint `y_lower` and upper endpoint `y_upper`

```
meintreg y_lower y_upper x || lev2:
```

As above, but with random coefficients for `x`

```
meintreg y_lower y_upper x || lev2: x
```

Three-level random-intercept model with `lev2` nested within `lev3`

```
meintreg y_lower y_upper x || lev3: || lev2:
```

Crossed-effects model with two-way crossed random effects by factors `a` and `b`

```
meintreg y_lower y_upper x || _all:R.a || b:
```

Menu

Statistics > Multilevel mixed-effects models > Interval regression

Syntax

```
meintreg depvarlower depvarupper fe_equation [ || re_equation ] [ || re_equation ... ]  
[ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R. varname
```

The values in *depvar*_{lower} and *depvar*_{upper} should have the following form:

Type of data		<i>depvar</i> _{lower}	<i>depvar</i> _{upper}
point data	$a = [a, a]$	a	a
interval data	$[a, b]$	a	b
left-censored data	$(-\infty, b]$	.	b
right-censored data	$[a, +\infty)$	a	.
missing		.	.

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

fe_options

Description

Model

```
noconstant                suppress constant term from the fixed-effects equation  
offset(varname)           include varname in model with coefficient constrained to 1
```

re_options

Description

Model

```
covariance(vartype)       variance–covariance structure of the random effects  
noconstant                suppress constant term from the random-effects equation  
fweight(varname)          frequency weights at higher levels  
iweight(varname)          importance weights at higher levels  
pweight(varname)          sampling weights at higher levels
```


<i>options</i>	Description
Model	
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<code>mvaghermite</code>	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<code>mcaghermite</code>	mode-curvature adaptive Gauss–Hermite quadrature
<code>ghermite</code>	nonadaptive Gauss–Hermite quadrature
<code>laplace</code>	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar_{lower}, *depvar_{upper}*, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists. *bayes*, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: meintreg. *vce()* and weights are not allowed with the *svy* prefix; see [SVY] svy.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: **independent**, **exchangeable**, **identity**, **unstructured**, **fixed(*matname*)**, or **pattern(*matname*)**.

covariance(independent) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is **covariance(independent)** unless a crossed random-effects model is fit, in which case the default is **covariance(identity)**.

covariance(exchangeable) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(identity) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(unstructured) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(fixed(*matname*)) and **covariance(pattern(*matname*))** covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a **fixed(*matname*)** covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a **pattern(*matname*)** covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if *matname*[i, j] = *matname*[k, l].

fweight(*varname*) specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, [**fw**=*fwivar*]. *varname* can be any valid Stata variable name, and you can specify **fweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) frequency weights, and **wt2** would hold the second-level (the school-level) frequency weights.

iweight(*varname*) specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, [**iw**=*iwtvar*]. *varname* can be any valid Stata variable name, and you can specify **iweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) importance weights, and **wt2** would hold the second-level (the school-level) importance weights.

pweight(*varname*) specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, [**pw**=*pwivar*]. *varname* can be any valid Stata variable name, and you can specify **pweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable **wt1** would hold the first-level (the observation-level) sampling weights, and **wt2** would hold the second-level (the school-level) sampling weights.

constraints(*constraints*), **collinear**; see [\[R\] estimation options](#).

SE/Robust

vce(*vcetype*) specifies the type of standard error reported, which includes types that are derived from asymptotic theory (**oim**), that are robust to some kinds of misspecification (**robust**), and that allow for intragroup correlation (**cluster** *clustvar*); see [\[R\] vce_option](#). If **vce**(**robust**) is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

level(#), **nocnsreport**; see [\[R\] estimation options](#).

notable suppresses the estimation table, either at estimation or upon replay.

noheader suppresses the output header, either at estimation or upon replay.

nogroup suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: **nocl**, **nopvalues**, **noomitted**, **vsquish**, **noemptycells**, **baselevels**, **allbaselevels**, **nofvlabel**, **fvwrap**(#), **fvwrapon**(*style*), **cformat**(%*fmt*), **pformat**(%*fmt*), **sformat**(%*fmt*), and **nolstretch**; see [\[R\] estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meintreg` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `meintreg` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[(gridspec)]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects interval regression is regression for censored data containing both fixed effects and random effects. `meintreg` fits mixed-effects regression models that account for left-, right-, and interval-censoring. Thus, it is a generalization of the models fit by `metobit`. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

Interval data arise naturally in many contexts, such as wage data where often you know only that a person's salary is between two values. If one of the interval's endpoints is unknown, the observation is censored. Interval data and right-censored data also arise in the area of survival analysis. `meintreg`

can fit models for data where each observation represents interval data, left-censored data, right-censored data, or point data. Regardless of the type of observation, the data should be stored in the dataset as interval data; see [Syntax](#).

Regardless of the type of censoring, the expected value of the underlying dependent variable—say, y^* —is modeled using the following linear prediction:

$$E(y^*|\mathbf{X}, \mathbf{u}) = \mathbf{X}\beta + \mathbf{Z}\mathbf{u} \quad (1)$$

$\mathbf{X}$ is an $n \times p$ design/covariate matrix, analogous to the covariates you would find in a standard linear regression model, with regression coefficients (fixed effects) β . $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. This linear prediction also contains the offset when `offset()` is specified.

The columns of matrix $\mathbf{Z}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercepts model, $\mathbf{Z}$ is simply the scalar 1. The random effects $\mathbf{u}$ are realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of Σ , known as variance components. One special case of (1) places $\mathbf{Z} = \mathbf{X}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean β and variance Σ .

Below we present a short example of mixed-effects censored regression; refer to [\[ME\] me](#) and [\[ME\] meglm](#) for additional examples of random-effects models. A two-level interval regression model can also be fit using `xtintreg`; see [\[XT\] xtintreg](#). In the absence of random effects, mixed-effects censored regression reduces to standard censored regression; see [\[R\] intreg](#).

► Example 1: Three-level random-intercept model

Mastitis is a disease affecting dairy cows, consisting of an inflammatory reaction of the udder tissue. Our fictional study was performed on 10 farms using a sample of 10 dairy cows taken from each farm, and time to infection was recorded for each udder quarter for each cow in the sample. The four udder quarters are clustered within the cow, and cows are nested within farms. This is loosely based on nonfictional studies by [Goethals et al. \(2009\)](#) and [Elghafghuf et al. \(2014\)](#).

Cows were examined periodically. Thus, if a cow developed an infection, we do not know the exact day the infection occurred; we only know that it occurred between the last infection-free examination and the first examination where the infection was present. Some udder quarters did not develop an infection by the end of the study, so these observations are right-censored. We include a binary covariate, `multiparous`, which is equal to 1 for cows that have experienced more than one calving, and 0 for cows with only one calving.

To fit a log-normal model to the data, which assumes that the outcome is always positive, we take the log of our dependent variables and then use `meintreg` to apply a multilevel Gaussian model for interval- and right-censored data.

```
. use http://www.stata-press.com/data/r15/mastitis
(Fictional udder infection data)
. generate lnleft = ln(left)
(5 missing values generated)
. generate lnright = ln(right)
(82 missing values generated)
```

```
. meintreg lnleft lnright i.multiparous || farm: || cow:
```

Fitting fixed-effects model:

```
Iteration 0:  log likelihood = -912.93005
Iteration 1:  log likelihood = -901.90184
Iteration 2:  log likelihood = -901.48206
Iteration 3:  log likelihood = -901.48176
Iteration 4:  log likelihood = -901.48176
```

Refining starting values:

```
Grid node 0:  log likelihood = -897.92167
```

Fitting full model:

```
Iteration 0:  log likelihood = -897.92167 (not concave)
Iteration 1:  log likelihood = -863.2033 (not concave)
Iteration 2:  log likelihood = -857.45304 (not concave)
Iteration 3:  log likelihood = -855.18135
Iteration 4:  log likelihood = -850.84325
Iteration 5:  log likelihood = -846.31976
Iteration 6:  log likelihood = -846.24446
Iteration 7:  log likelihood = -846.24426
Iteration 8:  log likelihood = -846.24426
```

Mixed-effects interval regression

```
Number of obs      =      400
Uncensored         =         0
Left-censored      =         5
Right-censored     =        82
Interval-cens.    =       313
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
farm	10	40	40.0	40
cow	100	4	4.0	4

Integration method: mvaghermite

```
Integration pts. =      7
```

Log likelihood = -846.24426

```
Wald chi2(1) =      8.75
```

```
Prob > chi2 =    0.0031
```

	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
1.multiparous	-.5689113	.1923729	-2.96	0.003	-.9459552	-.1918674
_cons	5.644119	.1896383	29.76	0.000	5.272435	6.015803
farm						
var(_cons)	.0246795	.0258621			.0031648	.1924544
farm>cow						
var(_cons)	.2481394	.0497735			.1674773	.367651
var(e.lnleft)	.2626232	.0257671			.2166796	.3183084

LR test vs. interval model: chi2(2) = 110.47

```
Prob > chi2 = 0.0000
```

Note: LR test is conservative and provided only for reference.

We see that infection was observed in 318 udder quarters, the 5 observations that are left-censored and the 313 that are interval censored. The coefficient for *multiparous* is negative, which means that the time to infection will be about 56.9% shorter for cows that experienced multiple calvings.

The within-cow variance is 0.248, and the residual variance is 0.263, while the within-farm variance is smaller, about 0.025. A likelihood-ratio test comparing the model to an interval regression model

without random effects is provided under the table and indicates that the three-level interval regression model is preferred.



Stored results

meintreg stores the following in **e()**:

Scalars

e(N)	number of observations
e(N_unc)	number of uncensored observations
e(N_lrc)	number of left-censored observations
e(N_rc)	number of right-censored observations
e(N_int)	number of interval-censored observations
e(k)	number of parameters
e(k_dv)	number of dependent variables
e(k_eq)	number of equations in e(b)
e(k_eq_model)	number of equations in overall model test
e(k_f)	number of fixed-effects parameters
e(k_r)	number of random-effects parameters
e(k_rs)	number of variances
e(k_rc)	number of covariances
e(df_m)	model degrees of freedom
e(ll)	log likelihood
e(N_clust)	number of clusters
e(chi2)	χ^2
e(p)	significance
e(ll_c)	log likelihood, comparison model
e(chi2_c)	χ^2 , comparison test
e(df_c)	degrees of freedom, comparison test
e(p_c)	significance, comparison test
e(rank)	rank of e(V)
e(ic)	number of iterations
e(rc)	return code
e(converged)	1 if converged, 0 otherwise

Macros

e(cmd)	meglm
e(cmd2)	meintreg
e(cmdline)	command as typed
e(depvar)	names of dependent variables
e(wtype)	weight type
e(wexp)	weight expression (first-level weights)
e(fweightk)	fweight variable for <i>k</i> th highest level, if specified
e(iweightk)	iweight variable for <i>k</i> th highest level, if specified
e(pweightk)	pweight variable for <i>k</i> th highest level, if specified
e(covariates)	list of covariates
e(ivars)	grouping variables
e(model)	interval
e(title)	title in estimation output
e(link)	identity
e(family)	gaussian
e(clustvar)	name of cluster variable
e(offset)	offset
e(intmethod)	integration method
e(n_quad)	number of integration points
e(chi2type)	Wald ; type of model χ^2
e(vce)	vcetype specified in vce()
e(vcetype)	title used to label Std. Err.
e(opt)	type of optimization

e(which)	max or min; whether optimizer is to perform maximization or minimization
e(ml_method)	type of ml method
e(user)	name of likelihood-evaluator program
e(technique)	maximization technique
e(datasignature)	the checksum
e(datasignaturevars)	variables used in calculation of checksum
e(properties)	b V
e(estat_cmd)	program used to implement estat
e(predict)	program used to implement predict
e(marginsnotok)	predictions disallowed by margins
e(marginswtype)	weight type for margins
e(marginswexp)	weight expression for margins
e(footnote)	program used to implement the footnote display
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved
Matrices	
e(b)	coefficient vector
e(Cns)	constraints matrix
e(ilog)	iteration log (up to 20 iterations)
e(gradient)	gradient vector
e(N_g)	group counts
e(g_min)	group-size minimums
e(g_avg)	group-size averages
e(g_max)	group-size maximums
e(V)	variance-covariance matrix of the estimators
e(V_modelbased)	model-based variance
Functions	
e(sample)	marks estimation sample

Methods and formulas

Without a loss of generality, consider a two-level regression model

$$E(\mathbf{y}_j^* | \mathbf{X}_j, \mathbf{u}_j) = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j, \quad \mathbf{y}^* \sim \text{normal}$$

for $j = 1, \dots, M$ clusters, with the j th cluster consisting of n_j observations, where, for the j th cluster, $\mathbf{y}_j^*$ is the $n_j \times 1$ response vector, $\mathbf{X}_j$ is the $n_j \times p$ matrix of fixed predictors, $\mathbf{Z}_j$ is the $n_j \times q$ matrix of random predictors, $\mathbf{u}_j$ is the $q \times 1$ vector of random effects, and $\boldsymbol{\beta}$ is the $p \times 1$ vector of regression coefficients on the fixed predictors. The random effects, $\mathbf{u}_j$, are assumed to be multivariate normal with mean $\mathbf{0}$ and variance $\boldsymbol{\Sigma}$.

Let $\boldsymbol{\eta}_j$ be the linear predictor, $\boldsymbol{\eta}_j = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j$, that also includes the offset variable when `offset()` is specified. Let y_{ij}^* and η_{ij} be the i th individual elements of $\mathbf{y}_j^*$ and $\boldsymbol{\eta}_j$, $i = 1, \dots, n_j$.

The dependent variable, y_{ij} , is a possibly left-, right-, or interval-censored version of y_{ij}^* , and it is recorded using two variables.

The conditional density function for the response at observation ij is then,

$$f(y_{ij}^L, y_{ij}^U | \eta_{ij}) = \begin{cases} (\sqrt{2\pi}\sigma_\epsilon)^{-1} \exp^{-(y_{ij}^L - \eta_{ij})^2 / (2\sigma_\epsilon^2)} & \text{if } (y_{ij}^L, y_{ij}^U) \in C \\ \Phi\left(\frac{y_{ij}^U - \eta_{ij}}{\sigma_\epsilon}\right) & \text{if } (y_{ij}^L, y_{ij}^U) \in L \\ 1 - \Phi\left(\frac{y_{ij}^L - \eta_{ij}}{\sigma_\epsilon}\right) & \text{if } (y_{ij}^L, y_{ij}^U) \in R \\ \Phi\left(\frac{y_{ij}^U - \eta_{ij}}{\sigma_\epsilon}\right) - \Phi\left(\frac{y_{ij}^L - \eta_{ij}}{\sigma_\epsilon}\right) & \text{if } (y_{ij}^L, y_{ij}^U) \in I \end{cases}$$

where C is the set of uncensored observations ($y_{ij}^L = y_{ij}^U$ and both nonmissing), L is the set of left-censored observations (y_{ij}^L missing and y_{ij}^U nonmissing), R is the set of right-censored observations (y_{ij}^L nonmissing and y_{ij}^U missing), I is the set of interval-censored observations ($y_{ij}^L < y_{ij}^U$ and both nonmissing), and $\Phi(\cdot)$ is the cumulative normal distribution.

Because the observations are assumed to be conditionally independent, the conditional log density function for cluster j is

$$\log f(\mathbf{y}_j | \boldsymbol{\eta}_j) = \sum_{i=1}^{n_i} \log f(y_{ij} | \eta_{ij})$$

and the likelihood function for cluster j is given by

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int_{\mathbb{R}^q} f(\mathbf{y}_j | \boldsymbol{\eta}_j) \exp\left(-\frac{1}{2} \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j\right) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int_{\mathbb{R}^q} \exp\left\{\log f(\mathbf{y}_j | \boldsymbol{\eta}_j) - \frac{1}{2} \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j\right\} d\mathbf{u}_j \end{aligned} \quad (2)$$

where $\mathbb{R}$ denotes the set of values on the real line and $\mathbb{R}^q$ is the analog in q -dimensional space.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

meintreg supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

References

- Elghafghuf, A., S. Dufour, K. Reyher, I. Dohoo, and H. Stryhn. 2014. Survival analysis of clinical mastitis data using a nested frailty Cox model fit as a mixed-effects Poisson model. *Preventive Veterinary Medicine* 117: 456–468.
- Goethals, K., B. Ampe, D. Berkvens, H. Laevens, P. Janssen, and L. Duchateau. 2009. Modeling interval-censored, clustered cow udder quarter infection times through the shared gamma frailty model. *Journal of Agricultural, Biological, and Environmental Statistics* 14: 1–14.

Also see

- [ME] **meintreg postestimation** — Postestimation tools for **meintreg**
- [ME] **metobit** — Multilevel mixed-effects tobit regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: meintreg** — Bayesian multilevel interval regression
- [R] **intreg** — Interval regression
- [SEM] **intro 5** — Tour of models ([Multilevel mixed-effects models](#))
- [ST] **stintreg** — Parametric models for interval-censored survival-time data
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtintreg** — Random-effects interval-data regression models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation commands are of special interest after `meintreg`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat icc</code>	estimate intraclass correlations
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as linear predictions, standard errors, probabilities, and expected values.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
eta	fitted linear predictor; the default
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
pr(<i>a</i> , <i>b</i>)	$\Pr(a < y < b)$
e(<i>a</i> , <i>b</i>)	$E(y \mid a < y < b)$
ystar(<i>a</i> , <i>b</i>)	$E(y^*), y^* = \max\{a, \min(y, b)\}$

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

where *a* and *b* may be numbers or variables; *a* missing (*a* ≥ .) means $-\infty$, and *b* missing (*b* ≥ .) means $+\infty$; see [U] [12.2.1 Missing values](#).

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
Integration	
<i>int_options</i>	integration options
<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only
<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<i>int_options</i>	integration options
<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`eta`, the default, calculates the fitted linear prediction.

`pr(a,b)` calculates estimates of $\Pr(a < y < b)$, which is the probability that y would be observed in the interval (a, b) .

a and b may be specified as numbers or variable names; lb and ub are variable names;

`pr(20,30)` calculates $\Pr(20 < y < 30)$;

`pr(lb,ub)` calculates $\Pr(lb < y < ub)$; and

`pr(20,ub)` calculates $\Pr(20 < y < ub)$.

a missing ($a \geq .$) means $-\infty$; `pr(.,30)` calculates $\Pr(-\infty < y < 30)$;

`pr(lb,30)` calculates $\Pr(-\infty < y < 30)$ in observations for which $lb \geq .$

(and calculates $\Pr(lb < y < 30)$ elsewhere).

b missing ($b \geq .$) means $+\infty$; `pr(20,.)` calculates $\Pr(+\infty > y > 20)$;
`pr(20,ub)` calculates $\Pr(+\infty > y > 20)$ in observations for which $ub \geq .$
(and calculates $\Pr(20 < y < ub)$ elsewhere).
`e(a,b)` calculates estimates of $E(y \mid a < y < b)$, which is the expected value of y conditional on y being in the interval (a,b) , meaning that y is truncated. a and b are specified as they are for `pr()`.
`ystar(a,b)` calculates estimates of $E(y^*)$, where $y^* = a$ if $y \leq a$, $y^* = b$ if $y \geq b$, and $y^* = y$ otherwise, meaning that y^* is the censored version of y . a and b are specified as they are for `pr()`.
`xb`, `stdp`, `scores`, `conditional()`, `marginal`, and `nooffset`; see [ME] **meglm postestimation**.
`reffects`, `ebmeans`, `ebmodes`, and `reses()`; see [ME] **meglm postestimation**.

Integration

`intpoints()`, `iterate()`, `tolerance()`; see [ME] **meglm postestimation**.

margins

Description for margins

`margins` estimates margins of response for linear predictions, probabilities, and expected values.

Menu for margins

Statistics > Postestimation

Syntax for margins

`margins` [*marginlist*] [*, options*]

`margins` [*marginlist*] , predict(*statistic ...*) [predict(*statistic ...*) ...] [*options*]

<i>statistic</i>	Description
<code>eta</code>	fitted linear predictor; the default
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>pr(a,b)</code>	$\Pr(a < y < b)$
<code>e(a,b)</code>	$E(y \mid a < y < b)$
<code>ystar(a,b)</code>	$E(y^*)$, $y^* = \max\{a, \min(y, b)\}$
<code>stdp</code>	not allowed with <code>margins</code>

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.
For the full syntax, see [R] **margins**.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects interval regression model with `meintreg`.

The `predict` command allows us to compute marginal and conditional predictions. Unless stated differently, we use the word “conditional” to mean “conditional on the empirical Bayes predictions of the random effects”. The default prediction is the linear prediction, `eta`, which is the expected value of the unobserved censored variable. Predictions of expected values for censored and truncated versions of the response are also available.

► Example 1: Obtaining conditional and marginal probabilities

In [example 1](#) of [\[ME\]](#) `meintreg`, we fit a three-level mixed-effects interval regression to model log time to udder tissue infection in dairy cows.

```
. use http://www.stata-press.com/data/r15/mastitis
(Simulated data on udder infection of dairy cows)

. generate lnleft = ln(left)
(5 missing values generated)

. generate lnright = ln(right)
(82 missing values generated)

. meintreg lnleft lnright i.multiparous || farm: || cow:
(output omitted)
```


Let’s assume that we want to predict the probability of infection within the first 90 days. Because our dependent variable is $\log(y)$, we need to compute

$$\Pr(0 < y < 90) = \Pr\{-\infty < \log(y) < \log(90)\}$$

We can use the `pr()` option for `predict` to compute the probability that our dependent variable lies in the interval $[-\infty, \log(90)]$.

We first compute the probability conditional on the random effects. Because the lower level on which we are conditioning on is cow, and we have only cow-level covariates, these predictions will be constant within cow. We can see that all the predicted probabilities for farm 3 are below 0.21, while the probabilities for farms 2 and 6 reach above 0.70 in some cases.

```
. predict pr_cond, pr(.,log(90))
(predictions based on fixed effects and posterior means of random effects)
. twoway scatter pr_cond farm, ylabel(0(.1).8) xlabel(1(1)10)
```



Now, we compute the marginal probabilities of infection within the first 90 days.

```
. predict pr_marg, pr(.,log(90)) marginal
. tabulate pr_marg multiparous
```

Marginal Pr(y<log(90))	=1 if cows experienced more than one calving, 0 otherwise		Total
	0	1	
.0589298	40	0	40
.2158333	0	360	360
Total	40	360	400

Marginal predictions depend only on the covariate pattern (including covariates in the random-effects part, if present in the model). Because we included only a binary covariate in the model, there are only two predicted values, one for each value of the covariate. We see that the probability of developing an infection in the first 90 days is higher for multiparous cows.

Alternatively, we can use `margins` to calculate the marginal probabilities. One advantage of using `margins` is that we can obtain confidence intervals for the probabilities and the difference between them.

```
. margins multiparous, predict(pr(.,log(90)))
Adjusted predictions          Number of obs      =          400
Model VCE      : OIM
Expression    : Pr(y<log(90)), predict(pr(.,log(90)))
```

	Delta-method		z	P> z	[95% Conf. Interval]	
	Margin	Std. Err.				
multiparous						
0	.0589298	.0305541	1.93	0.054	-.0009551	.1188147
1	.2158333	.0314158	6.87	0.000	.1542595	.2774071

```
. margins, dydx(multiparous) predict(pr(.,log(90)))
Conditional marginal effects          Number of obs      =          400
Model VCE      : OIM
Expression    : Pr(y<log(90)), predict(pr(.,log(90)))
dy/dx w.r.t.  : 1.multiparous
```

	Delta-method		z	P> z	[95% Conf. Interval]	
	dy/dx	Std. Err.				
1.multipar~s	.1569036	.0396889	3.95	0.000	.0791148	.2346923

Note: dy/dx for factor levels is the discrete change from the base level.

The default option for `predict, eta`, computes the fitted linear prediction; we can use this option to perform predictions for the uncensored unobserved response. We compute the conditional and marginal predictions for the log time to infection.

```
. predict eta_cond
(option eta assumed)
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)
. predict eta_marg, marginal
(option eta assumed)
. sort cow
. list cow multiparous eta_cond eta_marg in 1/8, sepby(cow)
```

	cow	multip~s	eta_cond	eta_marg
1.	1	0	5.486386	5.644119
2.	1	0	5.486386	5.644119
3.	1	0	5.486386	5.644119
4.	1	0	5.486386	5.644119
5.	2	1	5.101668	5.075207
6.	2	1	5.101668	5.075207
7.	2	1	5.101668	5.075207
8.	2	1	5.101668	5.075207


Comparing the conditional and marginal predictions, we see that the predicted log time to infection for the first cow is slightly shorter than the one expected for a cow with this covariate pattern, and the log time to infection for the second cow is slightly longer.

➤ **Example 2: Calculating transformed predictions**

Because our dependent variable is log transformed, we might want to compute predictions on the original scale. To do that, we need to obtain predictions for the exponentiated dependent variable.

This exercise is helpful to understand the distribution of the different statistics. If we want to predict the individual conditional time to infection, we need to obtain the conditional mean for $\exp(y)$. Because the conditional distribution of $\exp(y)$ is lognormal with location parameter equal to $\hat{\eta}$ and scale parameter equal to σ_{ϵ} (residual variance), then its (conditional) expected value is equal to $\exp(\hat{\eta} + \sigma_{\epsilon}^2/2)$. Here we calculate the conditional time to infection and plot kernel densities for multiparous and uniparous cows.

```
. generate time_cond = exp(eta_cond + _b[/var(e.lnleft)]/2)
. kdensity time_cond if multiparous == 0, xlabel(0(200)800) name(gr1)
. kdensity time_cond if multiparous == 1, xlabel(0(200)800) name(gr2)
. graph combine gr1 gr2
```



The density estimator of the time to infection shows that multiparous cows tend to have shorter times to infection than uniparous cows.

The marginal distribution of y is lognormal with location parameter $\mathbf{x}\beta$ and the scale parameter equal to the marginal variance; see [Methods and formulas](#) of [\[ME\] metobit postestimation](#) for the description of the marginal variance. Thus the marginal expected value of the time to infection is calculated as

```
. predict xb, xb
. generate time_marg = exp( xb + (_b[/var(_cons[farm])]) +
> _b[/var(_cons[farm>cow])]) + _b[/var(e.lnleft)]/2)
. tabulate time_marg multiparous
```

time_marg	=1 if cows experienced more than one calving, 0 otherwise		Total
	0	1	
209.1242	0	360	360
369.3851	40	0	40
Total	40	360	400

As before, we see that the unconditional expected value for the time to infection is shorter for multiparous cows.



Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [ME] [metobit postestimation](#).

Also see

[ME] [meintreg](#) — Multilevel mixed-effects interval regression

[U] [20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`melogit` fits mixed-effects models for binary and binomial responses. The conditional distribution of the response given the random effects is assumed to be Bernoulli, with success probability determined by the logistic cumulative distribution function.

`melogit` performs optimization using the original metric of variance components. When variance components are near the boundary of the parameter space, you may consider using the `meqrlogit` command, which provides alternative parameterizations of variance components; see [\[ME\]](#) `meqrlogit`.

Quick start

Without weights

Two-level logistic regression of `y` on `x` with random intercepts by `lev2`

```
melogit y x || lev2:
```

Mixed-effects model adding random coefficients for `x`

```
melogit y x || lev2: x
```

As above, but allow the random effects to be correlated

```
melogit y x || lev2: x, covariance(unstructured)
```

Three-level random-intercept model of `y` on `x` with `lev2` nested within `lev3`

```
melogit y x || lev3: || lev2:
```

Crossed-effects model of `y` on `x` with two-way crossed random effects by factors `a` and `b`

```
melogit y x || _all:R.a || b:
```

With weights

Two-level logistic regression of `y` on `x` with random intercepts by `lev2` and observation-level frequency weights `wvar1`

```
melogit y x [fweight=wvar1] || lev2:
```

Two-level random-intercept model from a two-stage sampling design with PSUs identified by `psu` using PSU-level and observation-level sampling weights `wvar2` and `wvar1`, respectively

```
melogit y x [pweight=wvar1] || psu:, pweight(wvar2)
```

Add secondary sampling stage with units identified by `ssu` having weights `wvar2` and PSU-level weights `wvar3` for a three-level random-intercept model

```
melogit y x [pw=wvar1] || psu:, pw(wvar3) || ssu:, pw(wvar2)
```

Same as above, but `svyset` data first

```
svyset psu, weight(wvar3) || ssu, weight(wvar2) || _n, weight(wvar1)
svy: melogit y x || psu: || ssu:
```

Menu

Statistics > Multilevel mixed-effects models > Logistic regression

Syntax

```
melogit depvar fe_equation [|| re_equation] [|| re_equation ...] [, options]
```

where the syntax of *fe_equation* is

```
[indepvars] [if] [in] [weight] [, fe_options]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [, re_options]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1
<code>asis</code>	retain perfect predictor variables

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>binomial</u> (<i>varname</i> #)	set binomial trials if data are in binomial form
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
or	report fixed-effects coefficients as odds ratios
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<i>maximize_options</i>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<code>mvaghermite</code>	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<code>mcaghermite</code>	mode-curvature adaptive Gauss–Hermite quadrature
<code>ghermite</code>	nonadaptive Gauss–Hermite quadrature
<code>laplace</code>	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] *bayes*: *melogit*.

vce() and weights are not allowed with the *svy* prefix; see [SVY] *svy*.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 *weight*. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

offset(varname) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

asis forces retention of perfect predictor variables and their associated, perfectly predicted observations and may produce instabilities in maximization; see [R] *probit*.

covariance(vartype) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, *unstructured*, *fixed(matname)*, or *pattern(matname)*.

covariance(independent) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is *covariance(independent)* unless a crossed random-effects model is fit, in which case the default is *covariance(identity)*.

covariance(exchangeable) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(identity) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(unstructured) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(fixed(matname)) and *covariance(pattern(matname))* covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted.

In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of `matname`. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if `matname[i, j] = matname[k, l]`.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. `varname` can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`binomial(varname | #)` specifies that the data are in binomial form; that is, `depvar` records the number of successes from a series of binomial trials. This number of trials is given either as `varname`, which allows this number to vary over the observations, or as the constant `#`. If `binomial()` is not specified (the default), `depvar` is treated as Bernoulli, with any nonzero, nonmissing values indicating positive responses.

`constraints(constraints)`, `collinear`; see [\[R\] estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [\[R\] vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] [estimation options](#).

or reports estimated fixed-effects coefficients transformed to odds ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated. or may be specified either at estimation or upon replay.

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrntolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `melogit` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `melogit` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[ (gridspec) ]`, `noestimate`, and `dnumerical`; see [ME] **meglm**.

`coeflegend`; see [R] **estimation options**.

Remarks and examples

For a general introduction to `me` commands, see [ME] **me**.

`melogit` is a convenience command for `meglm` with a `logit` link and a `bernoulli` or `binomial` family; see [ME] **meglm**.

Remarks are presented under the following headings:

Introduction
Two-level models
Other covariance structures
Three-level models
Crossed-effects models

Introduction

Mixed-effects logistic regression is logistic regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

`melogit` allows for many levels of random effects. However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = 1 | \mathbf{x}_{ij}, \mathbf{u}_j) = H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are the binary-valued y_{ij} , and we follow the standard Stata convention of treating $y_{ij} = 1$ if `depvarij` $\neq 0$ and treating $y_{ij} = 0$ otherwise. The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard logistic regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Finally, because this is logistic regression, $H(\cdot)$ is the logistic cumulative distribution function, which maps the linear predictor to the probability of a success ($y_{ij} = 1$), with $H(v) = \exp(v) / \{1 + \exp(v)\}$.

Model (1) may also be stated in terms of a latent linear response, where only $y_{ij} = I(y_{ij}^* > 0)$ is observed for the latent

$$y_{ij}^* = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

The errors ϵ_{ij} are distributed as logistic with mean 0 and variance $\pi^2/3$ and are independent of $\mathbf{u}_j$.

A two-level logistic model can also be fit using `xtlogit` with the `re` option; see [XT] [xtlogit](#). In the absence of random effects, mixed-effects logistic regression reduces to standard logistic regression; see [R] [logit](#).

Two-level models

► Example 1: Two-level random-intercept model

Ng et al. (2006) analyze a subsample of data from the 1989 Bangladesh fertility survey (Huq and Cleland 1990), which polled 1,934 Bangladeshi women on their use of contraception.

```
. use http://www.stata-press.com/data/r15/bangladesh
(Bangladesh Fertility Survey, 1989)

. describe
Contains data from http://www.stata-press.com/data/r15/bangladesh.dta
obs:           1,934           Bangladesh Fertility Survey,
                               1989
vars:           7             28 May 2016 20:27
size:          19,340         (_dta has notes)
```

variable name	storage type	display format	value label	variable label
district	byte	%9.0g		District
c_use	byte	%9.0g	yesno	Use contraception
urban	byte	%9.0g	urban	Urban or rural
age	float	%6.2f		Age, mean centered
child1	byte	%9.0g		1 child
child2	byte	%9.0g		2 children
child3	byte	%9.0g		3 or more children

Sorted by: district

The women sampled were from 60 districts, identified by the variable `district`. Each district contained either urban or rural areas (variable `urban`) or both. The variable `c_use` is the binary response, with a value of 1 indicating contraceptive use. Other covariates include mean-centered `age` and three indicator variables recording number of children.

Consider a standard logistic regression model, amended to have random effects for each district. Defining $\pi_{ij} = \Pr(c\_use_{ij} = 1)$, we have

$$\text{logit}(\pi_{ij}) = \beta_0 + \beta_1 \text{urban}_{ij} + \beta_2 \text{age}_{ij} + \beta_3 \text{child1}_{ij} + \beta_4 \text{child2}_{ij} + \beta_5 \text{child3}_{ij} + u_j \quad (2)$$

for $j = 1, \dots, 60$ districts, with $i = 1, \dots, n_j$ women in district j .

```
. melogit c_use urban age child* || district:
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1229.5485
Iteration 1:   log likelihood = -1228.5268
Iteration 2:   log likelihood = -1228.5263
Iteration 3:   log likelihood = -1228.5263
Refining starting values:
Grid node 0:   log likelihood = -1219.2681
Fitting full model:
Iteration 0:   log likelihood = -1219.2681 (not concave)
Iteration 1:   log likelihood = -1207.5978
Iteration 2:   log likelihood = -1206.8428
Iteration 3:   log likelihood = -1206.8322
Iteration 4:   log likelihood = -1206.8322
Mixed-effects logistic regression
Group variable:      district
Number of obs       =      1,934
Number of groups    =        60
Obs per group:
    min =           2
    avg =          32.2
    max =          118
Integration method: mvaghermite
Integration pts.    =           7
Wald chi2(5)       =      109.60
Prob > chi2        =       0.0000
Log likelihood = -1206.8322
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	.7322765	.1194857	6.13	0.000	.4980888	.9664641
age	-.0264981	.0078916	-3.36	0.001	-.0419654	-.0110309
child1	1.116001	.1580921	7.06	0.000	.8061465	1.425856
child2	1.365895	.1746691	7.82	0.000	1.02355	1.70824
child3	1.344031	.1796549	7.48	0.000	.9919139	1.696148
_cons	-1.68929	.1477591	-11.43	0.000	-1.978892	-1.399687
district var(_cons)	.215618	.0733222			.1107208	.4198954

```
LR test vs. logistic model: chibar2(01) = 43.39      Prob >= chibar2 = 0.0000
```

The estimation table reports the fixed effects and the estimated variance components. The fixed effects can be interpreted just as you would the output from `logit`. You can also specify the `or` option at estimation or on replay to display the fixed effects as odds ratios instead. If you did display results as odds ratios, you would find urban women to have roughly double the odds of using contraception as that of their rural counterparts. Having any number of children will increase the odds from three- to fourfold when compared with the base category of no children. Contraceptive use also decreases with age.

Underneath the fixed effect, the table shows the estimated variance components. The random-effects equation is labeled `district`, meaning that these are random effects at the `district` level. Because we have only one random effect at this level, the table shows only one variance component. The estimate of σ_u^2 is 0.22 with standard error 0.07.

A likelihood-ratio test comparing the model with ordinary logistic regression is provided and is highly significant for these data.

We now store our estimates for later use.

```
. estimates store r_int
```

In what follows, we will be extending (2), focusing on the variable `urban`. Before we begin, to keep things short we restate (2) as

$$\text{logit}(\pi_{ij}) = \beta_0 + \beta_1 \text{urban}_{ij} + \mathcal{F}_{ij} + u_j$$

where $\mathcal{F}_{ij}$ is merely shorthand for the portion of the fixed-effects specification having to do with age and children.

➤ Example 2: Two-level random-slope model

Extending (2) to allow for a random slope on the indicator variable `urban` yields the model

$$\text{logit}(\pi_{ij}) = \beta_0 + \beta_1 \text{urban}_{ij} + \mathcal{F}_{ij} + u_j + v_j \text{urban}_{ij} \quad (3)$$

which we can fit by typing

```
. melogit c_use urban age child* || district: urban
(output omitted)
. estimates store r_urban
```

Extending the model was as simple as adding `urban` to the random-effects specification so that the model now includes a random intercept and a random coefficient on `urban`. We dispense with the output because, although this is an improvement over the random-intercept model (2),

```
. lrtest r_int r_urban
Likelihood-ratio test                LR chi2(1) =      3.66
(Assumption: r_int nested in r_urban) Prob > chi2 =    0.0558

Note: The reported degrees of freedom assumes the null hypothesis is not on
the boundary of the parameter space.  If this is not true, then the
reported test is conservative.
```

we find the default covariance structure for (u_j, v_j) , `covariance(independent)`,

$$\Sigma = \text{Var} \begin{bmatrix} u_j \\ v_j \end{bmatrix} = \begin{bmatrix} \sigma_u^2 & 0 \\ 0 & \sigma_v^2 \end{bmatrix}$$

to be inadequate. We state that the random-coefficient model is an “improvement” over the random-intercept model because the null hypothesis of the likelihood-ratio comparison test ($H_0: \sigma_v^2 = 0$) is on the boundary of the parameter test. This makes the reported p -value, 5.6%, an upper bound on the actual p -value, which is actually half of that; see [Distribution theory for likelihood-ratio test in \[ME\] me](#).

We see below that we can reject this model in favor of one that allows correlation between u_j and v_j .

```
. melogit c_use urban age child* || district: urban, covariance(unstructured)
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1229.5485
Iteration 1:   log likelihood = -1228.5268
Iteration 2:   log likelihood = -1228.5263
Iteration 3:   log likelihood = -1228.5263
Refining starting values:
Grid node 0:   log likelihood = -1215.8592
Fitting full model:
Iteration 0:   log likelihood = -1215.8592 (not concave)
Iteration 1:   log likelihood = -1201.0652
Iteration 2:   log likelihood = -1199.6394
Iteration 3:   log likelihood = -1199.3157
Iteration 4:   log likelihood = -1199.315
Iteration 5:   log likelihood = -1199.315
Mixed-effects logistic regression
Group variable:      district
Number of obs       =      1,934
Number of groups    =       60
Obs per group:
    min =           2
    avg =          32.2
    max =          118
Integration method: mvaghermite
Integration pts.    =       7
Wald chi2(5)       =      97.50
Prob > chi2        =     0.0000
Log likelihood = -1199.315
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	.8157875	.1715519	4.76	0.000	.4795519	1.152023
age	-.026415	.008023	-3.29	0.001	-.0421398	-.0106902
child1	1.13252	.1603285	7.06	0.000	.818282	1.446758
child2	1.357739	.1770522	7.67	0.000	1.010723	1.704755
child3	1.353827	.1828801	7.40	0.000	.9953882	1.712265
_cons	-1.71165	.1605618	-10.66	0.000	-2.026345	-1.396954
district						
var(urban)	.6663237	.3224689			.258074	1.720387
var(_cons)	.3897448	.1292463			.203473	.7465413
district						
cov(urban, _cons)	-.4058861	.1755414	-2.31	0.021	-.7499408	-.0618313

```
LR test vs. logistic model: chi2(3) = 58.42
Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
. estimates store r_urban_corr
. lrtest r_urban r_urban_corr
Likelihood-ratio test
(Assumption: r_urban nested in r_urban_corr)
LR chi2(1) = 11.38
Prob > chi2 = 0.0007
```

By specifying `covariance(unstructured)` above, we told `melogit` to allow correlation between random effects at the district level; that is,

$$\Sigma = \text{Var} \begin{bmatrix} u_j \\ v_j \end{bmatrix} = \begin{bmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{bmatrix}$$

► Example 3: Alternative parameterization of random slopes

The purpose of introducing a random coefficient on the binary variable `urban` in (3) was to allow for separate random effects, within each district, for the urban and rural areas of that district. Hence, if we have the binary variable `rural` in our data such that $\text{rural}_{ij} = 1 - \text{urban}_{ij}$, then we can reformulate (3) as

$$\text{logit}(\pi_{ij}) = \beta_0 \text{rural}_{ij} + (\beta_0 + \beta_1) \text{urban}_{ij} + \mathcal{F}_{ij} + u_j \text{rural}_{ij} + (u_j + v_j) \text{urban}_{ij} \quad (3a)$$

where we have translated both the fixed portion and the random portion to be in terms of `rural` rather than a random intercept. Translating the fixed portion is not necessary to make the point we make below, but we do so anyway for uniformity.

Translating the estimated random-effects parameters from the previous output to ones appropriate for (3a), we get $\text{Var}(u_j) = \hat{\sigma}_u^2 = 0.39$,

$$\begin{aligned} \text{Var}(u_j + v_j) &= \hat{\sigma}_u^2 + \hat{\sigma}_v^2 + 2\hat{\sigma}_{uv} \\ &= 0.39 + 0.67 - 2(0.41) = 0.24 \end{aligned}$$

$$\text{and } \text{Cov}(u_j, u_j + v_j) = \hat{\sigma}_u^2 + \hat{\sigma}_{uv} = 0.39 - 0.41 = -0.02.$$

An alternative that does not require remembering how to calculate variances and covariances involving sums—and one that also gives you standard errors—is to let Stata do the work for you:

```
. generate byte rural = 1 - urban
. melogit c_use rural urban age child*, noconstant || district: rural urban,
> noconstant cov(unstructured)
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1229.5485
Iteration 1:   log likelihood = -1228.5268
Iteration 2:   log likelihood = -1228.5263
Iteration 3:   log likelihood = -1228.5263
Refining starting values:
Grid node 0:   log likelihood = -1208.3922
Fitting full model:
Iteration 0:   log likelihood = -1208.3922 (not concave)
Iteration 1:   log likelihood = -1203.556 (not concave)
Iteration 2:   log likelihood = -1200.5896
Iteration 3:   log likelihood = -1199.7288
Iteration 4:   log likelihood = -1199.3373
Iteration 5:   log likelihood = -1199.3151
Iteration 6:   log likelihood = -1199.315
Mixed-effects logistic regression
Group variable:      district
Number of obs       =      1,934
Number of groups    =        60
Obs per group:
    min =           2
    avg =          32.2
    max =          118
Integration method: mvaghermite
Integration pts.    =           7
Wald chi2(6)       =      120.24
Prob > chi2        =       0.0000
Log likelihood = -1199.315
( 1) [c_use]_cons = 0
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
rural	-1.711652	.1605617	-10.66	0.000	-2.026347	-1.396956
urban	-.8958623	.1704954	-5.25	0.000	-1.230027	-.5616974
age	-.026415	.008023	-3.29	0.001	-.0421398	-.0106903
child1	1.13252	.1603285	7.06	0.000	.8182819	1.446758
child2	1.357739	.1770522	7.67	0.000	1.010724	1.704755
child3	1.353827	.18288	7.40	0.000	.9953883	1.712265
_cons	0 (omitted)					
district						
var(rural)	.3897485	.1292403			.2034823	.7465212
var(urban)	.2442899	.1450625			.0762871	.7822759
district						
cov(rural, urban)	-.0161411	.1057462	-0.15	0.879	-.2233999	.1911177

```
LR test vs. logistic model: chi2(3) = 58.42
Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

The above output demonstrates an equivalent fit to that we displayed for model (3), with the added benefit of a more direct comparison of the parameters for rural and urban areas.



□ Technical note

We used the binary variables `rural` and `urban` instead of the factor notation `i.urban` because, although supported in the fixed-effects specification of the model, such notation is not supported in random-effects specifications.

□

□ Technical note

Our model fits for (3) and (3a) are equivalent only because we allowed for correlation in the random effects for both. Had we used the default independent covariance structure, we would be fitting different models; in (3) we would be making the restriction that $\text{Cov}(u_j, v_j) = 0$, whereas in (3a) we would be assuming that $\text{Cov}(u_j, u_j + v_j) = 0$.

The moral here is that although `melogit` will do this by default, one should be cautious when imposing an independent covariance structure, because the correlation between random effects is not invariant to model translations that would otherwise yield equivalent results in standard regression models. In our example, we remapped an intercept and binary coefficient to two complementary binary coefficients, something we could do in standard logistic regression without consequence but that here required more consideration.

Rabe-Hesketh and Skrondal (2012, sec. 11.4) provide a nice discussion of this phenomenon in the related case of recentering a continuous covariate.

□

Other covariance structures

In the above examples, we demonstrated the `independent` and `unstructured` covariance structures. Also available are `identity` (seen previously in output but not directly specified), which restricts random effects to be uncorrelated and share a common variance, and `exchangeable`, which assumes a common variance and a common pairwise covariance.

You can also specify multiple random-effects equations at the same level, in which case the above four covariance types can be combined to form more complex blocked-diagonal covariance structures. This could be used, for example, to impose an equality constraint on a subset of variance components or to otherwise group together a set of related random effects.

Continuing the previous example, typing

```
. melogit c_use urban age child* || district: child*, cov(exchangeable) ||
> district:
```

would fit a model with the same fixed effects as (3) but with random-effects structure

$$\text{logit}(\pi_{ij}) = \beta_0 + \cdots + u_{1j}\text{child1}_{ij} + u_{2j}\text{child2}_{ij} + u_{3j}\text{child3}_{ij} + v_j$$

That is, we have random coefficients on each indicator variable for children (the first `district:` specification) and an overall district random intercept (the second `district:` specification). The above syntax fits a model with overall covariance structure

$$\Sigma = \text{Var} \begin{bmatrix} u_{1j} \\ u_{2j} \\ u_{3j} \\ v_j \end{bmatrix} = \begin{bmatrix} \sigma_u^2 & \sigma_c & \sigma_c & 0 \\ \sigma_c & \sigma_u^2 & \sigma_c & 0 \\ \sigma_c & \sigma_c & \sigma_u^2 & 0 \\ 0 & 0 & 0 & \sigma_v^2 \end{bmatrix}$$

reflecting the relationship among the random coefficients for children. We did not have to specify `noconstant` on the first `district:` specification. `melogit` automatically avoids collinearity by including an intercept on only the final specification among repeated-level equations.

Of course, if we fit the above model, we would heed our own advice from the previous technical note and make sure that not only our data but also our specification characterization of the random effects permitted the above structure. That is, we would check the above against a model that had an `unstructured` covariance for all four random effects and then perhaps against a model that assumed an `unstructured` covariance among the three random coefficients on children, coupled with independence with the random intercept. All comparisons can be made by storing estimates (command `estimates store`) and then using `lrtest`, as demonstrated previously.

Three-level models

► Example 4: Three-level random-intercept model

Rabe-Hesketh, Touloupoulou, and Murray (2001) analyzed data from a study measuring the cognitive ability of patients with schizophrenia compared with their relatives and control subjects. Cognitive ability was measured as the successful completion of the “Tower of London”, a computerized task, measured at three levels of difficulty. For all but one of the 226 subjects, there were three measurements (one for each difficulty level). Because patients’ relatives were also tested, a family identifier, `family`, was also recorded.

```
. use http://www.stata-press.com/data/r15/towerlondon, clear
(Tower of London data)

. describe

Contains data from http://www.stata-press.com/data/r15/towerlondon.dta
  obs:                677                Tower of London data
  vars:                 5                31 May 2016 10:41
  size:               4,739              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
family	int	%8.0g		Family ID
subject	int	%9.0g		Subject ID
dtlm	byte	%9.0g		1 = task completed
difficulty	byte	%9.0g		Level of difficulty: -1, 0, or 1
group	byte	%8.0g		1: controls; 2: relatives; 3: schizophrenics

Sorted by: family subject

We fit a logistic model with response `dtlm`, the indicator of cognitive function, and with covariates `difficulty` and a set of indicator variables for `group`, with the controls (`group==1`) being the base category. We allow for random effects due to families and due to subjects within families, and we request to see odds ratios.

tended not to perform as well as the control subjects. Of course, we would make similar conclusions from a standard logistic model fit to the same data, but the odds ratios would differ somewhat.

❑ **Technical note**

In the [previous example](#), the subjects are coded with unique values between 1 and 251 (with some gaps), but such coding is not necessary to produce nesting within families. Once we specified the nesting structure to `melogit`, all that was important was the relative coding of `subject` within each unique value of `family`. We could have coded `subjects` as the numbers 1, 2, 3, and so on, restarting at 1 with each new family, and `melogit` would have produced the same results.

Group identifiers may also be coded using string variables.

The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by `||`. The order of nesting goes from left to right as the groups go from biggest (highest level) to smallest (lowest level).

Crossed-effects models

➤ **Example 5: Crossed-effects model**

[Rabe-Hesketh and Skrondal \(2012, 443–460\)](#) perform an analysis on school data from Fife, Scotland. The data, originally from [Paterson \(1991\)](#), are from a study measuring students’ attainment as an integer score from 1 to 10, based on the Scottish school exit examination taken at age 16. The study comprises 3,435 students who first attended any one of 148 primary schools and then any one of 19 secondary schools.

```
. use http://www.stata-press.com/data/r15/fifeschool
(School data from Fife, Scotland)

. describe

Contains data from http://www.stata-press.com/data/r15/fifeschool.dta
  obs:                3,435      School data from Fife, Scotland
 vars:                  5        28 May 2016 10:08
size:                24,045      (_dta has notes)
```

variable name	storage type	display format	value label	variable label
pid	int	%9.0g		Primary school ID
sid	byte	%9.0g		Secondary school ID
attain	byte	%9.0g		Attainment score at age 16
vrq	int	%9.0g		Verbal-reasoning score from final year of primary school
sex	byte	%9.0g		1: female; 0: male

```
Sorted by:

. generate byte attain_gt_6 = attain > 6
```

To make the analysis relevant to our present discussion, we focus not on the attainment score itself but instead on whether the score is greater than 6. We wish to model this indicator as a function of the fixed effect `sex` and of random effects due to primary and secondary schools.

For this analysis, it would make sense to assume that the random effects are not nested, but instead crossed, meaning that the effect due to primary school is the same regardless of the secondary school attended. Our model is thus

$$\text{logit}\{\Pr(\text{attain}_{ijk} > 6)\} = \beta_0 + \beta_1 \text{sex}_{ijk} + u_j + v_k \quad (4)$$

for student i , $i = 1, \dots, n_{jk}$, who attended primary school j , $j = 1, \dots, 148$, and then secondary school k , $k = 1, \dots, 19$.

Because there is no evident nesting, one solution would be to consider the data as a whole and fit a two-level, one-cluster model with random-effects structure

$$\mathbf{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_{148} \\ v_1 \\ \vdots \\ v_{19} \end{bmatrix} \sim N(\mathbf{0}, \Sigma); \quad \Sigma = \begin{bmatrix} \sigma_u^2 \mathbf{I}_{148} & \mathbf{0} \\ \mathbf{0} & \sigma_v^2 \mathbf{I}_{19} \end{bmatrix}$$

We can fit such a model by using the group designation `_all:`, which tells `melogit` to treat the whole dataset as one cluster, and the `R.varname` notation, which mimics the creation of indicator variables identifying schools:

```
. melogit attain_gt_6 sex || _all:R.pid || _all:R.sid, or
```

But we do not recommend fitting the model this way because of high total dimension ($148 + 19 = 167$) of the random effects. This would require working with matrices of column dimension 167, which is probably not a problem for most current hardware, but would be a problem if this number got much larger.

An equivalent way to fit (4) that has a smaller dimension is to treat the clusters identified by primary schools as nested within all the data, that is, as nested within the `_all` group.

```
. melogit attain_gt_6 sex || _all:R.sid || pid:, or
note: crossed random-effects model specified; option intmethod(laplace)
implied

Fitting fixed-effects model:
Iteration 0:  log likelihood = -2320.2374
Iteration 1:  log likelihood = -2317.9062
Iteration 2:  log likelihood = -2317.9059
Iteration 3:  log likelihood = -2317.9059

Refining starting values:
Grid node 0:  log likelihood = -2234.6403

Fitting full model:
Iteration 0:  log likelihood = -2234.6403 (not concave)
Iteration 1:  log likelihood = -2227.9507 (not concave)
Iteration 2:  log likelihood = -2227.9287 (not concave)
Iteration 3:  log likelihood = -2227.9265 (not concave)
Iteration 4:  log likelihood = -2227.9263
Iteration 5:  log likelihood = -2221.6885
Iteration 6:  log likelihood = -2220.3474
Iteration 7:  log likelihood = -2220.2912
Iteration 8:  log likelihood = -2220.0342
Iteration 9:  log likelihood = -2220.0047
Iteration 10: log likelihood = -2220.0035
Iteration 11: log likelihood = -2220.0035

Mixed-effects logistic regression                                Number of obs      =      3,435



| Group Variable | No. of Groups | Observations per Group<br>Minimum | Average | Maximum |
|----------------|---------------|-----------------------------------|---------|---------|
| _all           | 1             | 3,435                             | 3,435.0 | 3,435   |
| pid            | 148           | 1                                 | 23.2    | 72      |



Integration method:      laplace

Log likelihood = -2220.0035                                Wald chi2(1)      =      14.31
                                                                Prob > chi2       =      0.0002



| attain_gt_6            | Odds Ratio | Std. Err. | z     | P> z  | [95% Conf. Interval] |
|------------------------|------------|-----------|-------|-------|----------------------|
| sex                    | 1.325116   | .0985998  | 3.78  | 0.000 | 1.145295 1.533171    |
| _cons                  | .5311581   | .0622658  | -5.40 | 0.000 | .4221244 .6683549    |
| _all>sid<br>var(_cons) | .1239665   | .0692734  |       |       | .0414617 .3706485    |
| pid<br>var(_cons)      | .4520427   | .0952823  |       |       | .2990624 .6832775    |



Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline odds (conditional on zero random effects).
LR test vs. logistic model: chi2(2) = 195.80                      Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

Choosing the primary schools as those to nest was no accident; because there are far fewer secondary schools than primary schools, the above required only 19 random coefficients for the secondary schools and one random intercept at the primary school level, for a total dimension of 20. Our data also include a measurement of verbal reasoning, the variable `vrq`. Adding a fixed effect due to `vrq` in (4) would negate the effect due to secondary school, a fact we leave to you to verify as an exercise.

See [ME] **mixed** for a similar discussion of crossed effects in the context of linear mixed models. Also see [Rabe-Hesketh and Skrondal \(2012\)](#) for more examples of crossed-effects models, including models with random interactions, and for more techniques on how to avoid high-dimensional estimation.

□ Technical note

The estimation in the previous example was performed using a Laplacian approximation, even though we did not specify this. Whenever the R. notation is used in random-effects specifications, estimation reverts to the Laplacian method because of the high dimension induced by having the R. variables.

In the above example, through some creative nesting, we reduced the dimension of the random effects to 20, but this is still too large to permit estimation via adaptive Gaussian quadrature; see [Computation time and the Laplacian approximation](#) in [ME] **me**. Even with two quadrature points, our rough formula for computation time would contain within it a factor of $2^{20} = 1,048,576$.

The `intmethod(laplace)` option is therefore assumed when you use R. notation. If the number of distinct levels of your R. variables is small enough (say, five or fewer) to permit estimation via quadrature, you can override the imposition of `laplace` by specifying a different integration method in the `intmethod()` option.

□

Stored results

`melogit` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>melogit</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	fweight variable for <i>k</i> th highest level, if specified
<code>e(iweightk)</code>	iweight variable for <i>k</i> th highest level, if specified
<code>e(pweightk)</code>	pweight variable for <i>k</i> th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivals)</code>	grouping variables
<code>e(model)</code>	<code>logistic</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>logit</code>
<code>e(family)</code>	<code>bernoulli</code> or <code>binomial</code>
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	<code>offset</code>
<code>e(binomial)</code>	binomial number of trials
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<i>vcetype</i> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>
<code>e(marginswexp)</code>	weight expression for <code>margins</code>
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance–covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

`melogit` is a convenience command for `meglm` with a `logit` link and a `bernoulli` or `binomial` family; see [\[ME\]](#) [meglm](#).

Model (1) assumes Bernoulli data, a special case of the binomial. Because binomial data are also supported by `melogit` (option `binomial()`), the methods presented below are in terms of the more general binomial mixed-effects model.

For a two-level binomial model, consider the response y_{ij} as the number of successes from a series of r_{ij} Bernoulli trials (replications). For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$\begin{aligned} f(\mathbf{y}_j | \mathbf{u}_j) &= \prod_{i=1}^{n_j} \left[\binom{r_{ij}}{y_{ij}} \{H(\boldsymbol{\eta}_{ij})\}^{y_{ij}} \{1 - H(\boldsymbol{\eta}_{ij})\}^{r_{ij}-y_{ij}} \right] \\ &= \exp \left(\sum_{i=1}^{n_j} \left[y_{ij} \boldsymbol{\eta}_{ij} - r_{ij} \log \{1 + \exp(\boldsymbol{\eta}_{ij})\} + \log \binom{r_{ij}}{y_{ij}} \right] \right) \end{aligned}$$

for $\boldsymbol{\eta}_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \text{offset}_{ij}$ and $H(v) = \exp(v)/\{1 + \exp(v)\}$.

Defining $\mathbf{r}_j = (r_{j1}, \dots, r_{jn_j})'$ and

$$c(\mathbf{y}_j, \mathbf{r}_j) = \sum_{i=1}^{n_j} \log \binom{r_{ij}}{y_{ij}}$$

where $c(\mathbf{y}_j, \mathbf{r}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j | \mathbf{u}_j) = \exp \left[\mathbf{y}_j' \boldsymbol{\eta}_j - \mathbf{r}_j' \log \{1 + \exp(\boldsymbol{\eta}_j)\} + c(\mathbf{y}_j, \mathbf{r}_j) \right]$$

where $\boldsymbol{\eta}_j$ is formed by stacking the row vectors $\boldsymbol{\eta}_{ij}$. We extend the definitions of the functions $\log(\cdot)$ and $\exp(\cdot)$ to be vector functions where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{y}_j | \mathbf{u}_j) \exp(-\mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp \{c(\mathbf{y}_j, \mathbf{r}_j)\} (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int \exp \{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \tag{2}$$

where

$$h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) = \mathbf{y}_j' \boldsymbol{\eta}_j - \mathbf{r}_j' \log \{1 + \exp(\boldsymbol{\eta}_j)\} - \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] `meglm` for details.

`melogit` supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] `meglm` for details.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Harbord, R. M., and P. Whiting. 2009. [metandi: Meta-analysis of diagnostic accuracy using hierarchical logistic regression](#). *Stata Journal* 9: 211–229.
- Huq, N. M., and J. Cleland. 1990. *Bangladesh Fertility Survey 1989 (Main Report)*. National Institute of Population Research and Training.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Ng, E. S.-W., J. R. Carpenter, H. Goldstein, and J. Rasbash. 2006. Estimation in generalised linear mixed models with binary outcomes by simulated maximum likelihood. *Statistical Modelling* 6: 23–42.
- Paterson, L. 1991. Socio-economic status and educational attainment: A multidimensional and multilevel study. *Evaluation and Research in Education* 5: 97–121.
- Rabe-Hesketh, S., and A. Skrondal. 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Rabe-Hesketh, S., T. Touloupoulou, and R. M. Murray. 2001. Multilevel modeling of cognitive function in schizophrenic patients and their first degree relatives. *Multivariate Behavioral Research* 36: 279–298.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.

Also see

- [ME] **melogit postestimation** — Postestimation tools for melogit
- [ME] **mecloglog** — Multilevel mixed-effects complementary log-log regression
- [ME] **meprobit** — Multilevel mixed-effects probit regression
- [ME] **meqrlogit** — Multilevel mixed-effects logistic regression (QR decomposition)
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: melogit** — Bayesian multilevel logistic regression
- [SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtlogit** — Fixed-effects, random-effects, and population-averaged logit models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples
Also see

[predict](#)
Methods and formulas

[margins](#)
References

Postestimation commands

The following postestimation commands are of special interest after `melogit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat icc</code>	estimate intraclass correlations
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
density	predicted density function
distribution	predicted distribution function
pearson	Pearson residuals
deviance	deviance residuals
anscombe	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
Integration	
<i>int_options</i>	integration options
pearson, deviance, anscombe may not be combined with <code>marginal</code> .	
<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only
<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<i>int_options</i>	integration options
<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main
mu, the default, calculates the predicted mean, that is, the probability of a positive outcome.
eta, xb, stdp, density, distribution, pearson, deviance, anscombe, scores, conditional(), marginal, and nooffset; see [ME] meglm postestimation .
reffects, ebmeans, ebmodes, and reses(); see [ME] meglm postestimation .
Integration
intpoints(), iterate(), and tolerance(); see [ME] meglm postestimation .

margins

Description for margins

margins estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
<u>density</u>	not allowed with margins
<u>distribution</u>	not allowed with margins
<u>pearson</u>	not allowed with margins
<u>deviance</u>	not allowed with margins
<u>anscombe</u>	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options conditional(ebmeans) and conditional(ebmodes) are not allowed with margins.
Option marginal is assumed where applicable if conditional(fixedonly) is not specified.

Statistics not allowed with margins are functions of stochastic quantities other than e(b).
For the full syntax, see [R] margins.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a logistic mixed-effects model with `melogit`. For the most part, calculation centers around obtaining estimates of the subject/group-specific random effects. Random effects are not provided as estimates when the model is fit but instead need to be predicted after estimation. Calculation of intraclass correlations, estimating the dependence between latent linear responses for different levels of nesting, may also be of interest.

➤ Example 1: Estimating the intraclass correlation

Following [Rabe-Hesketh and Skrondal \(2012, chap. 10\)](#), we consider a two-level mixed-effects model for onycholysis (separation of toenail plate from nail bed) among those who contract toenail fungus. The data are obtained from [De Backer et al. \(1998\)](#) and were also studied by [Lesaffre and Spiessens \(2001\)](#). The onycholysis outcome is dichotomously coded as 1 (moderate or severe onycholysis) or 0 (none or mild onycholysis). Fixed-effects covariates include treatment (0: itraconazole; 1: terbinafine), the month of measurement, and their interaction.

We fit the two-level model with `melogit`:

```
. use http://www.stata-press.com/data/r15/toenail
(Onycholysis severity)
. melogit outcome treatment month trt_month || patient:, intpoints(30)
(iteration log omitted)

Mixed-effects logistic regression      Number of obs    =      1,908
Group variable:      patient          Number of groups  =       294
                                   Obs per group:
                                   min =           1
                                   avg  =          6.5
                                   max  =           7

Integration method: mvaghermite        Integration pts.  =       30
                                   Wald chi2(3)      =    150.61
Log likelihood = -625.38557              Prob > chi2       =     0.0000
```

outcome	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
treatment	-.1608934	.5802058	-0.28	0.782	-1.298076	.9762891
month	-.3911056	.0443906	-8.81	0.000	-.4781097	-.3041016
trt_month	-.1368286	.0680213	-2.01	0.044	-.2701479	-.0035093
_cons	-1.620355	.4322382	-3.75	0.000	-2.467526	-.7731834
patient						
var(_cons)	16.0841	3.062625			11.07431	23.36021

```
LR test vs. logistic model: chibar2(01) = 565.24      Prob >= chibar2 = 0.0000
```

It is of interest to determine the dependence among responses for the same subject (between-subject heterogeneity). Under the latent-linear-response formulation, this dependence can be obtained with the intraclass correlation. Formally, in a two-level random-effects model, the intraclass correlation corresponds to the correlation of latent responses within the same individual and also to the proportion of variance explained by the individual random effect.

We use `estat icc` to estimate the residual intraclass correlation:

```
. estat icc
```

```
Residual intraclass correlation
```

Level	ICC	Std. Err.	[95% Conf. Interval]	
patient	.8301913	.0268433	.7709672	.8765531

In the presence of fixed-effects covariates, `estat icc` reports the residual intraclass correlation, which is the correlation between latent linear responses conditional on the fixed-effects covariates.

Conditional on treatment and month of treatment, we estimate that latent responses within the same patient have a large correlation of approximately 0.83. Further, 83% of the variance of a latent response is explained by the between-patient variability.

◀

► Example 2: Predicting random effects

In [example 3](#) of [\[ME\] melogit](#), we represented the probability of contraceptive use among Bangladeshi women by using the model (stated with slightly different notation here)

$$\text{logit}(\pi_{ij}) = \beta_0 \text{rural}_{ij} + \beta_1 \text{urban}_{ij} + \beta_2 \text{age}_{ij} + \beta_3 \text{child1}_{ij} + \beta_4 \text{child2}_{ij} + \beta_5 \text{child3}_{ij} + a_j \text{rural}_{ij} + b_j \text{urban}_{ij}$$

where π_{ij} is the probability of contraceptive use, $j = 1, \dots, 60$ districts, $i = 1, \dots, n_j$ women within each district, and a_j and b_j are normally distributed with mean 0 and variance–covariance matrix

$$\Sigma = \text{Var} \begin{bmatrix} a_j \\ b_j \end{bmatrix} = \begin{bmatrix} \sigma_a^2 & \sigma_{ab} \\ \sigma_{ab} & \sigma_b^2 \end{bmatrix}$$


```
. use http://www.stata-press.com/data/r15/bangladesh
(Bangladesh Fertility Survey, 1989)

. generate byte rural = 1 - urban

. melogit c_use rural urban age child*, noconstant || district: rural urban,
> noconstant cov(unstructured)

(iteration log omitted)

Mixed-effects logistic regression      Number of obs   =       1,934
Group variable:      district          Number of groups =         60
                                   Obs per group:
                                   min =           2
                                   avg =          32.2
                                   max =          118

Integration method: mvaghermite        Integration pts. =           7

Log likelihood = -1199.315              Wald chi2(6)    =       120.24
( 1)  [c_use]_cons = 0                  Prob > chi2     =       0.0000
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
rural	-1.711652	.1605617	-10.66	0.000	-2.026347	-1.396956
urban	-.8958623	.1704954	-5.25	0.000	-1.230027	-.5616974
age	-.026415	.008023	-3.29	0.001	-.0421398	-.0106903
child1	1.13252	.1603285	7.06	0.000	.8182819	1.446758
child2	1.357739	.1770522	7.67	0.000	1.010724	1.704755
child3	1.353827	.18288	7.40	0.000	.9953883	1.712265
_cons	0	(omitted)				
district						
var(rural)	.3897485	.1292403			.2034823	.7465212
var(urban)	.2442899	.1450625			.0762871	.7822759
district						
cov(rural, urban)	-.0161411	.1057462	-0.15	0.879	-.2233999	.1911177

LR test vs. logistic model: $\chi^2(3) = 58.42$ Prob > $\chi^2 = 0.0000$
 Note: LR test is conservative and provided only for reference.

The purpose of using this particular model was to allow for district random effects that were specific to the rural and urban areas of that district and that could be interpreted as such. We can obtain predictions of these random effects and their corresponding standard errors,

```
. predict re_rural re_urban, reffects reses(se_rural se_urban)
(calculating posterior means of random effects)
(using 7 quadrature points)
```

The order in which we specified the variables to be generated corresponds to the order in which the variance components are listed in `melogit` output. If in doubt, a simple `describe` will show how these newly generated variables are labeled just to be sure.

Having generated estimated random effects and standard errors, we can now list them for the first 10 districts:

```
. by district, sort: generate tolist = (_n==1)
. list district re_rural se_rural re_urban se_urban if district <= 10 & tolist,
> sep(0)
```

	district	re_rural	se_rural	re_urban	se_urban
1.	1	-.9432691	.3156852	-.558359	.2332665
118.	2	-.0427011	.3822029	.0017684	.493834
138.	3	-.0149571	.6242161	.2263715	.4698407
140.	4	-.2864846	.3990107	.5869046	.3988538
170.	5	.0688648	.3102899	.0046016	.4685461
209.	6	-.3979315	.2762392	.2761181	.4204175
274.	7	-.1910399	.4046772	.0079117	.4938647
292.	8	.034071	.3180097	.2266263	.4689558
329.	9	-.3737992	.3923573	.0764026	.4569863
352.	10	-.5640147	.4769353	.0233582	.4939753



□ Technical note

When these data were first introduced in [ME] [melogit](#), we noted that not all districts contained both urban and rural areas. This fact is somewhat demonstrated by the random effects that are nearly 0 in the above. A closer examination of the data would reveal that district 3 has no rural areas, and districts 2, 7, and 10 have no urban areas.

The estimated random effects are not exactly 0 in these cases because of the correlation between urban and rural effects. For instance, if a district has no urban areas, it can still yield a nonzero (albeit small) random-effects estimate for a nonexistent urban area because of the correlation with its rural counterpart.

Had we imposed an independent covariance structure in our model, the estimated random effects in the cases in question would be exactly 0.



□ Technical note

The estimated standard errors produced above with the `reses()` option are conditional on the values of the estimated model parameters: β and the components of Σ . Their interpretation is therefore not one of standard sample-to-sample variability but instead one that does not incorporate uncertainty in the estimated model parameters; see [Methods and formulas](#).

That stated, conditional standard errors can still be used as a measure of relative precision, provided that you keep this caveat in mind.



▷ Example 3: Obtaining predicted probabilities

Continuing with [example 2](#), we can obtain predicted probabilities, the default prediction:

```
. predict p
(predictions based on fixed effects and posterior means of random effects)
(option mu assumed)
(using 7 quadrature points)
```

These predictions are based on a linear predictor that includes both the fixed effects and the random effects due to district. Specifying the conditional(fixedonly) option gives predictions that set the random effects to their prior mean of 0. Below we compare both over the first 20 observations:

```
. predict p_fixed, conditional(fixedonly)
(option mu assumed)

. list c_use p p_fixed age child* in 1/20
```

	c_use	p	p_fixed	age	child1	child2	child3
1.	no	.3572114	.4927182	18.44	0	0	1
2.	no	.21293	.3210403	-5.56	0	0	0
3.	no	.4664207	.6044016	1.44	0	1	0
4.	no	.4198625	.5584863	8.44	0	0	1
5.	no	.2504834	.3687281	-13.56	0	0	0
6.	no	.2406963	.3565185	-11.56	0	0	0
7.	no	.3572114	.4927182	18.44	0	0	1
8.	no	.4984106	.6345998	-3.56	0	0	1
9.	no	.4564025	.594723	-5.56	1	0	0
10.	no	.465447	.6034657	1.44	0	0	1
11.	yes	.2406963	.3565185	-11.56	0	0	0
12.	no	.1999512	.3040173	-2.56	0	0	0
13.	no	.4498569	.5883406	-4.56	1	0	0
14.	no	.439278	.5779263	5.44	0	0	1
15.	no	.4786124	.6160359	-0.56	0	0	1
16.	yes	.4457945	.584356	4.44	0	0	1
17.	no	.21293	.3210403	-5.56	0	0	0
18.	yes	.4786124	.6160359	-0.56	0	0	1
19.	yes	.4629632	.6010735	-6.56	1	0	0
20.	no	.4993888	.6355067	-3.56	0	1	0



► Example 4: Intraclass correlations for higher-level models

Continuing with [example 2](#), we can also compute intraclass correlations for the model using estat icc; see [\[ME\] estat icc](#).

In the presence of random-effects covariates, the intraclass correlation is no longer constant and depends on the values of the random-effects covariates. In this case, estat icc reports conditional intraclass correlations assuming 0 values for all random-effects covariates. For example, in a two-level model, this conditional correlation represents the correlation of the latent responses for two measurements on the same subject, both of which have random-effects covariates equal to 0. Similarly to the interpretation of intercept variances in random-coefficient models ([Rabe-Hesketh and Skrondal 2012](#), chap. 16), interpretation of this conditional intraclass correlation relies on the usefulness of the 0 baseline values of random-effects covariates. For example, mean centering of the covariates is often used to make a 0 value a useful reference.

Estimation of the conditional intraclass correlation in the Bangladeshi contraceptive study setting of [example 2](#) is of interest. In [example 2](#), the random-effects covariates rural and urban for the random level district are mutually exclusive indicator variables and can never be simultaneously 0. Thus we could not use estat icc to estimate the conditional intraclass correlation for this model, because estat icc requires that the random intercept is included in all random-effects specifications.

Instead, we consider an alternative model for the Bangladeshi contraceptive study. In [example 2](#) of [\[ME\] melogit](#), we represented the probability of contraceptive use among Bangladeshi women with fixed-effects for urban residence (`urban`), age (`age`), and the number of children (`child1`–`child3`). The random effects for urban and rural residence are represented as a random slope for urban residence and a random intercept at the district level.

We fit the model with `melogit`:

```
. use http://www.stata-press.com/data/r15/bangladesh, clear
(Bangladesh Fertility Survey, 1989)

. melogit c_use urban age child* || district: urban, covariance(unstructured)
(iteration log omitted)
```

Mixed-effects logistic regression	Number of obs	=	1,934
Group variable: district	Number of groups	=	60
	Obs per group:		
	min	=	2
	avg	=	32.2
	max	=	118
Integration method: mvaghermite	Integration pts.	=	7
Log likelihood = -1199.315	Wald chi2(5)	=	97.50
	Prob > chi2	=	0.0000

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	.8157875	.1715519	4.76	0.000	.4795519	1.152023
age	-.026415	.008023	-3.29	0.001	-.0421398	-.0106902
child1	1.13252	.1603285	7.06	0.000	.818282	1.446758
child2	1.357739	.1770522	7.67	0.000	1.010723	1.704755
child3	1.353827	.1828801	7.40	0.000	.9953882	1.712265
_cons	-1.71165	.1605618	-10.66	0.000	-2.026345	-1.396954
district						
var(urban)	.6663237	.3224689			.258074	1.720387
var(_cons)	.3897448	.1292463			.203473	.7465413
district						
cov(urban, _cons)	-.4058861	.1755414	-2.31	0.021	-.7499408	-.0618313

LR test vs. logistic model: $\chi^2(3) = 58.42$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

We use `estat icc` to estimate the intraclass correlation conditional on urban being equal to 0:

```
. estat icc
Conditional intraclass correlation
```

Level	ICC	Std. Err.	[95% Conf. Interval]	
district	.1059201	.0314045	.058246	.1849518

Note: ICC is conditional on zero values of random-effects covariates.

This estimate suggests that the latent responses are not strongly correlated for rural residents (`urban == 0`) within the same district, conditional on the fixed-effects covariates.

`estat icc` reports two intraclass correlations for this three-level nested model. The first is the level-3 intraclass correlation at the family level, the correlation between latent measurements of the cognitive ability in the same family. The second is the level-2 intraclass correlation at the subject-within-family level, the correlation between the latent measurements of cognitive ability in the same subject and family.

There is not a strong correlation between individual realizations of the latent response, even within the same subject.



Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [ME] [meglm postestimation](#).

References

- De Backer, M., C. De Vroey, E. Lesaffre, I. Scheys, and P. De Keyser. 1998. Twelve weeks of continuous oral therapy for toenail onychomycosis caused by dermatophytes: A double-blind comparative trial of terbinafine 250 mg/day versus itraconazole 200 mg/day. *Journal of the American Academy of Dermatology* 38: S57–S63.
- Lesaffre, E., and B. Spiessens. 2001. On the effect of the number of quadrature points in a logistic random-effects model: An example. *Journal of the Royal Statistical Society, Series C* 50: 325–335.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.

Also see

- [ME] [melogit](#) — Multilevel mixed-effects logistic regression
- [ME] [meglm postestimation](#) — Postestimation tools for meglm
- [U] [20 Estimation and postestimation commands](#)

Title

menbreg — Multilevel mixed-effects negative binomial regression

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`menbreg` fits mixed-effects negative binomial models to count data. The conditional distribution of the response given random effects is assumed to follow a Poisson-like process, except that the variation is greater than that of a true Poisson process.

Quick start

Mixed-effects negative binomial regression of `y` on `x` with random intercepts by `v1`

```
menbreg y x || v1:
```

Add `evvar` measuring exposure

```
menbreg y x, exposure(evvar) || v1:
```

As above, but report incidence-rate ratios instead of coefficients

```
menbreg y x, exposure(evvar) || v1:, irr
```

Add random coefficients for `x`

```
menbreg y x, exposure(evvar) || v1: x, irr
```

Three-level random-intercept model of `y` on `x` with `v1` nested within `v2`

```
menbreg y x || v2: || v1:
```

Menu

Statistics > Multilevel mixed-effects models > Negative binomial regression

Syntax

```
menbreg depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress the constant term from the fixed-effects equation
<code>exposure(<i>varname_e</i>)</code>	include $\ln(\text{varname}_e)$ in model with coefficient constrained to 1
<code>offset(<i>varname_o</i>)</code>	include <i>varname_o</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>dispersion</u> (<i>dispersion</i>)	parameterization of the conditional overdispersion; <i>dispersion</i> may be <code>mean</code> (default) or <code>constant</code>
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <u><code>robust</code></u> , or <u><code>cluster</code></u> <i>clustvar</i>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>irr</u>	report fixed-effects coefficients as incidence-rate ratios
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<code>mvaghermite</code>	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<code>mcaghermite</code>	mode-curvature adaptive Gauss–Hermite quadrature
<code>ghermite</code>	nonadaptive Gauss–Hermite quadrature
<code>laplace</code>	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: menbreg.

vce() and weights are not allowed with the *svy* prefix; see [SVY] svy.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

`noconstant` suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

`exposure(varnamee)` specifies a variable that reflects the amount of exposure over which the *depvar* events were observed for each observation; $\ln(\text{varname}_e)$ is included in the fixed-effects portion of the model with the coefficient constrained to be 1.

`offset(varnameo)` specifies that *varname_o* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

`covariance(vartype)` specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: `independent`, `exchangeable`, `identity`, `unstructured`, `fixed(matname)`, or `pattern(matname)`.

`covariance(independent)` covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is `covariance(independent)` unless a crossed random-effects model is fit, in which case the default is `covariance(identity)`.

`covariance(exchangeable)` structure specifies one common variance for all random effects and one common pairwise covariance.

`covariance(identity)` is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

`covariance(unstructured)` allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

`covariance(fixed(matname))` and `covariance(pattern(matname))` covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names

of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if $\text{matname}[i, j] = \text{matname}[k, l]$.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwivar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`dispersion(mean | constant)` specifies the parameterization of the conditional overdispersion given random effects. `dispersion(mean)`, the default, yields a model where the conditional overdispersion is a function of the conditional mean given random effects. For example, in a two-level model, the conditional overdispersion is equal to $1 + \alpha E(y_{ij} | \mathbf{u}_j)$. `dispersion(constant)` yields a model where the conditional overdispersion is constant and is equal to $1 + \delta$. α and δ are the respective conditional overdispersion parameters.

`constraints(constraints)`, `collinear`; see [R] [estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] [estimation options](#).

`irr` reports estimated fixed-effects coefficients transformed to incidence-rate ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated or stored. `irr` may be specified either at estimation or upon replay.

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrntolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `menbreg` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with **menbreg** but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[(gridspec)]`, `noestimate`, and `dnumerical`; see [\[ME\] meglm](#).

`coeflegend`; see [\[R\] estimation options](#).

Remarks and examples

Mixed-effects negative binomial regression is negative binomial regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

menbreg allows for many levels of random effects. However, for simplicity, consider a two-level model, where for a series of M independent clusters, and conditional on the latent variable ζ_{ij} and a set of random effects $\mathbf{u}_j$,

$$y_{ij} | \zeta_{ij} \sim \text{Poisson}(\zeta_{ij})$$

and

$$\zeta_{ij} | \mathbf{u}_j \sim \text{Gamma}(r_{ij}, p_{ij})$$

and

$$\mathbf{u}_j \sim N(\mathbf{0}, \Sigma)$$

where y_{ij} is the count response of the i th observation, $i = 1, \dots, n_j$, from the j th cluster, $j = 1, \dots, M$, and r_{ij} and p_{ij} have two different parameterizations, (2) and (3) below. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of Σ , known as variance components.

The probability that a random response y_{ij} takes the value y is then given by

$$\Pr(y_{ij} = y | \mathbf{u}_j) = \frac{\Gamma(y + r_{ij})}{\Gamma(y + 1)\Gamma(r_{ij})} p_{ij}^{r_{ij}} (1 - p_{ij})^y \quad (1)$$

where for convenience we suppress the dependence of the observable data y_{ij} on r_{ij} and p_{ij} .

Model (1) is an extension of the standard negative binomial model (see [\[R\] nbreg](#)) to incorporate normally distributed random effects at different hierarchical levels. (The negative binomial model itself can be viewed as a random-effects model, a Poisson model with a gamma-distributed random effect.) The standard negative binomial model is used to model overdispersed count data for which the variance is greater than that of a Poisson model. In a Poisson model, the variance is equal to the mean, and thus overdispersion is defined as the extra variability compared with the mean. According to this definition, the negative binomial model presents two different parameterizations of the overdispersion: the mean parameterization, where the overdispersion is a function of the mean, $1 + \alpha E(Y | \mathbf{x})$, $\alpha > 0$; and the constant parameterization, where the overdispersion is a constant function, $1 + \delta$, $\delta \geq 0$. We refer to α and δ as conditional overdispersion parameters.

Let $\mu_{ij} = E(y_{ij} | \mathbf{x}, \mathbf{u}_j) = \exp(\mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j)$, where $\mathbf{x}_{ij}$ is the $1 \times p$ row vector of the fixed-effects covariates, analogous to the covariates you would find in a standard negative binomial regression model, with regression coefficients (fixed effects) β ; $\mathbf{z}_{ij}$ is the $1 \times q$ vector of the random-effects

covariates and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. One special case places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$, so that all covariate effects are essentially random and distributed as multivariate normal with mean β and variance Σ .

Similarly to the standard negative binomial model, we can consider two parameterizations of what we call the conditional overdispersion, the overdispersion conditional on random effects, in a random-effects negative binomial model. For the mean-overdispersion (or, more technically, mean-conditional-overdispersion) parameterization,

$$r_{ij} = 1/\alpha \text{ and } p_{ij} = \frac{1}{1 + \alpha\mu_{ij}} \quad (2)$$

and the conditional overdispersion is equal to $1 + \alpha\mu_{ij}$. For the constant-overdispersion (or, more technically, constant-conditional-overdispersion) parameterization,

$$r_{ij} = \mu_{ij}/\delta \text{ and } p_{ij} = \frac{1}{1 + \delta} \quad (3)$$

and the conditional overdispersion is equal to $1 + \delta$. In what follows, for brevity, we will use the term overdispersion parameter to mean conditional overdispersion parameter, unless stated otherwise.

In the context of random-effects negative binomial models, it is important to decide which model is used as a reference model for the definition of the overdispersion. For example, if we consider a corresponding random-effects Poisson model as a comparison model, the parameters α and δ can still be viewed as unconditional overdispersion parameters, as we show below, although the notion of a constant overdispersion is no longer applicable.

If we retain the definition of the overdispersion as the excess variation with respect to a Poisson process for which the variance is equal to the mean, we need to carefully distinguish between the marginal (unconditional) mean with random effects integrated out and the conditional mean given random effects.

In what follows, for simplicity, we omit the dependence of the formulas on $\mathbf{x}$. Contingent on random effects, the (conditional) dispersion $\text{Var}(y_{ij}|\mathbf{u}_j) = (1 + \alpha\mu_{ij})\mu_{ij}$ for the mean parameterization and $\text{Var}(y_{ij}|\mathbf{u}_j) = (1 + \delta)\mu_{ij}$ for the constant parameterization; the usual interpretation of the parameters holds (conditionally).

If we consider the marginal mean or, specifically, the marginal dispersion for, for example, a two-level random-intercept model, then

$$\text{Var}(y_{ij}) = [1 + \{\exp(\sigma^2)(1 + \alpha) - 1\}E(y_{ij})]E(y_{ij})$$

for the mean parameterization and

$$\text{Var}(y_{ij}) = [1 + \delta + \{\exp(\sigma^2) - 1\}E(y_{ij})]E(y_{ij})$$

for the constant parameterization, where σ^2 is the variance component corresponding to the random intercept.

A few things of interest compared with the standard negative binomial model. First, the random-effects negative binomial model is not strictly an overdispersed model. The combination of values of α and σ^2 can lead to an underdispersed model, a model with smaller variability than the Poisson variability. Underdispersed models are not as common in practice, so we will concentrate on the overdispersion in this entry. Second, α (or δ) no longer solely determine the overdispersion and thus cannot be viewed as unconditional overdispersion parameters. Overdispersion is now a function of both α (or δ) and σ^2 . Third, the notion of a constant overdispersion is not applicable.

Two special cases are worth mentioning. When $\sigma^2 = 0$, the dispersion reduces to that of a standard negative binomial model. When $\alpha = 0$ (or $\delta = 0$), the dispersion reduces to that of a two-level random-intercept Poisson model, which itself is, in general, an overdispersed model; see [Rabe-Hesketh and Skrondal \(2012, chap. 13.7\)](#) for more details. As such, α and δ retain the typical interpretation as dispersion parameters relative to a random-intercept Poisson model.

Below we present two short examples of mixed-effects negative binomial regression; refer to [\[ME\] me](#) and [\[ME\] meglm](#) for more examples including crossed-effects models.

➤ **Example 1: Two-level random-intercept model**

[Rabe-Hesketh and Skrondal \(2012, chap. 13.7\)](#) analyze the data from [Winkelmann \(2004\)](#) on the impact of the 1997 health reform in Germany on the number of doctor visits. The intent of policymakers was to reduce government expenditures on healthcare. We use a subsample of the data restricted to 1,158 women who were employed full time the year before or after the reform.

```
. use http://www.stata-press.com/data/r15/drvisits
. describe
Contains data from http://www.stata-press.com/data/r15/drvisits.dta
  obs:      2,227
  vars:      8
  size:     71,264
                                23 Jan 2016 18:39
```

variable name	storage type	display format	value label	variable label
id	float	%9.0g		person id
numvisit	float	%9.0g		number of doctor visits in the last 3 months before interview
age	float	%9.0g		age in years
educ	float	%9.0g		education in years
married	float	%9.0g		=1 if married, 0 otherwise
badh	float	%9.0g		self-reported health status, =1 if bad
loginc	float	%9.0g		log of household income
reform	float	%9.0g		=0 if interview before reform, =1 if interview after reform

Sorted by:

The dependent variable, `numvisit`, is a count of doctor visits. The covariate of interest is a dummy variable, `reform`, which indicates whether a doctor visit took place before or after the reform. Other covariates include a self-reported health status, age, education, marital status, and a log of household income.

We first fit a two-level random-intercept Poisson model. We specify the random intercept at the `id` level, that is, an individual-person level.

```
. mepoisson numvisit reform age educ married badh loginc || id:, irr
```

Fitting fixed-effects model:

```
Iteration 0:  log likelihood = -9326.8542
Iteration 1:  log likelihood = -5989.7308
Iteration 2:  log likelihood = -5942.7581
Iteration 3:  log likelihood = -5942.7243
Iteration 4:  log likelihood = -5942.7243
```

Refining starting values:

```
Grid node 0:  log likelihood = -4761.1257
```

Fitting full model:

```
Iteration 0:  log likelihood = -4761.1257
Iteration 1:  log likelihood = -4683.2239
Iteration 2:  log likelihood = -4646.9329
Iteration 3:  log likelihood = -4645.736
Iteration 4:  log likelihood = -4645.7371
Iteration 5:  log likelihood = -4645.7371
```

Mixed-effects Poisson regression

Group variable:	id	Number of obs	=	2,227
		Number of groups	=	1,518
		Obs per group:		
		min	=	1
		avg	=	1.5
		max	=	2

Integration method: mvaghermite

	Integration pts.	=	7
	Wald chi2(6)	=	249.37
	Prob > chi2	=	0.0000

Log likelihood = -4645.7371

numvisit	IRR	Std. Err.	z	P> z	[95% Conf. Interval]
reform	.9517026	.0309352	-1.52	0.128	.8929617 1.014308
age	1.005821	.002817	2.07	0.038	1.000315 1.011357
educ	1.008788	.0127394	0.69	0.488	.9841258 1.034068
married	1.082078	.0596331	1.43	0.152	.9712905 1.205503
badh	2.471857	.151841	14.73	0.000	2.191471 2.788116
loginc	1.094144	.0743018	1.32	0.185	.9577909 1.249909
_cons	.5216748	.2668604	-1.27	0.203	.191413 1.421766

id	var(_cons)		
	.8177932	.0503902	.724761 .9227673

Note: Estimates are transformed only in the first equation.

Note: _cons estimates baseline incidence rate (conditional on zero random effects).

LR test vs. Poisson model: `chibar2(01) = 2593.97` `Prob >= chibar2 = 0.0000`

```
. estimates store mepoisson
```

Because we specified the `irr` option, the parameters are reported as incidence-rate ratios. The healthcare reform seems to reduce the expected number of visits by 5% but without statistical significance.

Because we have only one random effect at the `id` level, the table shows only one variance component. The estimate of σ_u^2 is 0.82 with standard error 0.05. The reported likelihood-ratio test shows that there is enough variability between women to favor a mixed-effects Poisson regression over a standard Poisson regression; see [Distribution theory for likelihood-ratio test](#) in [ME] **me** for a discussion of likelihood-ratio testing of variance components.

It is possible that after conditioning on the person-level random effect, the counts of doctor visits are overdispersed. For example, medical problems occurring during the time period leading to the survey can result in extra doctor visits. We thus reexamine the data with `menbreg`.

```
. menbreg numvisit reform age educ married badh loginc || id:, irr
Fitting fixed-effects model:
Iteration 0:   log likelihood = -4610.7165
Iteration 1:   log likelihood = -4563.4682
Iteration 2:   log likelihood = -4562.3241
Iteration 3:   log likelihood = -4562.3238
Refining starting values:
Grid node 0:   log likelihood = -4643.5216
Fitting full model:
Iteration 0:   log likelihood = -4643.5216   (not concave)
Iteration 1:   log likelihood = -4555.961
Iteration 2:   log likelihood = -4518.7353
Iteration 3:   log likelihood = -4513.1951
Iteration 4:   log likelihood = -4513.1853
Iteration 5:   log likelihood = -4513.1853
Mixed-effects nbinomial regression
Overdispersion:      mean
Group variable:      id
Number of obs       =      2,227
Number of groups    =      1,518
Obs per group:
    min =           1
    avg =           1.5
    max =           2
Integration method:  mvaghermite
Integration pts.    =           7
Wald chi2(6)        =      237.35
Prob > chi2         =      0.0000
Log likelihood = -4513.1853
```

numvisit	IRR	Std. Err.	z	P> z	[95% Conf. Interval]	
reform	.9008536	.042022	-2.24	0.025	.8221449	.9870975
age	1.003593	.0028206	1.28	0.202	.9980799	1.009137
educ	1.007026	.012827	0.55	0.583	.9821969	1.032483
married	1.089597	.064213	1.46	0.145	.970738	1.223008
badh	3.043562	.2366182	14.32	0.000	2.613404	3.544523
loginc	1.136342	.0867148	1.67	0.094	.9784833	1.319668
_cons	.5017199	.285146	-1.21	0.225	.1646994	1.528377
/lnalpha	-.7962692	.1190614			-1.029625	-.5629132
id						
var(_cons)	.4740088	.0582404			.3725642	.6030754

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline incidence rate (conditional on zero random effects).
LR test vs. nbinomial model: $\chi^2(01) = 98.28$ Prob $\geq \chi^2 = 0.0000$

The estimated effect of the healthcare reform now corresponds to the reduction in the number of doctor visits by 10%—twice as much compared with the Poisson model—and this effect is significant at the 5% level.

The estimate of the variance component σ_u^2 drops down to 0.47 compared with `mepoisson`, which is not surprising given that now we have an additional parameter that controls the variability of the data.

Because the conditional overdispersion α is assumed to be greater than 0, it is parameterized on the log scale, and its log estimate is reported as `/lnalpha` in the output. In our model, $\hat{\alpha} = \exp(-0.80) = 0.45$. We can also compute the unconditional overdispersion in this model by using $\exp(.47) \times (1 + .45) - 1 = 1.32$.

The reported likelihood-ratio test shows that there is enough variability between women to favor a mixed-effects negative binomial regression over negative binomial regression without random effects.

We can also perform a likelihood-ratio test comparing the mixed-effects negative binomial model to the mixed-effects Poisson model. Because we are comparing two different estimators, we need to use the `force` option with `lrtest`. In general, there is no guarantee as to the validity or interpretability of the resulting likelihood-ratio test, but in our case we know the test is valid because the mixed-effects Poisson model is nested within the mixed-effects negative binomial model.

```
. lrtest mepoisson ., force
Likelihood-ratio test                                LR chi2(1) =    265.10
(Assumption: mepoisson nested in .)                 Prob > chi2 =    0.0000
Note: The reported degrees of freedom assumes the null hypothesis is not on
the boundary of the parameter space. If this is not true, then the
reported test is conservative.
```

The reported likelihood-ratio test favors the mixed-effects negative binomial model. The reported test is conservative because the test of $H_0: \alpha = 0$ occurs on the boundary of the parameter space; see [Distribution theory for likelihood-ratio test](#) in [ME] [me](#) for details.

◀

The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by `||`. The order of nesting goes from left to right as the groups go from biggest (highest level) to smallest (lowest level). To demonstrate a three-level model, we revisit [example 3](#) from [ME] [mepoisson](#).

► Example 2: Three-level random-intercept model

[Rabe-Hesketh and Skrondal \(2012, exercise 13.7\)](#) describe data from the *Atlas of Cancer Mortality in the European Economic Community* (EEC) ([Smans, Mair, and Boyle 1993](#)). The data were analyzed in [Langford, Bentham, and McDonald \(1998\)](#) and record the number of deaths among males due to malignant melanoma during 1971–1980.

```
. use http://www.stata-press.com/data/r15/melanoma
(Skin cancer (melanoma) data)
. describe
Contains data from http://www.stata-press.com/data/r15/melanoma.dta
   obs:                354                Skin cancer (melanoma) data
  vars:                  6                30 May 2016 17:10
 size:                4,956              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
nation	byte	%11.0g	n	Nation ID
region	byte	%9.0g		Region ID: EEC level-I areas
county	int	%9.0g		County ID: EEC level-II/level-III areas
deaths	int	%9.0g		No. deaths during 1971–1980
expected	float	%9.0g		No. expected deaths
uv	float	%9.0g		UV dose, mean-centered

Sorted by:

Nine European nations (variable `nation`) are represented, and data were collected over geographical regions defined by EEC statistical services as level I areas (variable `region`), with deaths being recorded for each of 354 counties, which are level II or level III EEC-defined areas (variable `county`, which identifies the observations). Counties are nested within regions, and regions are nested within nations.

The variable `deaths` records the number of deaths for each county, and `expected` records the expected number of deaths (the exposure) on the basis of crude rates for the combined countries. The variable `uv` is a measure of exposure to ultraviolet (UV) radiation.

In [example 3](#) of [\[ME\] mepoisson](#), we noted that because counties also identified the observations, we could model overdispersion by using a four-level Poisson model with a random intercept at the county level. Here we fit a three-level negative binomial model with the default mean-dispersion parameterization.

```
. menbreg deaths uv, exposure(expected) || nation: || region:
Fitting fixed-effects model:
Iteration 0:   log likelihood =  -1361.855
Iteration 1:   log likelihood = -1230.0211
Iteration 2:   log likelihood = -1211.049
Iteration 3:   log likelihood = -1202.5641
Iteration 4:   log likelihood = -1202.5329
Iteration 5:   log likelihood = -1202.5329
Refining starting values:
Grid node 0:   log likelihood = -1209.6951
Fitting full model:
Iteration 0:   log likelihood = -1209.6951  (not concave)
(output omitted)
Iteration 11:  log likelihood = -1086.3902
Mixed-effects nbinoomial regression          Number of obs      =          354
Overdispersion:              mean
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
nation	9	3	39.3	95
region	78	1	4.5	13

```
Integration method: mvaghermite          Integration pts. =          7
Wald chi2(1) =          8.73
Log likelihood = -1086.3902              Prob > chi2 =          0.0031
```

deaths	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
uv	-.0335933	.0113725	-2.95	0.003	-.055883	-.0113035
_cons	-.0790606	.1295931	-0.61	0.542	-.3330583	.1749372
ln(expected)	1	(exposure)				
/lnalpha	-4.182603	.3415036			-4.851937	-3.513268
nation						
var(_cons)	.1283614	.0678971			.0455187	.3619758
nation>						
region						
var(_cons)	.0401818	.0104855			.0240938	.067012

```
LR test vs. nbinoomial model: chi2(2) = 232.29          Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

The estimates are very close to those of `mepoisson`. The conditional overdispersion in our model is $\hat{\alpha} = \exp(-4.18) = 0.0153$. It is in agreement with the estimate of the random intercept at the county level, 0.0146, in a four-level random-effects Poisson model reported by `mepoisson`. Because the negative binomial is a three-level model, we gained some computational efficiency over the four-level Poisson model.



Stored results

`menbreg` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>menbreg</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	<code>nbreg</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>log</code>
<code>e(family)</code>	<code>nbomial</code>
<code>e(clustvar)</code>	name of cluster variable
<code>e(dispersion)</code>	mean or constant
<code>e(offset)</code>	<code>offset</code>
<code>e(exposure)</code>	exposure variable
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.

<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement estat
<code>e(predict)</code>	program used to implement predict
<code>e(marginsnotok)</code>	predictions disallowed by margins
<code>e(marginswtype)</code>	weight type for margins
<code>e(marginswexp)</code>	weight expression for margins
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as asbalanced
<code>e(asobserved)</code>	factor variables <code>fvset</code> as asobserved

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

`menbreg` is a convenience command for `meglm` with a log link and an `nbinomial` family; see [\[ME\] `meglm`](#).

Without a loss of generality, consider a two-level negative binomial model. For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$ and the conditional overdispersion parameter α in a mean-overdispersion parameterization, is

$$\begin{aligned}
 f(\mathbf{y}_j | \mathbf{u}_j, \alpha) &= \prod_{i=1}^{n_j} \left\{ \frac{\Gamma(y_{ij} + r)}{\Gamma(y_{ij} + 1)\Gamma(r)} p_{ij}^r (1 - p_{ij})^{y_{ij}} \right\} \\
 &= \exp \left[\sum_{i=1}^{n_j} \{ \log \Gamma(y_{ij} + r) - \log \Gamma(y_{ij} + 1) - \log \Gamma(r) + c(y_{ij}, \alpha) \} \right]
 \end{aligned}$$

where $c(y_{ij}, \alpha)$ is defined as

$$-\frac{1}{\alpha} \log\{1 + \exp(\eta_{ij} + \log \alpha)\} - y_{ij} \log\{1 + \exp(-\eta_{ij} - \log \alpha)\}$$

and $r = 1/\alpha$, $p_{ij} = 1/(1 + \alpha\mu_{ij})$, and $\eta_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j$.

For the constant-overdispersion parameterization with the conditional overdispersion parameter δ , the conditional distribution of $\mathbf{y}_j$ is

$$\begin{aligned} f(\mathbf{y}_j|\mathbf{u}_j, \delta) &= \prod_{i=1}^{n_j} \left\{ \frac{\Gamma(y_{ij} + r_{ij})}{\Gamma(y_{ij} + 1)\Gamma(r_{ij})} p^{r_{ij}} (1-p)^{y_{ij}} \right\} \\ &= \exp \left[\sum_{i=1}^{n_j} \{ \log \Gamma(y_{ij} + r_{ij}) - \log \Gamma(y_{ij} + 1) - \log \Gamma(r_{ij}) + c(y_{ij}, \delta) \} \right] \end{aligned}$$

where $c(y_{ij}, \delta)$ is defined as

$$- \left(\frac{\mu_{ij}}{\delta} + y_{ij} \right) \log(1 + \delta) + y_{ij} \log \delta$$

and $r_{ij} = \mu_{ij}/\delta$ and $p = 1/(1 + \delta)$.

For conciseness, let γ denote either conditional overdispersion parameter. Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ , the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j, \gamma)$,

$$\begin{aligned} \mathcal{L}_j(\beta, \Sigma, \gamma) &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int f(\mathbf{y}_j|\mathbf{u}_j, \gamma) \exp(-\mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int \exp\{h(\beta, \Sigma, \mathbf{u}_j, \gamma)\} d\mathbf{u}_j \end{aligned} \quad (4)$$

where

$$h(\beta, \Sigma, \mathbf{u}_j, \gamma) = f(\mathbf{y}_j|\mathbf{u}_j, \gamma) - \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (4) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

menbreg supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

References

- Langford, I. H., G. Benthams, and A. McDonald. 1998. Multi-level modelling of geographically aggregated health data: A case study on malignant melanoma mortality and UV exposure in the European community. *Statistics in Medicine* 17: 41–57.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.
- Smans, M., C. S. Mair, and P. Boyle. 1993. *Atlas of Cancer Mortality in the European Economic Community*. Lyon, France: IARC Scientific Publications.
- Winkelmann, R. 2004. Health care reform and the number of doctor visits—An econometric analysis. *Journal of Applied Econometrics* 19: 455–472.

Also see

[ME] **menbreg postestimation** — Postestimation tools for menbreg

[ME] **mepoisson** — Multilevel mixed-effects Poisson regression

[ME] **meqrpoisson** — Multilevel mixed-effects Poisson regression (QR decomposition)

[ME] **me** — Introduction to multilevel mixed-effects models

[BAYES] **bayes: menbreg** — Bayesian multilevel negative binomial regression

[SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)

[SVY] **svy estimation** — Estimation commands for survey data

[XT] **xtnbreg** — Fixed-effects, random-effects, & population-averaged negative binomial models

[U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation command is of special interest after `menbreg`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike’s and Schwarz’s Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman’s specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
<code>mu</code>	mean response; the default
<code>eta</code>	fitted linear predictor
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code>density</code>	predicted density function
<code>distribution</code>	predicted distribution function
<code>pearson</code>	Pearson residuals
<code>deviance</code>	deviance residuals
<code>anscombe</code>	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

options	Description
Main	
<code>conditional(ctype)</code>	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<code>marginal</code>	compute <i>statistic</i> marginally with respect to the random effects
<code>nooffset</code>	make calculation ignoring offset or exposure
Integration	
<code>int_options</code>	integration options
pearson, deviance, anscombe may not be combined with <code>marginal</code> .	
ctype	Description
<code>ebmeans</code>	empirical Bayes means of random effects; the default
<code>ebmodes</code>	empirical Bayes modes of random effects
<code>fixedonly</code>	prediction for the fixed portion of the model only
re_options	Description
Main	
<code>ebmeans</code>	use empirical Bayes means of random effects; the default
<code>ebmodes</code>	use empirical Bayes modes of random effects
<code>reses(stub* newvarlist)</code>	calculate standard errors of empirical Bayes estimates
Integration	
<code>int_options</code>	integration options
int_options	Description
<code>intpoints(#)</code>	use # quadrature points to compute marginal predictions and empirical Bayes means
<code>iterate(#)</code>	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<code>tolerance(#)</code>	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main
mu, the default, calculates the predicted mean, that is, the predicted number of events.
eta, xb, stdp, density, distribution, pearson, deviance, anscombe, scores, conditional(), marginal, and nooffset; see [ME] meglm postestimation .
reffects, ebmeans, ebmodes, and reses(); see [ME] meglm postestimation .
Integration
intpoints(), iterate(), and tolerance(); see [ME] meglm postestimation .

margins

Description for margins

`margins` estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
density	not allowed with margins
distribution	not allowed with margins
pearson	not allowed with margins
deviance	not allowed with margins
anscombe	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options `conditional(ebmeans)` and `conditional(ebmodes)` are not allowed with `margins`.
Option `marginal` is assumed where applicable if `conditional(fixedonly)` is not specified.

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.
For the full syntax, see [\[R\] margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects negative binomial model with `menbreg`. For the most part, calculation centers around obtaining estimates of the subject/group-specific random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation.

Here we show a short example of predicted counts and predicted random effects; refer to [\[ME\] meglm postestimation](#) for additional examples applicable to mixed-effects generalized linear models.

► Example 1: Predicting counts and random effects

In [example 2](#) of [\[ME\] menbreg](#), we modeled the number of deaths among males in nine European nations as a function of exposure to ultraviolet radiation (uv). We used a three-level negative binomial model with random effects at the nation and region levels.

```
. use http://www.stata-press.com/data/r15/melanoma
(Skin cancer (melanoma) data)
. menbreg deaths uv, exposure(expected) || nation: || region:
(output omitted)
```

We can use `predict` to obtain the predicted counts as well as the estimates of the random effects at the nation and region levels.

```
. predict mu
(predictions based on fixed effects and posterior means of random effects)
(option mu assumed)
(using 7 quadrature points)
. predict re_nat re_reg, reffects
(calculating posterior means of random effects)
(using 7 quadrature points)
```

Stata displays a note that the predicted values of `mu` are based on the posterior means of random effects. You can use option `modes` to obtain predictions based on the posterior modes of random effects.

Here we list the data for the first nation in the dataset, which happens to be Belgium:

```
. list nation region deaths mu re_nat re_reg if nation==1, sepby(region)
```

	nation	region	deaths	mu	re_nat	re_reg
1.	Belgium	1	79	64.4892	-.0819939	.2937711
2.	Belgium	2	80	77.64736	-.0819939	.024005
3.	Belgium	2	51	44.56528	-.0819939	.024005
4.	Belgium	2	43	53.10434	-.0819939	.024005
5.	Belgium	2	89	65.35963	-.0819939	.024005
6.	Belgium	2	19	35.18457	-.0819939	.024005
7.	Belgium	3	19	8.770186	-.0819939	-.3434432
8.	Belgium	3	15	43.95521	-.0819939	-.3434432
9.	Belgium	3	33	34.17878	-.0819939	-.3434432
10.	Belgium	3	9	7.332448	-.0819939	-.3434432
11.	Belgium	3	12	12.93873	-.0819939	-.3434432

We can see that the predicted random effects at the nation level, `re_nat`, are the same for all the observations. Similarly, the predicted random effects at the region level, `re_reg`, are the same within each region.

◀

Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [\[ME\] meglm postestimation](#).

Also see

[ME] [menbreg](#) — Multilevel mixed-effects negative binomial regression

[ME] [meglm postestimation](#) — Postestimation tools for meglm

[U] [20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`menl` fits nonlinear mixed-effects models in which some or all fixed and random effects enter nonlinearly. These models are also known as multilevel nonlinear models or hierarchical nonlinear models. The overall error distribution of the nonlinear mixed-effects model is assumed to be Gaussian. Different covariance structures are provided to model random effects and to model heteroskedasticity and correlations within lowest-level groups.

Quick start

Nonlinear mixed-effects regression of y on x_1 and x_2 with random intercepts B_0 by id

```
menl y = {a}*(1-exp(-(b0+{b1}*x1+{b2}*x2+{B0[id]})))
```

As above, but using the more efficient specification of the linear combination

```
menl y = {a}*(1-exp(-{xb: x1 x2 B0[id]}))
```

As above, but using `define()` to specify the linear combination

```
menl y = {a}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id])
```

As above, but perform restricted maximum-likelihood estimation instead of the default maximum-likelihood estimation

```
menl y = {a}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id]) reml
```

Specify your own initial values for fixed effects, but use the default expectation-maximization (EM) method to obtain initial values for random-effects parameters

```
menl y = {a}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id])          ///
initial({a} 1 {xb:x1} 1 {xb:x2} 0.5 {xb:_cons} 2, fixed)
```

Include random intercepts A_0 by id to allow parameter a to vary between levels of id and specify the xb suboption to indicate that a contains a linear combination rather than a scalar parameter

```
menl y = {a:}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id])          ///
define(a: A0[id], xb)
```

Include a random slope on continuous variable x_2 in the linear combination, and allow correlation between random slopes B_1 and intercepts B_0

```
menl y = {a}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id] c.x2#B1[id]) ///
covariance(B0 B1, unstructured)
```

Specify a heteroskedastic within-subject error variance that varies as a power of x_2

```
menl y = {a}*(1-exp(-{xb:})), define(xb: x1 x2 B0[id] c.x2#B1[id]) ///
covariance(B0 B1, unstructured) resvariance(power x2)
```

Display random-effects and within-group error parameters as standard deviations and correlations

```
menl, stddeviations
```

Three-level nonlinear regression of y on variable `time` and factor variable `f` with random intercepts `S0` by `lev3` and `W0` by `lev2` nested within `lev3`, using an AR(1) correlation structure for the residuals

```
menl y = {phi1:}+{phi2:}*exp(-{phi3}*time), ///  
      define(phi1: i.f S0[lev3]) define(phi2: i.f W0[lev3>lev2]) ///  
      rescorrelation(ar 1, t(time))
```

Three-level nonlinear regression of y on x_1 with random intercepts `W0` and slopes `W1` on continuous x_1 by `lev3` and with random intercepts `S0` and slopes `S1` on x_1 by `lev2` nested within `lev3`, using unstructured covariance for `W0` and `W1` and exchangeable covariance for `S0` and `S1`

```
menl y = {phi1:}+{b1}*cos({b2}*x1), ///  
      define(phi1:x1 W0[lev3] S0[lev3>lev2] ///  
            c.x1#(W1[lev3] S1[lev3>lev2])) ///  
      covariance(W0 W1, unstructured) ///  
      covariance(S0 S1, exchangeable) ///
```

As above, but assume that residuals are independent but have different variances for males and females

```
menl y = {phi1:}+{b1}*cos({b2}*x1), ///  
      define(phi1:x1 W0[lev3] S0[lev3>lev2] ///  
            c.x1#(W1[lev3] S1[lev3>lev2])) ///  
      covariance(W0 W1, unstructured) ///  
      covariance(S0 S1, exchangeable) ///  
      rescovariance(identity, by(female))
```

Menu

Statistics > Multilevel mixed-effects models > Nonlinear regression

Syntax

```
menl depvar = <menexpr> [if] [in] [, options]
```

<menexpr> defines a nonlinear regression function as a substitutable expression that contains model parameters and random effects specified in braces {}, as in `exp({b}+{U[id]})`; see [Random-effects substitutable expressions](#) for details.

<i>options</i>	Description
Model	
<u>m</u> le	fit model via maximum likelihood; the default
<u>r</u> eml	fit model via restricted maximum likelihood
<u>d</u> efine(<i>name</i> :<resubexpr>)	define a function of model parameters; this option may be repeated
<u>c</u> ovariance(<i>covspec</i>)	variance–covariance structure of the random effects; this option may be repeated
<u>i</u> nitial(<i>initial_values</i>)	initial values for parameters
Residuals	
<u>r</u> escovariance(<i>rescovspec</i>)	covariance structure for within-group errors
<u>r</u> esvariance(<i>resvarspec</i>)	heteroskedastic variance structure for within-group errors
<u>r</u> escorrelation(<i>rescorrspec</i>)	correlation structure for within-group errors
Reporting	
<u>l</u> evel(#)	set confidence level; default is <code>level(95)</code>
<u>v</u> ariance	show random-effects and within-group error parameter estimates as variances and covariances; the default
<u>s</u> tddeviations	show random-effects and within-group error parameter estimates as standard deviations and correlations
<u>n</u> oretable	suppress random-effects table
<u>n</u> ofetable	suppress fixed-effects table
<u>e</u> stmetric	show parameter estimates in the estimation metric
<u>n</u> olegend	suppress table expression legend
<u>n</u> oheader	suppress output header
<u>n</u> ogroup	suppress table summarizing groups
<u>n</u> ostderr	do not estimate standard errors of random-effects parameters
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
EM options	
<u>e</u> miterate(#)	number of EM iterations; default is <code>emiterate(25)</code>
<u>e</u> mtolerance(#)	EM convergence tolerance; default is <code>emtolerance(1e-10)</code>
<u>e</u> mlog	show EM iteration log
Maximization	
<i>menlmaxopts</i>	control the maximization process
<u>c</u> oeflegend	display legend instead of statistics

`coeflegend` does not appear in the dialog box.

See [\[U\] 20 Estimation and postestimation commands](#) for more capabilities of estimation commands.

The syntax of *covspec* is

```
rename1 rename2 [...], vartype
```

<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect; all covariances are 0; the default
<u>exchangeable</u>	equal variances for random effects and one common pairwise covariance
<u>identity</u>	equal variances for random effects; all covariances are 0
<u>unstructured</u>	all variances and covariances to be distinctly estimated

The syntax of *rescovspec* is

```
rescov [, rescovopts]
```

<i>rescov</i>	Description
<u>identity</u>	uncorrelated within-group errors with one common variance; the default
<u>independent</u>	uncorrelated within-group errors with distinct variances
<u>ar</u> [#]	within-group errors with autoregressive (AR) structure of order #, AR(#); ar 1 is implied by ar
<u>ma</u> [#]	within-group errors with moving-average (MA) structure of order #, MA(#); ma 1 is implied by ma
<u>ctar1</u>	within-group errors with continuous-time AR(1) structure

The syntax of *resvarspec* is

```
resvarfunc [, resvaropts]
```

<i>resvarfunc</i>	Description
<u>identity</u>	equal within-group error variances; the default
<u>linear</u> <i>varname</i>	within-group error variance varies linearly with <i>varname</i>
<u>power</u> <i>varname</i> _yhat	variance function is a power of <i>varname</i> or of predicted mean
<u>exponential</u> <i>varname</i> _yhat	variance function is exponential of <i>varname</i> or of predicted mean

The syntax of *rescorrspec* is

```
rescorr [, rescorropts]
```

<i>rescorr</i>	Description
<u>identity</u>	uncorrelated within-group errors; the default
<u>ar</u> [#]	within-group errors with AR(#) structure; ar 1 is implied by ar
<u>ma</u> [#]	within-group errors with MA(#) structure; ma 1 is implied by ma
<u>ctar1</u>	within-group errors with continuous-time AR(1) structure

Options

Model

`mle` and `reml` specify the statistical method for fitting the model.

`mle`, the default, specifies that the model be fit using maximum likelihood (ML).

`reml` specifies that the model be fit using restricted maximum likelihood (REML), also known as residual maximum likelihood.

`define(name:<resubexpr>)` defines a function of model parameters, *<resubexpr>*, and labels it as *name*. This option can be repeated to define multiple functions. The `define()` option is useful for expressions that appear multiple times in the main nonlinear specification *menlexpr*: you define the expression once and then simply refer to it by using `{name:}` in the nonlinear specification. This option can also be used for notational convenience. See *Random-effects substitutable expressions* for how to specify *<resubexpr>*.

`covariance(rename1 rename2 [...], vartype)` specifies the structure of the covariance matrix for the random effects. *rename1*, *rename2*, and so on, are the names of the random effects to be correlated (see *Random effects*), and *vartype* is one of the following: `independent`, `exchangeable`, `identity`, or `unstructured`. Instead of *renames*, you can specify *restub** to refer to random effects that share the same *restub* in their names.

`independent` allows for a distinct variance for each random effect and assumes that all covariances are 0; the default.

`exchangeable` specifies one common variance for all random effects and one common pairwise covariance.

`identity` is short for “multiple of the identity”; that is, all variances are equal, and all covariances are 0.

`unstructured` allows for all variances and covariances to be distinct. If *p* random effects are specified, the unstructured covariance matrix will have $p(p+1)/2$ unique parameters.

`initial(initial_values)` specifies the initial values for model parameters. You can specify a $1 \times k$ matrix, where *k* is the total number of parameters in the model, or you can specify a parameter name, its initial value, another parameter name, its initial value, and so on. For example, to initialize `{alpha}` to 1.23 and `{delta}` to 4.57, you would type

```
. menl ..., initial(alpha 1.23 delta 4.57) ...
```

To initialize multiple parameters that have the same group name, for example, `{y:x1}` and `{y:x2}`, with the same initial value, you can simply type

```
. menl ..., initial({y:} 1) ...
```

For the full specification, see *Specifying initial values*.

Residuals

`menl` provides two ways to model the within-group error covariance structure, sometimes also referred to as residual covariance structure in the literature. You can model the covariance directly by using the `rescovariance()` option or indirectly by using the `resvariance()` and `rescorrelation()` options.

`rescovariance(rescov [, rescovopts])` specifies the *within-group errors* covariance structure or covariance structure of the residuals within the *lowest-level group* of the nonlinear mixed-effects model. For example, if you are modeling random effects for classes nested within schools, then

`rescovariance()` refers to the residual variance–covariance structure of the observations within classes, the lowest-level groups.

rescov is one of the following: `identity`, `independent`, `ar` `[#]`, `ma` `[#]`, or `ctar1`. Below, we describe each *rescov* with its specific options *rescovopts*:

`identity` `[, by(byvar)]`, the default, specifies that all within-group errors be independent and identically distributed (i.i.d.) with one common error variance σ_ϵ^2 . When combined with `by(byvar)`, independence is still assumed, but you estimate a distinct variance for each category of *byvar*.

`independent`, `index(varname)` specifies that within-group errors are independent with distinct variances for each value (`index`) of *varname*. `index(varname)` is required.

`ar` `[#]`, `t(timevar)` `[by(byvar)]` assumes that within-group errors have an AR(`#`) structure. If `#` is omitted, `ar 1` is assumed. `t(timevar)` is required.

`ma` `[#]`, `t(timevar)` `[by(byvar)]` assumes that within-group errors have an MA(`#`) structure. If `#` is omitted, `ma 1` is assumed. `t(timevar)` is required.

`ctar1`, `t(timevar)` `[by(byvar)]` assumes that within-group errors have a continuous-time AR(1) structure. This is a generalization of the AR covariance structure that allows for unequally spaced and noninteger time values. `t(timevar)` is required.

rescovopts are `index(varname)`, `t(timevar)`, and `by(byvar)`.

`index(varname)` is used within the `rescovariance()` option with *rescov* `independent`. *varname* is a nonnegative-integer-valued variable that identifies the observations within the lowest-level groups (for example, `obsid`). The groups may be unbalanced in that different groups may have different `index()` values, but you may not have repeated `index()` values within any particular group.

`t(timevar)` is used within the `rescovariance()` option to specify a time variable for the `ar`, `ma`, and `ctar1` structures.

With *rescov* `ar` and `ma`, *timevar* is an integer-valued time variable used to order the observations within the lowest-level groups and to determine the lags between successive observations. Any nonconsecutive time values will be treated as gaps. For the `ar` or `ma` structure, `# + 1` parameters are estimated: `#` AR or `#` MA coefficients and one overall error variance σ_ϵ^2 .

With *rescov* `ctar1`, *timevar* is a real-valued time variable. The correlation between two error terms is the parameter ρ , raised to a power equal to the absolute value of the difference between the `t()` values for those errors. For the `ctar1` structure, two parameters are estimated: the correlation parameter ρ and one overall error variance σ_ϵ^2 .

`by(byvar)` is for use within the `rescovariance()` option and specifies that a set of distinct within-group error covariance parameters be estimated for each category of *byvar*. In other words, you can use `by()` to model heteroskedasticity. *byvar* must be nonnegative-integer valued and constant within the lowest-level groups.

`resvariance(resvarfunc` `[, resvaropts]`) specifies a heteroskedastic variance structure of the within-group errors. It can be used with the `rescorrelation()` option to specify flexible within-group error covariance structures. The heteroskedastic variance structure is modeled as $\text{Var}(\epsilon_{ij}) = \sigma^2 g^2(\delta, v_{ij})$, where σ is an unknown scale parameter, $g(\cdot)$ is a function that models heteroskedasticity (also known as variance function in the literature), δ is a vector of unknown parameters of the variance function, and v_{ij} 's are the values of a fixed covariate x_{ij} or of the predicted mean $\hat{\mu}_{ij}$.

resvarfunc, omitting the arguments, is one of the following: *identity*, *linear*, *power*, or *exponential*, and *resvaropts* are options specific to each variance function.

identity, the default, specifies a homoskedastic variance structure for the within-group errors; $g(\delta, v_{ij}) = 1$, so that $\text{Var}(\epsilon_{ij}) = \sigma^2 = \sigma_\epsilon^2$.

linear varname specifies that the within-group error variance vary linearly with *varname*; that is, $g(\delta, v_{ij}) = \sqrt{\text{varname}_{ij}}$, so that $\text{Var}(\epsilon_{ij}) = \sigma^2 \text{varname}_{ij}$. *varname* must be positive.

power varname | *_yhat* [, *strata(stratavar)* *noconstant*] specifies that the within-group error variance, or more precisely the variance function, be expressed in terms of a power of either *varname* or the predicted mean *_yhat*, plus a constant term; $g(\delta, v_{ij}) = |v_{ij}|^{\delta_1} + \delta_2$. If *noconstant* is specified, the constant term δ_2 is suppressed. In general, three parameters are estimated: a scale parameter σ , the power δ_1 , and the constant term δ_2 . When *strata(stratavar)* is specified, the power and constant parameters (but not the scale) are distinctly estimated for each stratum. A total number of $2L + 1$ parameters are estimated (L power parameters, L constant parameters, and scale σ), where L is the number of strata defined by variable *stratavar*.

exponential varname | *_yhat* [, *strata(stratavar)*] specifies that the within-group error variance vary exponentially with *varname* or with the predicted mean *_yhat*; $g(\gamma, v_{ij}) = \exp(\gamma v_{ij})$. Two parameters are estimated: a scale parameter σ and an exponential parameter γ . When *strata(stratavar)* is specified, the exponential parameter γ (but not scale σ) is distinctly estimated for each stratum. A total number of $L + 1$ parameters are estimated (L exponential parameters and scale σ), where L is the number of strata defined by variable *stratavar*.

resvaropts are *strata(stratavar)* and *noconstant*.

strata(stratavar) is used within the *resvariance()* option with *resvarfunc* *power* and *exponential*. *strata()* specifies that the parameters of the variance function $g(\cdot)$ be distinctly estimated for each stratum. The scale parameter σ remains constant across strata. In contrast, *rescovariance()*'s *by(byvar)* suboption specifies that all covariance parameters, including σ (whenever applicable), be estimated distinctly for each category of *byvar*. *stratavar* must be nonnegative-integer valued and constant within the lowest-level groups.

noconstant is used within the *resvariance()* option with *resvarfunc* *power*. *noconstant* specifies that the constant parameter be suppressed in the expression of the variance function $g(\cdot)$.

rescorrelation(rescorr [, *rescoropts*]) specifies a correlation structure of the within-group errors. It can be used with the *resvariance()* option to specify flexible within-group error covariance structures.

rescorr is one of the following: *identity*, *ar* [#], *ma* [#], or *ctar1*.

identity, the default, specifies that all within-group error correlations be zeros.

ar [#], *t(timevar)* [*by(byvar)*] assumes that within-group errors have an AR(#) correlation structure. If # is omitted, *ar* 1 is assumed. The *t(timevar)* option is required.

ma [#], *t(timevar)* [*by(byvar)*] assumes that within-group errors have an MA(#) correlation structure. If # is omitted, *ma* 1 is assumed. The *t(timevar)* option is required.

ctar1, *t(timevar)* [*by(byvar)*] assumes that within-group errors have a continuous-time AR(1) correlation structure. The *t(timevar)* option is required.

rescorropts are `t(timevar)` and `by(byvar)`.

`t(timevar)` is used within the `rescorrelation()` option to specify a time variable for the `ar`, `ma`, and `ctar1` structures.

With *rescorr ar* and *ma*, *timevar* is an integer-valued time variable used to order the observations within the lowest-level groups and to determine the lags between successive observations. Any nonconsecutive time values will be treated as gaps. For the *ar* or *ma* structure, # AR or MA coefficients are estimated.

With *rescorr ctar1*, *timevar* is a real-valued time variable. For the continuous AR(1) correlation model, the correlation between two errors is the parameter ρ , raised to a power equal to the absolute value of the difference between the `t()` values for those errors.

`by(byvar)` is used within the `rescorrelation()` option and specifies that a set of distinct within-group error correlation parameters be estimated for each category of *byvar*. *byvar* must be nonnegative-integer valued and constant within the lowest-level groups.

Reporting

`level(#)`; see [R] [estimation options](#).

`variance`, the default, displays the random-effects and within-group error parameter estimates as variances and covariances.

`stddeviations` displays the random-effects and within-group error parameter estimates as standard deviations and correlations.

`noretable` suppresses the random-effects table from the output.

`nofetable` suppresses the fixed-effects table from the output.

`estmetric` displays all parameter estimates in the estimation metric. Fixed-effects estimates are unchanged from those normally displayed, but random-effects parameter estimates are displayed as log-standard deviations and hyperbolic arctangents of correlations. Within-group error parameter estimates are also displayed in their original estimation metric.

`nolegend` suppresses the expression legend that appears before the fixed-effects estimation table when functions of parameters or named substitutable expressions are specified in the main equation or in the `define()` options.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

`nostderr` prevents *menl* from calculating standard errors for the estimated random-effects parameters, although standard errors are still provided for the fixed-effects parameters. Specifying this option will speed up computation times.

display_options: `noci`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

EM options

These options control the EM iterations that occur before estimation switches to the Lindstrom–Bates method. EM is used to obtain starting values.

`emiterate(#)` specifies the number of EM iterations to perform. The default is `emiterate(25)`.

`emtolerance(#)` specifies the convergence tolerance for the EM algorithm. The default is `emtolerance(1e-10)`. EM iterations will be halted once the log (restricted) likelihood changes by a relative amount less than `#`. At that point, optimization switches to the Lindstrom–Bates method.

`emlog` specifies that the EM iteration log be shown. The EM iteration log is not displayed by default.

Maximization

menlmaxopts: `iterate(#)`, `tolerance(#)`, `ltolerance(#)`, `pnlsopts()`, `lmeopts()`, `[no]log`. The convergence is declared when either `tolerance()` or `ltolerance()` is satisfied; see [Stopping rules](#) for details.

`iterate(#)` specifies the maximum number of iterations for the Lindstrom–Bates alternating algorithm. One iteration of the alternating algorithm involves $\#_{\text{pnls}}$ penalized least-squares (PNLS) iterations as specified in `pnlsopts()`’s `iterate()` suboption and $\#_{\text{lme}}$ linear mixed-effects (LME) iterations as specified in `lmeopts()`’s `iterate()` suboption. The default is the current value of `set maxiter`, which is `iterate(16000)` by default.

`tolerance(#)` specifies the tolerance for the parameter vector. When the relative change in the parameter vector from one alternating algorithm iteration to the next is less than or equal to `tolerance()`, the parameter convergence is satisfied. The default is `tolerance(1e-6)`.

`ltolerance(#)` specifies the tolerance for the linearization log likelihood in the Lindstrom–Bates alternating algorithm. The linearization log likelihood is the log likelihood from the LME optimization step in the last iteration. When the relative change in the linearization log likelihood from one alternating algorithm iteration to the next is less than or equal to `ltolerance()`, the log-likelihood convergence is satisfied. The default is `ltolerance(1e-7)`.

`pnlsopts(pnlsopts)` controls the PNLS optimization step of the Lindstrom–Bates alternating algorithm. *pnlsopts* include any of the following: `iterate(#)`, `ltolerance(#)`, `tolerance(#)`, `nrtolerance(#)`, and *maximize_options*. The convergence of this step within each alternating iteration is declared when `nrtolerance()` and one of `tolerance()` or `ltolerance()` are satisfied.

`iterate(#)` specifies the maximum number of iterations for the PNLS optimization step within the alternating algorithm. The default is `iterate(5)`.

`ltolerance(#)` specifies the tolerance for the objective function in the PNLS optimization step. When the relative change in the objective function from one PNLS iteration to the next is less than or equal to `ltolerance()`, the objective-function convergence is satisfied. The default is `ltolerance(1e-7)`.

`tolerance(#)` specifies the tolerance for the vector of fixed-effects parameters. When the relative change in the coefficient vector from one PNLS iteration to the next is less than or equal to `tolerance()`, the parameter convergence criterion is satisfied. The default is `tolerance(1e-6)`.

`nrtolerance(#)` specifies the tolerance for the scaled gradient in the PNLS optimization step. Convergence is declared when $g(-H^{-1})g'$ is less than `nrtolerance(#)`, where g is the gradient row vector and H is the approximated Hessian matrix from the current iteration. The default is `nrtolerance(1e-5)`.

maximize_options are `[no]log`, `trace`, `showtolerance`, `nonnrtolerance`; see [\[R\] maximize](#).

`lmeopts(lmeopts)` controls the LME optimization step of the Lindstrom–Bates alternating algorithm. *lmeopts* include any of the following: `iterate(#)`, `ltolerance(#)`, `tolerance(#)`, `nrtolerance(#)`, and *maximize_options*. The convergence of this step within each alternating iteration is declared when `nrtolerance()` and one of `tolerance()` or `ltolerance()` are satisfied.

`iterate(#)` specifies the maximum number of iterations for the LME optimization step within the alternating algorithm. The default is `iterate(5)`.

`ltolerance(#)` specifies the tolerance for the log likelihood in the LME optimization step. When the relative change in the log likelihood from one LME iteration to the next is less than or equal to `ltolerance()`, the log-likelihood convergence is satisfied. The default is `ltolerance(1e-7)`.

`tolerance(#)` specifies the tolerance for the random-effects and within-group error covariance parameters. When the relative change in the vector of parameters from one LME iteration to the next is less than or equal to `tolerance()`, the convergence criterion for covariance parameters is satisfied. The default is `tolerance(1e-6)`.

`nrtolerance(#)` specifies the tolerance for the scaled gradient in the LME optimization step. Convergence is declared when $g(-H^{-1})g'$ is less than `nrtolerance(#)`, where g is the gradient row vector and H is the approximated Hessian matrix from the current iteration. The default is `nrtolerance(1e-5)`.

maximize_options are `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `nonnrtolerance`; see [R] [maximize](#).

`[no]log`; see [R] [maximize](#).

The following option is available with `men1` but is not shown in the dialog box:

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Remarks are presented under the following headings:

[Introduction](#)

[Random-effects substitutable expressions](#)

[Substitutable expressions](#)

[Linear combinations](#)

[Linear forms versus linear combinations](#)

[Random effects](#)

[Multilevel specifications](#)

[Summary](#)

[Specifying initial values](#)

[Two-level models](#)

[Testing variance components](#)

[Random-effects covariance structures](#)

[Heteroskedastic within-group errors](#)

[Restricted maximum likelihood](#)

[Three-level models](#)

[Obtaining initial values](#)

[Linearization approach to finding initial values](#)

[Graphical approach to finding initial values](#)

[Smart regressions approach to finding initial values](#)

[Examples of specifying initial values](#)

Introduction

Nonlinear mixed-effects (NLME) models are models containing both fixed effects and random effects where some of, or all, the fixed and random effects enter the model nonlinearly. They can be viewed as a generalization of linear mixed-effects (LME) models (see [ME] [mixed](#)), in which the conditional mean of the outcome given the random effects is a nonlinear function of the coefficients and random effects. Alternatively, they can be considered as an extension of nonlinear regression models for independent data (see [R] [nl](#)), in which coefficients may incorporate random effects, allowing them to vary across different levels of hierarchy and thus inducing correlation within observations at the same level.

Why use NLME models? Can't we use higher-order polynomial LME models or generalized linear mixed-effects (GLME) models instead?

In principle, any smooth nonlinear function can be approximated by a higher-order polynomial. One may argue that we can use an LME (see [ME] [mixed](#)) polynomial model and increase the order of the polynomial until we get an accurate approximation of the desired nonlinear model. There are three problems with this approach. First, parameters in NLME models often have natural physical interpretations such as half-life and limiting growth. This is not the case in LME polynomial models. For example, what is the physical interpretation of the coefficient of `time`⁴? Second, NLME models typically use fewer parameters than the corresponding LME polynomial model, which provides a more parsimonious summarization of the data. Third, NLME models usually provide better predictions outside the range of the observed data than predictions based on LME higher-order polynomial models.

GLME models (see [ME] [meglm](#)) are also nonlinear, but in the restricted sense that the conditional mean response given random effects is a nonlinear function of the linear predictor that contains both fixed and random effects, and only indirectly nonlinear in fixed and random effects themselves. That is, the nonlinear function must be an invertible function of the linear predictor. However, many estimation methods for GLME and NLME models are similar because random effects enter both models nonlinearly.

Population pharmacokinetics, bioassays, and studies of biological and agricultural growth processes are just a few areas that use NLME models to analyze multilevel data such as longitudinal or repeated-measures data. Comprehensive treatments of both methodology and history of NLME models may be found in [Davidian and Giltinan \(1995\)](#), [Vonesh and Chinchilli \(1997\)](#), [Demidenko \(2013\)](#), and [Pinheiro and Bates \(2000\)](#). [Davidian and Giltinan \(2003\)](#) provide an excellent summary.

Consider a sample of M subjects from a population of interest, where n_j measurements, $y_{1j}, \dots, y_{n_j j}$, are observed on subject j at times $t_{1j}, \dots, t_{n_j j}$. By “subject”, we mean any distinct experimental unit, individual, panel, or cluster with two or more correlated observations. The basic nonlinear two-level model can be written as follows (in our terminology, a one-level NLME is just a nonlinear regression model for independent data),

$$y_{ij} = \mu(\mathbf{x}'_{ij}, \boldsymbol{\beta}, \mathbf{u}_j) + \epsilon_{ij} \quad i = 1, \dots, n_j; \quad j = 1, \dots, M \quad (1)$$

where $\mu(\cdot)$ is a real-valued function that depends on a $p \times 1$ vector of fixed effects $\boldsymbol{\beta}$, a $q \times 1$ vector of random effects $\mathbf{u}_j$, which are distributed as multivariate normal with mean $\mathbf{0}$ and variance–covariance matrix $\boldsymbol{\Sigma}$, and a covariate vector $\mathbf{x}_{ij}$ that contains both within-subject covariates $\mathbf{x}_{ij}^w$ and between-subject covariates $\mathbf{x}_{ij}^b$. The $n_j \times 1$ vector of errors $\boldsymbol{\epsilon}_j = (\epsilon_{1j}, \dots, \epsilon_{n_j j})'$ is assumed to be multivariate normal with mean $\mathbf{0}$ and variance–covariance matrix $\sigma^2 \boldsymbol{\Lambda}_j$, where depending on $\boldsymbol{\Lambda}_j$, σ^2 is either a within-group error variance σ_e^2 or a squared scale parameter σ^2 .

Parameters of NLME models often have scientifically meaningful interpretations, and research questions are formed based on them. To allow parameters to reflect phenomena of interest, (1) can be equivalently formulated as a two-stage hierarchical model as follows:

$$\begin{aligned} \text{Stage 1: Individual-level model } y_{ij} &= m(\mathbf{x}_{ij}^w, \phi_j) + \epsilon_{ij} & i = 1, \dots, n_j \\ \text{Stage 2: Group-level model } \phi_j &= \mathbf{d}(\mathbf{x}_j^b, \beta, \mathbf{u}_j) & j = 1, \dots, M \end{aligned} \quad (2)$$

In stage 1, we model the response by using a function $m(\cdot)$, which describes within-subject behavior. This function depends on subject-specific parameters ϕ_j 's, which have a natural physical interpretation, and a vector of within-subject covariates $\mathbf{x}_{ij}^w$. In stage 2, we use a known vector-valued function $\mathbf{d}(\cdot)$ to model between-subject behavior, that is, to model ϕ_j 's and to explain how they vary across subjects. The $\mathbf{d}(\cdot)$ function incorporates random effects and, optionally, a vector of between-subject covariates $\mathbf{x}_j^b$. The general idea is to specify a common functional form for each subject in stage 1 and then allow some parameters to vary randomly across subjects in stage 2.

To further illustrate (1) and (2), we consider the soybean plants data (Davidian and Giltinan 1995), in which we model the average leaf weight per soybean plant, y_{ij} , in plot j at t_{ij} days after planting. Let's first use (1):

$$\begin{aligned} y_{ij} &= \mu(\mathbf{x}_{ij}', \beta, \mathbf{u}_j) + \epsilon_{ij} \\ &= \frac{\beta_1 + u_{1j}}{1 + \exp[-\{t_{ij} - (\beta_2 + u_{2j})\} / (\beta_3 + u_{3j})]} + \epsilon_{ij} \end{aligned}$$

Here $\beta = (\beta_1, \beta_2, \beta_3)'$, $\mathbf{u}_j = (u_{1j}, u_{2j}, u_{3j})'$, and $\mathbf{x}_{ij}$ is simply t_{ij} .

Equivalently, we can use (2) to define our model,

$$\begin{aligned} \text{Stage 1: } y_{ij} &= m(\mathbf{x}_{ij}^w, \phi_j) + \epsilon_{ij} \\ &= \frac{\phi_{1j}}{1 + \exp\{-(t_{ij} - \phi_{2j}) / \phi_{3j}\}} + \epsilon_{ij} \\ \text{Stage 2: } \phi_{1j} &= \beta_1 + u_{1j} \\ \phi_{2j} &= \beta_2 + u_{2j} \\ \phi_{3j} &= \beta_3 + u_{3j} \end{aligned}$$

where $\mathbf{x}_{ij}^w = t_{ij}$, $\phi_j = (\phi_{1j}, \phi_{2j}, \phi_{3j})' = \mathbf{d}(\mathbf{x}_j^b, \beta, \mathbf{u}_j) = \beta + \mathbf{u}_j$. A key advantage of (2) is the interpretability. ϕ_j 's are parameters that characterize features of the trajectory. For example, ϕ_{1j} can be interpreted as the asymptotic average leaf weight per soybean plant in plot j when $t_{ij} \rightarrow \infty$ and ϕ_{2j} as the time at which half of ϕ_{1j} is reached; that is, if we set $t_{ij} = \phi_{2j}$, then $E(y_{ij}) = \phi_{1j}/2$. **menl** provides both representations.

The random effects $\mathbf{u}_j$ are not directly estimated (although they may be predicted) but instead are characterized by the elements of Σ , known as variance components, which are estimated together with the parameters of the within-group error variance–covariance matrix $\sigma^2 \mathbf{\Lambda}_j$. Correlation among repeated measures is induced either indirectly through the subject-specific random effects $\mathbf{u}_j$ or directly through specification of the within-subject covariance matrix $\sigma^2 \mathbf{\Lambda}_j$. Several covariance structures are available for Σ , similar to those allowed in **mixed**. In contrast to **mixed**, **menl** provides more flexible modeling of the within-subject variance and correlation structures.

menl uses the following decomposition of the Λ_j matrix,

$$\Lambda_j = \mathbf{S}_j \mathbf{C}_j \mathbf{S}_j \quad (3)$$

where $\mathbf{S}_j$ is diagonal with positive elements such that $\text{Var}(\epsilon_{ij}) = \sigma^2[\mathbf{S}_j]_{ii}^2$ and $\mathbf{C}_j$ is a correlation matrix such that $\text{corr}(\epsilon_{ij}, \epsilon_{kj}) = [\mathbf{C}_j]_{ik}$; $[A]_{ij}$ denotes the ij th element of matrix A . Decomposition (3) of Λ_j allows us to separately model the variance structure (heteroskedasticity) and the correlation structure by using disjoint sets of parameters for $\mathbf{C}_j$ and $\mathbf{S}_j$. This is different from how `mixed` handles within-subject correlation, where heteroskedasticity and correlation are determined by the type of the chosen residual covariance structure. For convenience, `menl` accommodates the behavior of the `mixed` command for specifying residual covariance structures via the `rescovariance()` option. The more flexible modeling of the residual structures according to (3) is available via the `resvariance()` and `rescorrelation()` options.

For LME models, because the random effects $\mathbf{u}_j$'s are unobserved, inference about β and the covariance parameters are based on the marginal likelihood obtained after integrating out the random effects. Unlike LME models, no closed-form solution is available because the random effects enter the model nonlinearly, making the integration analytically intractable in all but the simplest situations. There are two principal methods proposed in the literature for fitting NLME models. One is to use an adaptive Gauss–Hermite (AGH) quadrature to approximate the integral that appears in the expression of the marginal likelihood. The other one is to use the linearization method of [Lindstrom and Bates \(1990\)](#), also known as a conditional first-order linearization method, which is based on a first-order Taylor-series approximation of the mean function and essentially linearizes the mean function with respect to fixed and random effects. With the AGH method, the level of accuracy increases as the number of quadrature points increases but at the expense of increasing computational burden. The linearization method is computationally efficient because it avoids the intractable integration, but the approximation cannot be made arbitrarily accurate. Despite its potential limiting accuracy, the linearization method has proven the most popular in practice ([Fitzmaurice et al. 2009](#), sec. 5.4.8). The linearization method of [Lindstrom and Bates \(1990\)](#), with extensions from [Pinheiro and Bates \(1995\)](#), is the method of estimation in `menl`.

The linearization method uses a first-order Taylor-series expansion of the specified nonlinear mean function to approximate it with a linear function of fixed and random effects. Thus an NLME model is approximated by an LME model, in which the fixed-effects and random-effects design matrices involve derivatives of the nonlinear mean function with respect to fixed effects (coefficients) and random effects, respectively. As such, inference after the linearization method uses the computational machinery of the LME models. For example, estimates of random effects are computed as best linear unbiased predictors (BLUPs) of random effects from the approximating LME model. The accuracy of the inferential results will depend on the accuracy of the linearization method in approximating the original NLME model. In general, asymptotic inference for the NLME models based on the linearization method is only “approximately asymptotic”, making it less accurate than the corresponding asymptotic inference for true LME models. In practice, however, the linearization method was found to perform well in many situations (for example, [Pinheiro and Bates \[1995\]](#); [Wolfinger and Lin \[1997\]](#); [Plan et al. \[2012\]](#); and [Harring and Liu \[2016\]](#)).

Both ML and REML estimation are supported by `menl`. The ML estimates are based on the usual application of likelihood theory, given the distributional assumptions of the model. In small samples, ML estimation generally leads to small-sample bias in the estimated variance components. The REML method ([Thompson 1962](#)) reduces this bias by forming a set of linear contrasts of the response that do not depend on the fixed effects β but instead depend only on the variance components to be estimated. The likelihood is then formed based on the distribution of the linear contrasts, and standard ML methods are applied.

The next section describes how to specify nonlinear expressions containing random effects in `menl`.

Random-effects substitutable expressions

You define the nonlinear model to be fit by **menl** by using a random-effects substitutable expression, a substitutable expression that contains random effects. For example, `exp({b}+{U[id]})`, `{b1}/({b2}+{b3}*x+{U[id]})`, and `({b1}+{U1[id]})/(1+{b2}*x+{c.x#U2[id]})` are a few examples of such expressions. We describe them in more detail below.

Substitutable expressions

Let's first consider substitutable expressions without random effects. Substitutable expressions are just like any other mathematical expressions involving scalars and variables, such as those you would use with Stata's **generate** command, except that the parameters to be estimated are bound in braces. See [\[U\] 13.2 Operators](#) and [\[U\] 13.3 Functions](#) for more information on expressions.

Although **menl** requires that you include random effects in your model, for teaching purposes, we will start with simpler substitutable expressions that do not contain random effects. Suppose that we wish to fit the model

$$y_{ij} = \alpha \left(1 - e^{-(\beta_0 + \beta_1 x_{1ij} + \beta_2 x_{2ij})} \right) + \epsilon_{ij}$$

where α , β_0 , β_1 , and β_2 are the parameters to be estimated and ϵ_{ij} is an error term. We could simply type

```
. menl y = {a}*(1 - exp(-({b0}+{b1}*x1+{b2}*x2)))
```

Because **a**, **b0**, **b1**, and **b2** are enclosed in braces, **menl** knows that they are parameters in the model.

You can group several parameters together by assigning a group name (or equation name) to them. Parameters with the same group name, **lc** in the example below, will be grouped together in the output table:

```
. menl y = {a}*(1 - exp(-({lc:b0}+{lc:b1}*x1+{lc:b2}*x2)))
```

That is, parameters **b0**, **b1**, and **b2** will appear together in the output table in the equation labeled **lc**. Parameters without equation names will appear at the bottom of the output table.

Sometimes, it may be convenient to define subexpressions within the main expression. This can be done inside the expression itself or by using the **define()** option. For example,

```
. menl y = {a}*(1 - exp(-{xb:})), define(xb: {lc:b0}+{lc:b1}*x1+{lc:b2}*x2)
```

defines the linear predictor of the exponent in the **define()** option with label **xb** and then refers to it inside the exponent as **{xb:}**. You can define as many subexpressions as you like by using the **define()** option repeatedly. Defining subexpressions is also useful for later predictions; see, for instance, [example 13](#).

The above is equivalent to

```
. menl y = {a}*(1 - exp(-{xb: {lc:b0}+{lc:b1}*x1+{lc:b2}*x2}))
```

Parameters **{a}**, **{lc:b0}**, **{lc:b1}**, and **{lc:b2}** are what we call “free parameters”, meaning that they are not defined by a [linear form](#), which we describe in the next section. Free parameters are displayed with a forward slash in front of their names or their group names.

The general syntax for a free parameter is

```
{[ eqname: ] name}
```

Linear combinations

Nonlinear functions will often contain linear combinations of variables. Recall our nonlinear function from [Substitutable expressions](#):

$$y_{ij} = \alpha \left(1 - e^{-(\beta_0 + \beta_1 x_{1ij} + \beta_2 x_{2ij})} \right) + \epsilon_{ij}$$

Instead of explicitly specifying the linear combination that appears in the exponent, as we did in the previous section, we can use `menl`'s shorthand notation

```
. menl y = {a}*(1 - exp(-({lc: x1 x2})))
```

By specifying `{lc:x1 x2}`, you are telling `menl` that you are declaring a linear combination named `lc` that is a function of two variables, `x1` and `x2`. `menl` will create three parameters, named `{lc:_cons}`, `{lc:x1}`, and `{lc:x2}`.

Although both specifications produce the same results, the shorthand specification is more convenient.

The general syntax for defining a linear combination is

```
{eqname: varspec [ , xb noconstant ]}
```

where `varspec` includes a list of variables ([varlist](#)), a list of [random-effects terms](#), or both.

The `xb` option is used to distinguish between the linear combination that contains one variable and a free parameter that has the same name as the variable and the same group name as the linear combination. For example, `{lc: x1, xb}` is equivalent to `{lc:_cons} + {lc:x1}*x1`, whereas `{lc:x1}` refers to either a free parameter `x1` with a group name `lc` or the coefficient of the `x1` variable, if `{lc:}` has been previously defined in the expression as a linear combination that involves variable `x1`; see examples below. Thus the `xb` option indicates that the specification is a linear combination rather than a single parameter to be estimated.

When you define a linear combination, a constant term is included by default (a mathematician would argue that “affine combination” is the correct terminology!). The `noconstant` option suppresses the constant.

Having defined a linear combination such as `{lc:x1 x2}`, you can refer to its individual coefficients by using `{lc:x1}` and `{lc:x2}` or, more generally, `{eqname:varname}`. For example, suppose that we want to fit the following model, where the coefficient of `x1`, β_1 , appears in two places in the expression:

$$y_{ij} = \frac{1}{(1 + \beta_1 x_{1ij} + \beta_2 x_{2ij} + \beta_3 x_{3ij})} \exp \{ -(\alpha_0 + \alpha_1 z_{ij}) / (1 + \beta_1 x_{4ij}) \} + \epsilon_{ij}$$

We use `{lc1: x1 x2 x3, noconstant}` to specify the first linear combination, which appears in the denominator outside the exponentiated expression, and then use `{lc1:x1}` to refer to β_1 in the denominator inside the exponentiated expression. We also use the `xb` option, when we specify the second linear combination that contains only one covariate `z`. Below is the full specification:

```
. menl y = 1/({lc1: x1 x2 x3, noconstant})*exp(-{lc2: z, xb}/({lc1:x1}*x4))
```

You may also refer to a “subset” of a previously defined linear combination. For example, let’s modify our previous expression by substituting $\beta_1 x_{4ij}$ in the denominator in the exponent with the subset $\beta_1 x_{1ij} + \beta_3 x_{3ij}$ of the first linear combination:

$$y_{ij} = \frac{1}{(1 + \beta_1 x_{1ij} + \beta_2 x_{2ij} + \beta_3 x_{3ij})} \exp \{ -(\alpha_0 + \alpha_1 z_{ij}) / (1 + \beta_1 x_{1ij} + \beta_3 x_{3ij}) \} + \epsilon_{ij}$$

The coefficients for variables `x1` and `x3` are the same in the denominators inside and outside the exponent. We fit this model by typing

```
. menl y = 1/(1+{lc1: x1 x2 x3, nocons})*      ///
      exp(-{lc2: z, xb}/(1+{lc1: x1 x3, nocons}))
```

We used the same equation name, `lc1`, to constrain the coefficients to be the same between the two linear-combination specifications. If we used a different equation name, say, `lc3`, in the last linear combination, we would have specified $\beta_4 x_{1ij} + \beta_5 x_{3ij}$ instead of $\beta_1 x_{1ij} + \beta_3 x_{3ij}$ and estimated two extra parameters, β_4 named `{lc3:x1}` and β_5 named `{lc3:x3}`.

To refer to the entire linear combination that was already defined, you can simply refer to its name. For example, if both denominators included the same linear combination, $\beta_1 x_{1ij} + \beta_2 x_{2ij} + \beta_3 x_{3ij}$, the corresponding `menl` specification would be

```
. menl y = 1/(1+{lc1: x1 x2 x3, nocons})*exp(-{lc2: z, xb}/(1+{lc1:}))
```

Just like subexpressions, linear combinations can be defined in the `define()` option. For example, the above is equivalent to

```
. menl y = 1/(1+{lc1:})*exp(-{lc2:}/(1+{lc1:})), define(lc1: x1 x2 x3, nocons) ///
      define(lc2: z, xb)
```

Linear forms versus linear combinations

As we mentioned in [Linear combinations](#), the linear-combination specification is syntactically convenient. It can also be more computationally efficient when a linear combination is a linear form.

A linear combination is what we call a linear form as long as you do not refer to its coefficients or any subset of the linear combination anywhere in the expression. Linear forms are beneficial for some nonlinear commands such as `nl` because they make derivative computation faster and more accurate. Although `menl` does not fully utilize the linear-form specification in its computations, it is still important to learn to distinguish between linear forms and linear combinations.

For example, in [Linear combinations](#), the first linear combination `{lc:}`, the linear combination `{lc2:}`, and the linear combination `{lc1:}` in the last example are all linear forms. The linear combination `{lc1:}` in the examples where we referred to `{lc1:x1}` and `{lc1:x1 x3}` is not a linear form.

In contrast to free parameters, parameters of a linear form are displayed without forward slashes in the output. Rather, they are displayed as parameters within an equation whose name is the linear combination name. Parameters of linear combinations that are not linear forms are considered free parameters.

Random effects

So far, we have restricted our discussion to substitutable expressions that do not contain random effects. Examples of random effects specified within the `menl` syntax are `{U1[id]}`, `{U2[id1>id2]}`, `{c.x1#U3[id]}`, and `{2.f1#U4[id]}`. These represent a random intercept at the `id` level, a random intercept at the `id2`-within-`id1` level, a random slope for the continuous variable `x1`, and a random slope associated with the second level of the factor variable `f1`, respectively.

The general syntax for specifying random effects, *respec*, is provided below.

<i>respec</i>	Description
<code>{rename[<i>levelspec</i>]}</code>	Random intercepts <i>rename</i> at hierarchy <i>levelspec</i>
<code>{c.varname#rename[<i>levelspec</i>]}</code>	Random coefficients <i>rename</i> for continuous variable <i>varname</i>
<code>{#.fvvarname#rename[<i>levelspec</i>]}</code>	Random coefficients <i>rename</i> for the #th level of factor variable <i>fvvarname</i>

rename is a random-effects name. It is a Stata name that starts with a capital letter. *levelspec* defines the level of hierarchy and is described below.

<i>levelspec</i>	Description
<i>levelvar</i>	variable identifying the group structure for the random effect at that level
<i>lv2 > lv1</i>	two-level nesting: levels of variable <i>lv1</i> are nested within <i>lv2</i>
<i>lv3 > lv2 > lv1</i>	three-level nesting: levels of variable <i>lv1</i> are nested within <i>lv2</i> , which is nested within <i>lv3</i>
<i>... > lv3 > lv2 > lv1</i>	higher-level nesting

You can equivalently specify levels in the opposite order, from the lowest level to the highest; for example, *lv1 < lv2 < lv3*, but they will be displayed in the canonical order, from the highest level to the lowest.

Random effects can be specified within a linear-combination specification such as `{lc_u: x1 x2 U1[id1] U2[id2>id1]}`. In this case, the curly braces around each random effect are not needed.

Let us illustrate several random-effects specifications with `menl`. In this section, we concentrate on two-level nonlinear models; see [Multilevel specifications](#) for higher-level models.

Suppose that we want to fit the following model:

$$y_{ij} = \frac{\alpha z_{ij} + u_{0j}}{1 + \exp\{-(\beta_0 + \beta_1 x_{1ij})\}} + \epsilon_{ij}$$

Compared with models we considered in previous sections, this model includes random effects or, specifically, random intercepts u_{0j} . Suppose that these random intercepts correspond to the levels of the *id* variable. Then, we can include them in our model by using `{U0[id]}`, where *U0* will be their name.

```
. menl y = ({a}*z+{U0[id]} )/(1+exp(-({b0}+{b1}*x1)))
```

A more efficient specification is to use the linear-combination notation:

```
. menl y = {lc1: z U0[id], nocons}/(1+exp(-{lc2: x1, xb}))
```

The curly braces around `U0[id]` are removed when it is specified within a linear-combination specification.

If you need to refer to the random-effects term again in the expression, you can simply use its name. For example, suppose that our model includes the same random intercepts in both the numerator and the denominator.

$$y_{ij} = \frac{\alpha z_{ij} + u_{0j}}{1 + \exp\{-(\beta_0 + \beta_1 x_{1ij} + u_{0j})\}} + \epsilon_{ij}$$

We include random intercepts u_{0j} 's in the second linear combination by simply referring to their name, *U0*:

```
. menl y = {lc1: z U0[id], nocons}/(1+exp(-{lc2: x1 U0}))
```

If instead of u_{0j} 's, we had a different set of random intercepts, v_{0j} 's, in the denominator, we would need to specify a new set of random intercepts, say, `V0[id]`, with `men1`:

```
. men1 y = {lc1: z U0[id], nocons}/(1+exp(-{lc2: x1 V0[id]}))
```

The shorthand notation for referring to random effects only by name, that is, without the brackets and the *levels*spec, is also useful when specifying the `covariance()` option, especially for multilevel random effects with long-level specifications; see *Multilevel specifications*.

Let's now see how to include random slopes. Consider the following extension of the *first*, simpler model in this subsection:

$$y_{ij} = \frac{\alpha z_{ij} + u_{0j} + u_{1j} z_{ij}}{1 + \exp\{-(\beta_0 + \beta_1 x_{1ij})\}} + \epsilon_{ij}$$

Here u_{1j} is a random slope for a continuous variable z and is specified as `{c.z#U1[id]}` directly or without curly braces within a linear-combination specification.

```
. men1 y = {lc1: z U0[id] c.z#U1[id], nocons}/(1+exp(-{lc2: x1, xb}))
```

We can also include random slopes for factor variables. To demonstrate this, let's consider a different nonlinear model for variety. Consider the model below, where binary variables x_{1ij} and x_{2ij} correspond to the factor levels 1 and 2 of a factor variable x that takes on values 0, 1, and 2, with 0 being the base level.

$$y_{ij} = \alpha_0 + \alpha_1 z_{1ij} - \sqrt{w_{ij}^2 + \exp(\beta_0 + \beta_1 x_{1ij} + \beta_2 x_{2ij} + u_{0j} + u_{1j} x_{1ij} + u_{2j} x_{2ij})} + \epsilon_{ij}$$

There are three random-effects terms in this model: random intercepts u_{0j} , random slopes u_{1j} for x_{1ij} (level 1 of x), and random slopes u_{2j} for x_{2ij} (level 2 of x). In Stata, for a factor variable x , we can use the factor-variable notation ([\[U\] 11.4.3 Factor variables](#)) to refer to its levels, `1.x` for level 1 and `2.x` for level 2. So, to include the three random-effects terms in `men1`, we will use `U0[id]`, `1.x#U1[id]`, and `2.x#U2[id]`, respectively.

```
. men1 y = {lc1: z1, xb} - sqrt(c.w#c.w + ///
    exp({lc2: 1.x U0[id] 1.x#U1[id] 2.x#U2[id]}))
```

In the above specification, we used the factor-variable notations `1.x` to include fixed effects for all levels of x , except the base level, and `c.w#c.w` to include a square of w ; see [\[U\] 11.4.3 Factor variables](#) for details. The factor-variable specification `1.` or any other specification that refers to multiple levels of a factor variable is not allowed when specifying random coefficients, because each level will typically require a different set of random effects. For example, if we had specified `1.x#U[id]` in the above example, we would have received an error.

Multilevel specifications

In *Random effects*, we focused on specifying substitutable expressions containing random effects for two-level nonlinear mixed-effects models. Here we will consider higher-level models.

Suppose that we want to fit the following three-level nonlinear mixed-effects model,

$$y_{ijk} = \beta_0 + u_{0k}^{(3)} + u_{0jk}^{(2)} + \cos\left\{\left(\beta_1 + u_{1k}^{(3)}\right) x_{1ijk}\right\} + \epsilon_{ijk}$$

where first-level observations, indexed by i , are nested within second-level groups, indexed by j , which are nested within third-level groups, indexed by k .

There are three random-effects terms in this model: random intercepts, $u_{0k}^{(3)}$, and random slopes for x_1 , $u_{1k}^{(3)}$, at the third level (idk) and random intercepts $u_{0jk}^{(2)}$ at the second level (idj-nested-within-idk). We specify random intercepts and random slopes for x_1 at the highest hierarchical level just like we did in [Random effects](#) for two-level models. Specifically, we can use `U0[idk]` and `c.x1#U1[idk]`, respectively. To specify random intercepts $u_{0jk}^{(2)}$ at the idj-nested-within-idk level, we need to use one of the [levspec](#) specifications for two nested levels. For example, we can use `UU0[idk>idj]`. Below is the full specification:

```
. menl y = {lc1: U0[idk] UU0[idk>idj]} + cos({lc2: x1 c.x1#U1[idk], noconstant})
```

We can also include a random slope of the x_1 variable at the idj-within-idk level in the cosine function by specifying `c.x1#UU1[idk>idj]` inside the `cos()` function.

```
. menl y = {lc1: U0[idk] UU0[idk>idj]} + ///
cos({lc2: x1 c.x1#U1[idk] c.x1#UU1[idk>idj], noconstant})
```

We can shorten the above specification by writing `c.x1#U1[idk] c.x1#UU1[idk>idj]` more compactly as `c.x1#(U1[idk] UU1[idk>idj])`,

```
. menl y = {lc1: U0[idk] UU0[idk>idj]} + ///
cos({lc2: x1 c.x1#(U1[idk] UU1[idk>idj]), noconstant})
```

Similarly, if we had a four-level model with, say, a random intercept at the idj-within-idk-within-idl level, we could specify it as `W[idl>idk>idj]`; see [levspec](#) for other specifications.

Summary

To summarize, here are a few rules to keep in mind when defining substitutable expressions.

1. Model parameters and random effects are bound in braces if specified directly in the expression: `{b0}`, `{U0[id]}`, etc.
2. Model parameters can be assigned group names: `{slopes:x1}`, `{slopes:x2}`, etc.
3. Random-effects names must start with a capital letter as in `{U0[id]}`, `{c.x#U1[id]}`, `{V0[id2>id1]}`, `{1.z#V1[id2>id1]}`, etc.
4. The factor-variable specification `i.`, as in `{i.z#V1[id2>id1]}`, or any other specification that refers to multiple levels of a factor variable, as in `{i(1/4).z#V1[id2>id1]}`, is not allowed when specifying random coefficients.
5. Linear combinations of variables can be included using the specification

```
{eqname:varlist[, xb noconstant]}
```

For example, `{price: mpg weight i.rep78}` and `{lc: x1 x2, noconstant}`.

6. Random effects can be specified within a linear combination, in which case they should be included without curly braces, for example, `{lc_u: x1 x2 U[id]}`.
7. To specify a linear combination that contains only one variable, use the `xb` option, for example, `{lc: x1, xb}`.
8. To refer to the previously defined linear combination again in the expression, simply use its name `{eqname:}`, for example, `{lc:}` and `{lc_u:}`.
9. You can refer to individual parameters of the linear combination by using `{eqname:_cons}` and `{eqname:varname}`, for example, `{price:_cons}` and `{price:weight}`.

10. You can refer to a “subset” of the previously defined linear combination by using `{eqname:subset}`, where *subset* is a subset of the variables from *varlist* used to define *eqname*, as in `{price: mpg weight}`. To refer to the subset containing only one variable, use the `xb` option, as in `{price: weight, xb}`. If a linear combination contains only one random-effects term, the `xb` option is implied.
11. To refer to the previously defined random effects again in the expression or in the `covariance()` option, simply use their names, such as `{U0}` and `{U1}`.
12. You can define subexpressions, including linear combinations, inside the main expression or in the `define()` option, which can be repeated. For example,


```
. menl y = {number:}/{denom:}, define(number: z U0[id]) ///
      define(denom:1+exp(-{lc: x1, xb}))
```
13. Specify linear forms whenever possible for faster and more accurate computation of derivatives; see [Linear forms versus linear combinations](#).
14. Model parameters that are not defined by linear forms are considered free parameters. They are included in the output with a forward slash in front of their names or group names and displayed after linear forms in the estimation table.

Specifying initial values

By default, **menl** uses the EM algorithm to obtain initial values, but you may often need to specify your own. You specify your own initial values in the `initial()` option. For example, specifying the `initial(a 1.1 b -2)` option with **menl** initializes parameter `{a}` to 1.1 and parameter `{b}` to -2.

When you specify your own initial values, they are used for initialization, and the EM algorithm is not performed. When you specify initial values for only a subset of model parameters, the remaining parameters are initialized with some predetermined values such as zeros for fixed-effects parameters and correlations and ones for variances. You can specify the `iterate(0)` option to see the initial values that will be used by **menl** in the optimization.

Often, you may have good initial values for fixed-effects parameters but not for random-effects parameters. In this case, you can specify `initial()`’s `fixed` suboption to supply your own fixed-effects parameters, but use the EM algorithm to obtain initial values for the random-effects parameters.

There are three ways in which you can use the `initial(initial_values)` option: you can specify a vector of values, a list of values, or values for individual parameters and groups of parameters.

Specifically, *initial_values* is one of the following:

```
vectorname [ , skip copy fixed ]
# [#] [...], copy
paramlist [=]# [paramlist [=]# [...]] [ , fixed ]
```

skip specifies that any parameters found in the specified initialization vector, *vectorname*, that are not also found in the model be ignored. The default action is to issue an error message.

copy specifies that the initial values be copied into the initialization vector without checking for valid column names. **copy** must be specified when initial values are supplied as a list of numbers.

fixed specifies that initial estimates are being supplied for the fixed effects only and that **menl** should still perform the EM algorithm to refine initial values for variance components. The specified initial values are used for fixed-effects parameters during the EM algorithm. If you omit **fixed**, **menl** presumes that you are specifying starting values for all parameters in `e(b)`, and the EM algorithm will not be performed.

Examples of *paramlist* are `param`, `{param}`, `{param1} {param2}`, `{param1 param2}`, `{grp:param1} {grp:param2}`, `{grp:param3}`, `{grp:param1 param2}`, and `{grp:}`.

Let's describe each specification in more detail. You can specify the name of a vector containing the initial values, say, `initial(b0)`. Vector `b0` should be properly labeled with labels found in column names of `e(b)`, unless you specify the `copy` option. A properly labeled vector can have fewer elements than `e(b)` or, if `skip` is specified, even more elements. A vector without labels must be of the same dimension as `e(b)`.

Alternatively, you can supply a list of numbers to `initial()`, in which case `copy` must be specified. The list of numbers should be of length equal to the dimension of `e(b)`. For example, if `e(b)` has four parameters and you type `initial(1.1 0 3 -2, copy)`, then the four coefficients in `e(b)` will be initialized to 1.1, 0, 3, and -2 , respectively. If instead you specify, for example, only three initial values in your list, an error will be issued.

Finally, you can initialize parameters by referring to their names. You can specify a parameter name, its initial value, another parameter name, its initial value, and so on, for example, `initial(a 1.1 b -2)`. You can also assign the same initial value to a group of parameters. For example, `initial({a b c} 1)` will initialize parameters `{a}`, `{b}`, and `{c}` to 1 and `initial({lc:x1 x2 _cons} 0)` will initialize `{lc:x1}`, `{lc:x2}`, and `{lc:_cons}` to 0. You can assign the same initial value to all parameters with the same group name. For example, we can shorten the previous specification to `initial({lc:} 0)`.

Depending on the situation, it may also be beneficial to specify initial values for the NLS algorithm used by `menl` to obtain starting values for the EM algorithm. These initial values can be specified in the parameter definition such as `{a=0.5}`, in which case the NLS algorithm used during the initialization will use 0.5 as the starting value for parameter `a` instead of the default 0. Such initialization is particularly useful for parameters used in the denominators for which zero values may lead to an undefined value of the mean function.

See [Examples of specifying initial values](#) and [Obtaining initial values](#) for examples.

Two-level models

The sole purpose of this section and its examples is to highlight the syntax of `menl` and make you familiar with how to specify substitutable expressions in `menl` and with its output. Also see an introductory example in [Nonlinear models](#) in [ME] `me`.

We will use the data from the Longitudinal Study of Unicorn Health in Zootopia, which contain the brain weight (`weight`) of 30 newborn male unicorns and 30 newborn female unicorns. Measurements were collected at 13 occasions every 2 months over the first 2 years after birth (`time`). Based on previous studies, a model for unicorn brain shrinkage is believed to be

$$\text{weight}_{ij} = \beta_1 + (\beta_2 - \beta_1) \exp(-\beta_3 \text{time}_{ij}) + \epsilon_{ij} \quad i = 1, 2, \dots, 13; \quad j = 1, 2, \dots, 60$$

Parameter β_1 represents the average brain weight of unicorns as `timeij` increases to infinity. Parameter β_2 is the average brain weight at birth (at `timeij = 0`), and β_3 is a scale parameter that determines the rate at which the average brain weight of unicorns approaches the asymptotic weight β_1 (decay rate). This model can be fit with the `nl` command; see [R] `nl`.

We will start with a simple two-level model in which we allow the asymptote parameter β_1 to vary between unicorns by replacing β_1 in the above equation with $\beta_1 + u_{0j}$,

$$\text{weight}_{ij} = \beta_1 + u_{0j} + (\beta_2 - \beta_1 - u_{0j}) \exp(-\beta_3 \text{time}_{ij}) + \epsilon_{ij} \quad (4)$$

where β_1 , β_2 , and β_3 are fixed-effects parameters to be estimated and u_{0j} is a random intercept at the unicorn, `id`, level that follows a normal distribution with mean 0 and variance σ_u^2 .

Equivalently, the model defined by (4) can be written as a two-stage model,

$$\text{weight}_{ij} = \phi_{1j} + (\phi_{2j} - \phi_{1j}) \exp(-\phi_{3j}\text{time}_{ij}) + \epsilon_{ij} \tag{5}$$

with the following stage 2 specification:

$$\begin{aligned} \phi_{1j} &= \beta_1 + u_{0j} \\ \phi_{2j} &= \beta_2 \\ \phi_{3j} &= \beta_3 \end{aligned} \tag{6}$$

Parameters ϕ_{1j} , ϕ_{2j} , and ϕ_{3j} now describe the behavior of the j th unicorn. For example, ϕ_{1j} represents the brain weight of the j th unicorn as time_{ij} increases to infinity.

► **Example 1: Simple two-level model**

Let’s use `menl` to first fit a single-equation model defined by (4), described above.

```
. use http://www.stata-press.com/data/r15/unicorn
(Brain shrinkage of unicorns in the land of Zootopia)

. menl weight = {b1}+{U0[id]}+({b2}-{b1}-{U0[id]})*exp(-{b3}*time)

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:    linearization log likelihood = -56.9757597

Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      780
Group variable: id                        Number of groups   =      60
                                           Obs per group:
                                           min =      13
                                           avg  =     13.0
                                           max  =      13

Linearization log likelihood =  -56.97576
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b1	4.707954	.1414511	33.28	0.000	4.430715	4.985193
/b2	8.089432	.0260845	310.12	0.000	8.038307	8.140556
/b3	4.13201	.0697547	59.24	0.000	3.995293	4.268726

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(U0)	1.189809	.2180063	.8308322	1.703888
	var(Residual)	.0439199	.0023148	.0396095	.0486995

Notes:

1. The response variable `weight` is specified on the left-hand side of the equality sign, and parameters to be estimated are enclosed in curly braces `{b1}`, `{b2}`, and `{b3}` on the right-hand side.
2. By typing `{U0[id]}`, we specified a random intercept at the level identified by the group variable `id`, that is, the unicorn level (level two).

3. The estimation log consists of three parts:

- a. A set of EM iterations used to refine starting values. By default, the iterations themselves are not displayed, but you can display them by using the `emlog` option. NLME models may often have multiple solutions and converge to a local maximum. It is thus important to try different initial values to investigate the existence of multiple solutions and the convergence to a global maximum; see [Obtaining initial values](#).
 - b. A set of iterations displaying the value of the [linearization log likelihood](#) from the Lindstrom–Bates algorithm or alternating algorithm. The term “linearization” reflects the fact that the reported log likelihood corresponds to the linear mixed-effects model obtained after [linearization](#) of the specified nonlinear mean function with respect to fixed and random effects. See [Inference based on linearization](#) and [Stopping rules](#) for details about the algorithm.
 - c. The message “Computing standard errors”. This is just to inform you that `menl` has finished its iterative maximization and is now reparameterizing the variance components (see [Methods and formulas](#)) to their natural metric and computing their standard errors. If you are interested only in the fixed-effects estimates, you can use the `nstderr` option to bypass this step.
4. The output title, “Mixed-effects ML nonlinear regression”, informs us that our model was fit using ML, the default. For REML estimates, use the `reml` option.
5. The header information is similar to that of the `mixed` command, but unlike `mixed`, `menl` in general does not report a model χ^2 statistic in the header because a test of the joint significance of all fixed-effects parameters (except the constant term) may not be relevant in a nonlinear model.
6. The first estimation table reports the fixed effects. We estimate $\beta_1 = 4.71$, $\beta_2 = 8.09$, and $\beta_3 = 4.13$. Although z tests against zeros are reported automatically for all fixed-effects parameters, as part of standard estimation output, they may not always be of interest or even appropriate for parameters of nonlinear models. You can use the `test` command ([\[R\] test](#)) to test hypotheses of interest or reparameterize your model so that the tests of parameters against zeros are meaningful.
7. The second estimation table shows the estimated variance components. The first section of the table is labeled `id: Identity`. In general, this means that our model includes random effects at the `id` (unicorn) level and that their variance–covariance matrix, Σ , is the identity matrix (all random effects have the same variance). In our example, because we have only one random effect, u_{0j} , the random-effect covariance structure is irrelevant, and the variance of the random intercept, σ_u^2 , labeled as `var(U0)` in the output, is estimated as 1.19 with standard error 0.22.
8. The row labeled `var(Residual)` displays the estimated overall error variance or variance of the error term; that is, $\widehat{\text{Var}}(\epsilon_{ij}) = \hat{\sigma}_e^2 = 0.044$.

◀

► Example 2: Two-level model as a two-stage model, using the `define()` option

The model from [example 1](#) can also be specified as a two-stage model, as defined by (5) and (6), by using the `define()` option. The `define()` option is particularly useful when you have a complicated nonlinear expression, and you would like to break it down into smaller pieces. Parameters of interest that are functions of other parameters can be defined using the `define()` option. This will make it easier to predict them for each subject after estimation; see [\[ME\] menl postestimation](#).

Below we specify the asymptote parameter, ϕ_{1j} , by using `define()`. The colon (`:`) in `{phi1:}` instructs `menl` that `phi1` will be specified as a substitutable expression within the `define()` option. Parameters `{phi2}` and `{phi3}` are simple free parameters and thus do not need to be specified in `define()`.

```
. menl weight = {phi1:}+({phi2}-{phi1:})*exp(-{phi3}*time),
> define(phi1: {b1}+{U0[id]})
Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:    linearization log likelihood = -56.9757597
Computing standard errors:
Mixed-effects ML nonlinear regression          Number of obs    =      780
Group variable: id                            Number of groups  =      60
                                                Obs per group:
                                                min =           13
                                                avg  =          13.0
                                                max  =           13

Linearization log likelihood =  -56.97576
      phi1:  {b1}+{U0[id]}
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b1	4.707954	.1414511	33.28	0.000	4.430715	4.985193
/phi2	8.089432	.0260845	310.12	0.000	8.038307	8.140556
/phi3	4.13201	.0697547	59.24	0.000	3.995293	4.268726

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(U0)	1.189809	.2180058	.8308328	1.703887
	var(Residual)	.0439199	.0023148	.0396095	.0486995

The results are identical to those obtained in [example 1](#), but the estimation table now has a legend that lists the expression `phi1` defined in the model. We can suppress this legend by specifying the `nolegend` option.

We could have defined `phi1` directly in the main expression instead of the `define()` option,

```
. menl weight = {phi1:{b1}+{U0[id]}}+({phi2}-{phi1:})*exp(-{phi3}*time)
(output omitted)
```

but by using the `define()` option, we simplified the main expression.

► **Example 3: Two-level model containing covariates**

Reducing brain weight loss has been an active research area in Zootopia for the past two decades, and scientists believe that consuming rainbow cupcakes right after birth may help slow down brain shrinkage. Recall that the scale parameter ϕ_{3j} determines the rate at which the brain weight of the j th unicorn decreases to its asymptotic value ϕ_{1j} . Hence, covariate `cupcake`, which represents the number of rainbow cupcakes consumed right after birth, is added to the equation of ϕ_{3j} . Also, we would like to investigate whether the asymptote, ϕ_{1j} , is gender specific, so we include the factor variable `female` in the equation for ϕ_{1j} . `femalej` is 1 if the j th unicorn is a female and 0 otherwise.

The stage 2 specification of the model defined by (5) becomes

$$\begin{aligned}\phi_{1j} &= \beta_{10} + \beta_{11}\text{female}_j + u_{0j} \\ \phi_{2j} &= \beta_2 \\ \phi_{3j} &= \beta_{30} + \beta_{31}\text{cupcake}_j\end{aligned}\tag{7}$$

The `define()` option can be repeated, so we specify a separate `define()` option for ϕ_{1j} , ϕ_{2j} , and ϕ_{3j} . We could have left ϕ_{2j} as a free parameter `{phi2}` in our specification, but we wanted to closely match the stage 2 specification (7).

```
. menl weight = {phi1:}+({phi2:}-{phi1:})*exp(-{phi3:}*time),
> define(phi1: {b10}+{b11}*1.female+{U0[id]})
> define(phi2: {b2})
> define(phi3: {b30}+{b31}*cupcake)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -29.0149875

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	780
Group variable: id	Number of groups	=	60
	Obs per group:		
	min	=	13
	avg	=	13.0
	max	=	13

Linearization log likelihood = -29.014988

```
phi1: {b10}+{b11}*1.female+{U0[id]}
phi2: {b2}
phi3: {b30}+{b31}*cupcake
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b10	4.072752	.1627414	25.03	0.000	3.753785	4.39172
/b11	1.264407	.2299723	5.50	0.000	.8136694	1.715144
/b2	8.088102	.0255465	316.60	0.000	8.038032	8.138172
/b30	4.706926	.1325714	35.50	0.000	4.44709	4.966761
/b31	-.2007309	.0356814	-5.63	0.000	-.2706651	-.1307966

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(U0)	.7840578	.1438924	.5471839	1.123473
	var(Residual)	.0420763	.0022176	.0379468	.0466551

In the table legend, `/b10` and `/b11` correspond, respectively, to the constant term and coefficient of `1.female` in the equation for ϕ_{1j} . `/b2` is ϕ_{2j} , and `/b30` and `/b31` correspond, respectively, to the constant term and coefficient for `cupcake` in the equation for ϕ_{3j} .

Based on our results, consuming rainbow cupcakes after birth indeed slows down brain shrinkage: `/b31` is roughly -0.2 with a 95% CI of $[-0.271, -0.131]$.

◀

▶ Example 4: Specifying linear combinations

A more convenient way to specify the model in [example 3](#) is to use linear-combination specifications; see [Random-effects substitutable expressions](#).

For example, `define(phi1: {b10}+{b11}*1.female+{U0[id]})` can be replaced with `define(phi1: i.female U0[id]). menl` knows that we are defining ϕ_{1j} as a linear combination of `i.female` and `U0[id]` and thus will create a constant term and a coefficient for each level of factor

variable `female` and will use a coefficient of 1 for any random effect. Because `female` has only two levels, `menl` will create two coefficients for `0b.female` and `1.female`, respectively, but will constrain the coefficient of the base level, level 0, to be 0.

We now fit our model by using linear-combination specifications within the `define()` options. We explain the use of the second and third `define()` specifications following estimation.

```
. menl weight = {phi1:}+({phi2:}-{phi1:})*exp(-{phi3:}*time),
> define(phi1: i.female U0[id])
> define(phi2: _cons, xb)
> define(phi3: cupcake, xb)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -29.0149875

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	780
Group variable: id	Number of groups	=	60
	Obs per group:		
	min	=	13
	avg	=	13.0
	max	=	13
	Wald chi2(2)	=	61.78
	Prob > chi2	=	0.0000
Linearization log likelihood = -29.014987			
phi1: i.female U0[id]			
phi3: cupcake, xb			

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
female						
female	1.264407	.2299723	5.50	0.000	.8136694	1.715144
_cons	4.072752	.1627414	25.03	0.000	3.753785	4.39172
phi2						
_cons	8.088102	.0255465	316.60	0.000	8.038032	8.138172
phi3						
cupcake	-.2007309	.0356814	-5.63	0.000	-.2706651	-.1307966
_cons	4.706926	.1325714	35.50	0.000	4.44709	4.966761

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(U0)	.7840578	.1438928	.5471833	1.123475
	var(Residual)	.0420763	.0022176	.0379468	.0466551

By using linear-combination specifications within the `define()` options, we improved the readability of the coefficient table. For example, it is now clear that `_cons` in the equation labeled `phi3` corresponds to the constant term in the equation for ϕ_{3j} . This term was labeled `/b30` previously.

Notes:

1. The `define()` option interprets its specification as a [random-effects substitutable expression](#), so you do not need to specify the curly braces (`{}`) around the specification.
2. All [rules](#) for random-effects substitutable expressions apply to the specifications within `define()`.

3. Following one of the rules for random-effects substitutable expressions, we used the `xb` option within `define()`s for `phi2` and `phi3`, because their linear combinations contained only one term: `_cons` for `phi2` and `cupcake` for `phi3`.
4. Specification `define(phi2: _cons, xb)` is the same as `define(phi2:, xb)`. In other words, `_cons` is implied with any linear combination, unless the `noconstant` option is specified. We chose to include `_cons` directly for clarity.
5. We could have used a free parameter `{phi2}` instead of the linear combination `{phi2: _cons, xb}`, but we wanted to preserve the order in which `phi1`, `phi2`, and `phi3` appear in the estimation table. See [example 5](#), where we specify ϕ_{2j} as a free parameter `{phi2}`.
6. In the presence of linear combinations, `menl` reports a joint test of significance of all coefficients (except the constant term) across all linear combinations.
7. Linear combinations containing only a constant such as `{phi2:}` are not listed in the table expression legend for brevity.

◀

► Example 5: Including random coefficients

In previous examples, we only illustrated how to specify random intercepts such as `{U0[id]}`, and it is bad karma to end a unicorn story without showing how to specify random coefficients or random slopes.

Continuing with our model as defined by (5) and (7), let's suppose that the equation for the brain-weight scale parameter, ϕ_{3j} , is as follows:

$$\phi_{3j} = \beta_{30} + (\beta_{31} + u_{1j})\text{cupcake}_j$$

We incorporated a unicorn-specific random slope for variable `cupcake`. The random slope, u_{1j} , for a continuous variable `cupcake` can be specified in `menl` as `c.cupcake#U1[id]`, and by default, `menl` assumes that it is independent of the random intercept, u_{0j} . (See [example 9](#) for specifying other random-effects covariance structures.)


```
. menl weight = {phi1:}+({phi2}-{phi1:})*exp(-{phi3:}*time),
> define(phi1: i.female U0[id])
> define(phi3: cupcake c.cupcake#U1[id])

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:   linearization log likelihood = 165.417513
Iteration 2:   linearization log likelihood = 165.420082
Iteration 3:   linearization log likelihood = 165.420112
Iteration 4:   linearization log likelihood = 165.420101

Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      780
Group variable: id                       Number of groups   =      60
                                           Obs per group:
                                           min =      13
                                           avg  =     13.0
                                           max  =      13
                                           Wald chi2(2)      =     46.70
                                           Prob > chi2       =     0.0000

Linearization log likelihood =    165.4201
      phi1:  i.female U0[id]
      phi3:  cupcake c.cupcake#U1[id]
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
female	1.320623	.2215707	5.96	0.000	.8863522	1.754893
female	4.006823	.1568268	25.55	0.000	3.699448	4.314198
_cons						
phi3						
cupcake	-.219661	.0659984	-3.33	0.001	-.3490155	-.0903066
_cons	4.771466	.1128421	42.28	0.000	4.5503	4.992633
/phi2	8.087655	.0179406	450.80	0.000	8.052492	8.122818

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Independent	var(U0)	.727464	.1337149	.5074024	1.042967
	var(U1)	.1258914	.0309569	.0777471	.2038487
	var(Residual)	.0208202	.0011403	.018701	.0231795

In addition to the overall error variance and the random-intercept variance, we now have a random-slope variance, which is labeled `var(U1)` in the output. In this example, we also specified parameter ϕ_{2j} as a free parameter `{phi2}` instead of a linear combination as in [example 4](#). As we mentioned in [Summary](#), free parameters are displayed after linear combinations, so `phi2` is listed last in the estimation table.

Previous studies of unicorns considered a model that also incorporated gender-specific variation between unicorns in asymptotic weight ϕ_{1j} ,

$$\phi_{1j} = \beta_{10} + u_{0j} + (\beta_{11} + u_{2j})\text{female}_j$$

but found no statistical evidence of such variation.

If we wanted to verify this for our data, we could have fit the following model:

```
. menl weight = {phi1:}+({phi2}-{phi1:})*exp(-{phi3:}*time), ///  
  define(phi1: i.female U0[id] 1.female#U2[id])           ///  
  define(phi3: cupcake c.cupcake#U1[id])
```

Compared with our previous specification, we included a new term in the `define()` option for `phi1`—a random slope for level 1 of the factor variable `female`, `1.female#U2[id]`. To include random slopes for a factor variable, we must specify random effects for each level, except the base level, of the factor variable. The specification `i.fvvarname` for referring to all levels of a factor variable is not allowed in the context of random effects, because a different set of random effects must be defined for each level. For example, if we specified `i.female#U2[id]` in our example, we would have received an error.

◀

To summarize:

1. Use `{name}` to define free parameters such as `{b1}`.
2. Use, for example, `{U0[id]}` to define random intercepts at the `id` level, `{c.varname#U1[id]}` to define random slopes for continuous variable `varname` at the `id` level, and `{#.fvvarname#U1[id]}` for each level `#`, except the base level, of variable `fvvarname` to include random slopes for factor variable `fvvarname`. The specification `{i.fvvarname#U1[id]}` is not allowed.
3. Use linear-combination specifications whenever possible. Do not use `{}` around random effects when they are specified within a linear combination.
4. Use multiple `define()` options to specify parameters of interest that are functions of other parameters, and use linear-combination specifications within `define()` whenever possible.
5. Use the `xb` option within a linear combination or within `define()` whenever you specify one variable such as `define(phi1: cupcake, xb)`, one random effect such as `{phi2: U0[id], xb}`, or a constant-only linear combination such as `{phi2: _cons, xb}` or `{phi2: , xb}`. When you specify the `xb` option, the above specifications are interpreted by `menl`, respectively, as a linear combination `{phi1:_cons}+{phi1:cupcake}*cupcake`, a linear combination `{phi:_cons}+{U0[id]}`, and a constant term `{phi2:_cons}`.
6. Unicorns do exist in our world, they are just gray, fat, and called rhinos.

Testing variance components

Consider data on the intensity of 23 large earthquakes in western North America between 1940 and 1980 analyzed originally by Joyner and Boore (1981) and then also by Davidian and Giltinan (1995, sec. 11.4). The objective of the study was to model the maximum horizontal acceleration (in g units), `accel`, taken at the i th measuring station for the j th earthquake, as a function of the magnitude of the quake on the Richter scale, `richter`, and the distance (in km) of the measuring station from the quake epicenter, `distance`. We are also interested in the possible effect of the soil type `soil`, soil versus rock, at a given measuring station on acceleration. The results of this study are useful to understand the nature of the damage caused by a particular earthquake and to determine the location for sensitive installations such as nuclear facilities.

Let's consider one of the models studied by Davidian and Giltinan (1995) for these data,

$$\log_{10}(\text{accel}_{ij}) = \phi_{1j} - \log_{10} \sqrt{\text{distance}_{ij}^2 + \exp(\phi_{2j})} - \phi_{3ij} \sqrt{\text{distance}_{ij}^2 + \exp(\phi_{2j})} + \epsilon_{ij} \quad (8)$$

where

$$\begin{aligned}\phi_{1j} &= \beta_0 + \beta_1 \text{richter}_j + u_{1j} \\ \phi_{2j} &= \beta_2 \\ \phi_{3i} &= \beta_3 + u_{3j}\end{aligned}$$

(9)

and

$$\mathbf{u}_j = \begin{bmatrix} u_{1j} \\ u_{3j} \end{bmatrix} \sim N(\mathbf{0}, \mathbf{\Sigma}), \text{ diagonal } \mathbf{\Sigma} = \begin{bmatrix} \sigma_{u_1}^2 & 0 \\ 0 & \sigma_{u_3}^2 \end{bmatrix}, \text{ and } \epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$$

(10)

➤ Example 6: Fitting an NLME model for the earthquake data

Let’s fit the model defined by (8), (9), and (10) by using `menl`.

```
. use http://www.stata-press.com/data/r15/earthquake
(Earthquake intensity (Joyner and Boore, 1981))

. menl laccel = {phi1:}-log10(sqrt(c.distance#c.distance*exp({phi2})))
> -{phi3:}*sqrt(c.distance#c.distance*exp({phi2})),
> define(phi1: richter U1[quake]) define(phi3: U3[quake], xb)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = 2.4115811
Iteration 2: linearization log likelihood = 2.40751409
Iteration 3: linearization log likelihood = 2.40734699
Iteration 4: linearization log likelihood = 2.40734238
Iteration 5: linearization log likelihood = 2.40734121
Iteration 6: linearization log likelihood = 2.40734113
```

Computing standard errors:

```
Mixed-effects ML nonlinear regression      Number of obs      =      182
Group variable: quake                     Number of groups   =       23
                                           Obs per group:
                                           min =           1
                                           avg =          7.9
                                           max =          38

                                           Wald chi2(1)       =      26.26
                                           Prob > chi2        =      0.0000

Linearization log likelihood = 2.4073411
      phi1: richter U1[quake]
      phi3: U3[quake], xb
```

laccel		Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1	richter	.2310021	.0450804	5.12	0.000	.1426461	.319358
	_cons	-.8836537	.2826255	-3.13	0.002	-1.437589	-.329718
phi3	_cons	.004575	.0014192	3.22	0.001	.0017935	.0073566
	/phi2	4.063075	.4023386	10.10	0.000	3.274506	4.851644

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
quake: Independent				
var(U1)	.0056676	.0073404	.0004477	.071752
var(U3)	.000013	8.42e-06	3.66e-06	.0000463
var(Residual)	.0461647	.0054421	.0366409	.0581639

We also store our estimates for later use:

```
. estimates store E1
```

By default, `menl` assumes that the random effects u_{1j} and u_{3j} are independent, so there is no need to specify the `covariance()` option in this case. In other words, omitting the `covariance()` option is equivalent to specifying `covariance(U1 U3, independent)`.



Example 7: Likelihood-ratio test for variance components

Davidian and Giltinan (1995) did not include any random effects in the model for the ϕ_{2j} parameters. Let's check whether the random effects are needed in the equations for ϕ_{1j} and ϕ_{3j} parameters in (9).

One simple way to assess whether a random effect associated with a certain ϕ_j can be omitted, is to examine its coefficient of variation (CV), the ratio of the standard deviation to the mean. Let's compute the CV for ϕ_{3j} . For convenience, let's redisplay the results from example 6 as standard deviations for variance components.

```
. menl, stddeviations
Mixed-effects ML nonlinear regression      Number of obs      =      182
Group variable: quake                     Number of groups   =       23
                                           Obs per group:
                                           min =           1
                                           avg =          7.9
                                           max =           38
                                           Wald chi2(1)      =      26.26
                                           Prob > chi2       =      0.0000

Linearization log likelihood = 2.4073411
      phi1: richter U1[quake]
      phi3: U3[quake], xb
```

laccel	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]		
phi1	richter	.2310021	.0450804	5.12	0.000	.1426461	.319358
	_cons	-.8836537	.2826255	-3.13	0.002	-1.437589	-.329718
phi3	_cons	.004575	.0014192	3.22	0.001	.0017935	.0073566
	/phi2	4.063075	.4023386	10.10	0.000	3.274506	4.851644

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
quake: Independent					
	sd(U1)	.0752832	.0487517	.0211582	.2678656
	sd(U3)	.0036085	.0011673	.0019142	.0068026
	sd(Residual)	.2148596	.0126644	.1914181	.241172

The `stddeviations` option specifies that `menl` display random-effects and error standard deviations instead of variances. It will also display correlations instead of covariances whenever they are in the model. Because random-effects variances for these data are very small, we will use this option in all subsequent examples to display results in the standard deviation metric.

The interquake random variation in the ϕ_{3j} values about their mean is $CV = sd(U3)/\{\text{phi3:}_{\_cons}\} = 0.0036/0.0046 \approx 78\%$, and it appears reasonable to keep it in the model. You can perform a formal likelihood-ratio (LR) test of $H_0: \sigma_{u_3}^2 = 0$ to verify this, as we show below for the test of $H_0: \sigma_{u_1}^2 = 0$.

Let’s check whether we need random intercept u_{1j} to model ϕ_{1j} . Computing CV in this case to get an initial assessment is not simple because the mean of ϕ_{1j} depends on the j th quake through variable `richt`. Given the same main equation (8), we will use the LR test to compare the restricted model, with u_{1j} excluded, which is defined by (11) and (12) below, with the full model defined by (9) and (10).

The stage 2 specification of the restricted model is

$$\begin{aligned}\phi_{1j} &= \beta_0 + \beta_1 \text{richt}_{1j} \\ \phi_{2j} &= \beta_2 \\ \phi_{3ij} &= \beta_3 + u_{3j}\end{aligned}\tag{11}$$

where

$$u_{3j} \sim N(0, \sigma_{u_3}^2) \quad \text{and} \quad \epsilon_{ij} \sim N(0, \sigma_{\epsilon}^2)\tag{12}$$

We now fit the restricted model:

```
. menl laccel = {phi1:}-log10(sqrt(c.distance#c.distance*exp({phi2})))
> -{phi3:}*sqrt(c.distance#c.distance*exp({phi2})),
> define(phi1: richter, xb) define(phi3: U3[quake], xb)
> stddeviations
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = 2.12628622
Iteration 2: linearization log likelihood = 2.126043
Iteration 3: linearization log likelihood = 2.12603276
Iteration 4: linearization log likelihood = 2.12603003
Iteration 5: linearization log likelihood = 2.1260297
```

Computing standard errors:

```
Mixed-effects ML nonlinear regression      Number of obs      =      182
Group variable: quake                     Number of groups   =       23
Obs per group:
      min =          1
      avg =         7.9
      max =         38
Wald chi2(1) =      32.22
Prob > chi2   =      0.0000
Linearization log likelihood = 2.1260297
```

```
phi1: richter, xb
phi3: U3[quake], xb
```

	laccel	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1	richter	.2208878	.0389144	5.68	0.000	.1446169	.2971586
	_cons	-.7863293	.2503442	-3.14	0.002	-1.276995	-.2956637
phi3	_cons	.0054348	.0015661	3.47	0.001	.0023653	.0085044
	/phi2	4.228431	.3702251	11.42	0.000	3.502803	4.954059

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
quake: Identity					
	sd(U3)	.0042144	.0011309	.0024907	.0071309
	sd(Residual)	.2170084	.0122821	.1942231	.2424668

```
. estimates store E2
```

Next, we use `lrtest` to perform an LR test of the hypothesis:

$$H_0: \sigma_{u_1}^2 = 0 \quad \text{versus} \quad H_1: \sigma_{u_1}^2 \neq 0$$

```
. lrtest E1 E2, stats
Likelihood-ratio test                                LR chi2(1) =      0.56
(Assumption: E2 nested in E1)                       Prob > chi2 =      0.4532
Note: The reported degrees of freedom assumes the null hypothesis is not on
      the boundary of the parameter space.  If this is not true, then the
      reported test is conservative.
Akaike's information criterion and Bayesian information criterion
```

Model	Obs	ll(null)	ll(model)	df	AIC	BIC
E2	182	.	2.12603	6	7.747941	26.97198
E1	182	.	2.407341	7	9.185318	31.61336

Note: N=Obs used in calculating BIC; see **[R] BIC note**.

Because testing of $H_0: \sigma_{u_1}^2 = 0$ is on the boundary of the parameter space, `lrtest` reports a note that the provided LR test is conservative; that is, the actual p -value is smaller than the one reported. For a test of $H_0: \sigma_{u_1}^2 = 0$ in a two-level model, the true asymptotic distribution is not $\chi^2(1)$ but a mixture of $\chi^2(0)$ and $\chi^2(1)$ with equal weights, $0.5\chi^2(0) + 0.5\chi^2(1)$; thus the p -value is actually $0.4532/2 = 0.2266$ (see [Rabe-Hesketh and Skrondal 2012](#), sec 8.8). We do not have sufficient evidence to reject the null hypothesis, so we can omit random effect u_{1j} from the full model. AIC and BIC also favor a simpler, reduced model.



➤ **Example 8: Including within-subject covariates**

One of the questions of interest in the earthquake study was the potential effect of the soil type on acceleration. Variable `soil` is a within-subject covariate because the values `soilij` may vary within a subject (earthquake). We include variable `soil` in the equation for ϕ_{3ij} in (11),

$$\begin{aligned}\phi_{1j} &= \beta_0 + \beta_1 \text{richter}_j \\ \phi_{2j} &= \beta_2 \\ \phi_{3ij} &= \beta_3 + \beta_4 \text{soil}_{ij} + u_{3j}\end{aligned}$$

and fit the corresponding model:

```
. menl laccel = {phi1:}-log10(sqrt(c.distance#c.distance+exp({phi2})))
> -{phi3:}*sqrt(c.distance#c.distance+exp({phi2})),
> define(phi1: richter, xb) define(phi3: i.soil U3[quake]) stddeviations
Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:   linearization log likelihood = 3.56347788
Iteration 2:   linearization log likelihood = 3.56324717
Iteration 3:   linearization log likelihood = 3.56323391
Iteration 4:   linearization log likelihood = 3.56323036
Iteration 5:   linearization log likelihood = 3.56322978
Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      182
Group variable: quake                     Number of groups   =       23
                                           Obs per group:
                                           min =              1
                                           avg =              7.9
                                           max =              38
                                           Wald chi2(2)       =      34.20
                                           Prob > chi2         =      0.0000
Linearization log likelihood = 3.5632298
      phi1: richter, xb
      phi3: i.soil U3[quake]
```

	laccel	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1	richter	.2275944	.0395549	5.75	0.000	.1500683	.3051206
	_cons	-.8079826	.2548833	-3.17	0.002	-1.307545	-.3084205
phi3	soil	-.0011041	.0006441	-1.71	0.087	-.0023665	.0001583
	_cons	.0067347	.0017416	3.87	0.000	.0033213	.0101481
	/phi2	4.3212	.3653809	11.83	0.000	3.605067	5.037334

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
quake: Identity					
	sd(U3)	.0043088	.0011285	.0025788	.0071992
	sd(Residual)	.2147101	.0121424	.1921829	.2398779

The estimated coefficient for the soil type is -0.0011 with a 95% CI of $[-0.0024, 0.0002]$. The knowledge of the soil type at a particular site does not appear to add explanatory power to our model.

◀

Random-effects covariance structures


menl supports various covariance structures to model the random-effects covariance matrix. They are specified using the `covariance()` option. The `covariance()` option may be repeated. This is necessary to accommodate multilevel NLME models, where you may need to specify different covariance matrices for the random effects at different levels. Repeating this option may also be useful if you want to specify a block-diagonal covariance structure. See [example 19](#) for details.

► Example 9: Two-level model with correlated random effects

Davidian and Giltinan (1995, sec. 1.1.3 and 11.2) discuss a study of soybean plants that started in 1988 and spanned over three growing seasons, `year`. The central objective of the study was to compare the growth patterns of two genotypes of soybean plants, `variety`: a commercial variety of soybean, denoted by F, and an experimental variety, denoted by P. In each season, eight plots were planted using F variety and eight using P variety. To assess growth, researchers sampled each plot 8 to 10 times ($8 \leq n_j \leq 10$) at approximately weekly intervals, `time`. At each sampling time, six plants were taken from each plot at random. Leaves from the plants were weighed, and the resulting total weight was divided by six to yield a measure of the average leaf weight per plant (in g) for the plot for that week, `weight`. Plots are identified by the `plot` variable.

Let's plot the data first.

```
. use http://www.stata-press.com/data/r15/soybean
(Growth of soybean plants (Davidian and Giltinan, 1995))
. twoway connected weight time if year==2, connect(L) by(variety)
```



The graph shows the average leaf weights per plant over time for the eight plots with plants of each genotype in the 1989 growing season. Longitudinal growth measures for each plot are connected with solid lines. Apart from some intraplot variation, the growth profile of each plot follows roughly an S shape, according to which growth begins slowly, then shows a linear trend during the middle of the growing season, and then “levels off” at the end. Such pattern is typical for many growth studies.

The main goal of the study was to compare growth patterns over the growing season for the two soybean genotypes. Because the three growing seasons differed markedly in terms of precipitation—1988 was unusually dry, 1989 was wet, and 1990 was normal—contrasting these growth patterns across years was also of interest. The results of this study are useful, for example, for harvesting purposes.

A popular model for individual profiles that resemble an S shape is the logistic growth model:

$$\text{weight}_{ij} = \frac{\phi_{1j}}{1 + \exp \{-(\text{time}_{ij} - \phi_{2j}) / \phi_{3j}\}} + \epsilon_{ij} \tag{13}$$

ϕ_{1j} is the asymptotic average leaf weight per soybean plant in plot j as $\text{time}_{ij} \rightarrow \infty$. ϕ_{2j} is the time at which half of ϕ_{1j} is reached; that is, if $\text{time}_{ij} = \phi_{2j}$, then $E(\text{weight}_{ij}) = 0.5\phi_{1j}$. ϕ_{1j}

and ϕ_{2j} will henceforth be referred to as “the limiting growth” and “half-life”, respectively. ϕ_{3j} is a scale parameter, and it represents the number of days it takes for average leaf weight to grow from 50% (half-life) to about 73% of its limiting growth. That is, if we set $\text{time}_{ij} = t_{0.73} = \phi_{2j} + \phi_{3j}$, the right-hand side of (13), ignoring the error term, reduces to $\phi_{1j}/\{1 + \exp(-1)\} = 0.73\phi_{1j}$, and then $\phi_{3j} = t_{0.73} - \phi_{2j}$.

We will start with a simple stage 2 specification that does not contain any covariates. Also, because the number of soybean plots, 48, is large compared with the number of random effects, 3, we consider a general positive-definite, unstructured, random-effects covariance matrix:

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} + \begin{bmatrix} u_{1j} \\ u_{2j} \\ u_{3j} \end{bmatrix} \quad (14)$$

$$\mathbf{u}_j = \begin{bmatrix} u_{1j} \\ u_{2j} \\ u_{3j} \end{bmatrix} \sim N(\mathbf{0}, \Sigma), \Sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} \end{bmatrix}, \epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$$

To specify this covariance structure in `menl`, we specify `unstructured` in the `covariance()` option. The `covariance()` option also requires that we list the names of random effects to be correlated.

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3:})),
> define(phi1: U1[plot], xb) define(phi2: U2[plot], xb) define(phi3: U3[plot], xb)
> covariance(U1 U2 U3, unstructured)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -739.901421

Iteration 2: linearization log likelihood = -739.84929

(iteration log omitted)

Iteration 39: linearization log likelihood = -739.834521

Iteration 40: linearization log likelihood = -739.834448

Computing standard errors:

Mixed-effects ML nonlinear regression

Group variable: plot

Number of obs = 412

Number of groups = 48

Obs per group:

min = 8

avg = 8.6

max = 10

Linearization log likelihood = -739.83445

phi1: U1[plot], xb

phi2: U2[plot], xb

phi3: U3[plot], xb

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	19.25314	.8031811	23.97	0.000	17.67893	20.82734
phi2						
_cons	55.01999	.7272491	75.65	0.000	53.59461	56.44537
phi3						
_cons	8.403468	.3152551	26.66	0.000	7.78558	9.021357

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
plot: Unstructured				
var(U1)	27.05081	6.776516	16.55561	44.19929
var(U2)	17.61605	5.317899	9.748766	31.83227
var(U3)	1.972036	.9849825	.7409021	5.248904
cov(U1,U2)	15.73304	5.413365	5.123042	26.34304
cov(U1,U3)	5.193819	2.165586	.9493488	9.438289
cov(U2,U3)	5.649306	2.049458	1.632442	9.66617
var(Residual)	1.262237	.1111686	1.062119	1.50006

The expected limiting growth or expected maximum average weight, $\beta_1 = E(\phi_{1j})$, of soybean leaves is estimated to be around 19.25 grams. The expected half-life or the time at which the leaves reach half of their expected maximum average weight, $\beta_2 = E(\phi_{2j})$, is estimated to be around 55 days after planting. The expected time needed for the average leaf weight per plant to grow from 50% to 73% of the limiting growth, $\beta_3 = E(\phi_{3j})$, is about 8.4 days.

The estimates of the six random-effects variance–covariance parameters σ_{11} , σ_{22} , σ_{33} , σ_{12} , σ_{13} , and σ_{23} are displayed in the upper part of the random-effects parameters table. There is a plot-to-plot variation in the estimates of all three parameters of interest: β_1 , β_2 , and β_3 . Also, the plot-specific effects associated with the parameters of interest are positively correlated. For example, based on the estimate of 5.19 of $\text{cov}(U1,U3)$, plants with larger maximum weights tend to grow faster.

We store our estimates for later use:


```
. estimates store S1
```



➤ **Example 10: Residuals-vs-fitted plot to check for heteroskedasticity**

A popular tool for investigating within-cluster heteroskedasticity is the plot of residuals against the predicted values and other candidate variance covariates. For growth models, variance is often a function of the mean (predicted values). Below we construct the plot of residuals versus predicted values to evaluate the assumption of homoskedastic errors in [example 9](#).

```
. predict fitweight, yhat
. predict res, residuals
. scatter res fitweight
```



The plot reveals increasing variability with the predicted average leaf weights, which indicates that our within-cluster variance model is misspecified. In [Heteroskedastic within-group errors](#), we will show how to account for within-cluster heteroskedasticity by using the `resvariance()` option.

◀

Heteroskedastic within-group errors

Until now, we assumed that the within-group errors—the ϵ 's in the considered models—are i.i.d. Gaussian with common variance σ_ϵ^2 , labeled as `var(Residual)` by `menl` in the output.

To relax the assumptions of homoskedasticity and the independence of errors, `menl` provides two alternatives. You can model the within-group error variance–covariance matrix, $\sigma^2 \mathbf{\Lambda}_j$, directly by using the `rescovariance()` option. If you used the `mixed` command and its `residuals()` option before, you should be familiar with this approach. Alternatively, you can model the error variance–covariance matrix indirectly by modeling the heteroskedasticity structure with the `resvariance()` option and the correlation structure with the `rescorrelation()` option; see [Variance-components parameters](#). The latter approach offers more flexibility, particularly in modeling the heteroskedasticity structure. For example, many NLME models exhibit within-subject heterogeneity that is a power function of the mean. The `rescovariance()` option cannot model this, but `resvariance(power _yhat)` can.

If your error structure is simple and is similar to those encountered in `mixed`, you can use the `rescovariance()` option. Otherwise, use `resvariance()`, `rescorrelation()`, or both to model more flexible within-group error covariance structures.

► Example 11: Heteroskedastic power structure

Continuing with [example 9](#), for these types of growth data, we find it is common for the intraplot variance to increase systematically with the average leaf weight, as we saw in [example 10](#) from the residuals-versus-fitted plot. [Davidian and Giltinan \(1995\)](#) proposed a variance structure that models the within-group error variance as a power function of the mean to account for the intraplot variability.

To reduce the number of parameters to be estimated, the authors assume that the random effects are independent.

Stage 2 specification of the model defined by (13) becomes

$$\boldsymbol{\phi}_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} + \begin{bmatrix} u_{1j} \\ u_{2j} \\ u_{3j} \end{bmatrix} \tag{15}$$

where

$$\mathbf{u}_j = \begin{bmatrix} u_{1j} \\ u_{2j} \\ u_{3j} \end{bmatrix} \sim N(\mathbf{0}, \boldsymbol{\Sigma}), \text{ diagonal } \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{u_1}^2 & 0 & 0 \\ 0 & \sigma_{u_2}^2 & 0 \\ 0 & 0 & \sigma_{u_3}^2 \end{bmatrix}$$

and

$$\text{Var}(\epsilon_{ij}) = \sigma^2(\widehat{\text{weight}}_{ij})^{2\delta}$$

Parameter σ^2 in the above is no longer an overall error variance σ_ϵ^2 but a common multiplier or a (squared) scale parameter.

In `menl`, this type of heteroskedasticity is modeled by specifying `resvariance(power _yhat, noconstant)`. `_yhat` designates that the variance should be modeled as a function of predicted values, `weightij`. By default, variance function `power` includes a constant, which we suppress by specifying the `noconstant` option.

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3:})),  
> define(phi1: U1[plot], xb) define(phi2: U2[plot], xb) define(phi3: U3[plot], xb)  
> resvariance(power _yhat, noconstant)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = -364.0225  
Iteration 2: linearization log likelihood = -364.228352  
Iteration 3: linearization log likelihood = -364.43168  
Iteration 4: linearization log likelihood = -364.383194  
Iteration 5: linearization log likelihood = -364.389644  
Iteration 6: linearization log likelihood = -364.389152  
Iteration 7: linearization log likelihood = -364.389201  
Iteration 8: linearization log likelihood = -364.389204
```

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	412
Group variable: plot	Number of groups	=	48
	Obs per group:		
	min	=	8
	avg	=	8.6
	max	=	10

```

Linearization log likelihood = -364.3892
      phi1: U1[plot], xb
      phi2: U2[plot], xb
      phi3: U3[plot], xb

```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	16.82289	.6030526	27.90	0.000	15.64093	18.00485
phi2						
_cons	51.74669	.4579629	112.99	0.000	50.8491	52.64428
phi3						
_cons	7.54537	.0856321	88.11	0.000	7.377534	7.713206

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
plot: Independent					
	var(U1)	11.32134	2.831139	6.934848	18.48241
	var(U2)	2.68911	.9344037	1.36093	5.31351
	var(U3)	5.44e-12	.	.	.
Residual variance:					
Power _yhat					
	sigma2	.0509223	.004422	.0429527	.0603706
	delta	.9339856	.0244477	.886069	.9819023

The near-zero estimate of the variance component of u_{3j} , $\text{var}(U3)$, suggests that the random-effects model is overparameterized. The within-group heteroskedasticity structure appears to explain enough variability in our data, and we no longer need random effects specific to ϕ_{3j} . This is quite common in mixed-effects models: the random-effects covariance structure and the within-group error covariance structure compete with each other, in the sense that fewer random effects are needed when the within-group error covariance structure is present, and vice versa.

Let's omit u_{3j} from (15) but now assume an unstructured covariance matrix for u_{1j} and u_{2j} . The EM algorithm used by `menl` to obtain initial values produces the starting values for variance components that are, in general, close to the final estimates upon convergence. Thus it can be used as a tool to help us detect potential convergence problems because of an overparameterized random-effects structure at an earlier stage. For example, we can check whether an unstructured covariance matrix is a reasonable choice for the random effects u_{1j} and u_{2j} for these data by displaying estimates after a few iterations. This can be done by specifying the `iterate(#)` option, where `#` is a small number of iterations, say, between 1 and 4. Below we specify `iterate(3)` to perform only three iterations and the `stddeviations` option to obtain standard deviations and correlations instead of variances and covariances for easier interpretability:

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3})),
> define(phi1: U1[plot], xb) define(phi2: U2[plot], xb)
> covariance(U*, unstructured) resvariance(power _yhat, noconstant)
> iterate(3) stddeviations
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1:   linearization log likelihood = -379.663433
Iteration 2:   linearization log likelihood = -362.909205
Iteration 3:   linearization log likelihood = -361.923807
```

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	412
Group variable: plot	Number of groups	=	48
	Obs per group:		
	min	=	8
	avg	=	8.6
	max	=	10

Linearization log likelihood = -361.94127

```
phi1:  U1[plot], xb
phi2:  U2[plot], xb
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	16.92772	.5676773	29.82	0.000	15.81509	18.04035
phi2						
_cons	51.81715	.4484055	115.56	0.000	50.93829	52.69601
/phi3	7.54089	.0869002	86.78	0.000	7.370569	7.711211

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
plot: Unstructured				
sd(U1)	2.904992	.4071122	2.207271	3.823263
sd(U2)	1.282251	.255516	.8676664	1.89493
corr(U1,U2)	-.99999	.0033707	-1	1
Residual variance:				
Power _yhat				
sigma	.225488	.0095078	.2076023	.2449146
delta	.955391	.0230645	.9101854	1.000597

Warning: convergence not achieved

The `U*` in `covariance(U*, unstructured)` is a shorthand notation to reference all random effects starting with `U`, that is, `U1` and `U2` in this example. The correlation between u_{1j} and u_{2j} is near -1 with a 95% CI of $[-1, 1]$, which indicates that the random-effects model may still be overparameterized. If you try to fit this model without the `iteration(3)` option, it would keep iterating without convergence.

Therefore, we further simplify the random-effects covariance structure by assuming independence between u_{1j} and u_{2j} . Stage 2 specification of the model defined by (13) is now

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 + u_{2j} \\ \beta_3 \end{bmatrix}$$

(16)

where

$$\mathbf{u}_j = \begin{bmatrix} u_{1j} \\ u_{2j} \end{bmatrix} \sim N(\mathbf{0}, \Sigma), \text{ diagonal } \Sigma = \begin{bmatrix} \sigma_{u_1}^2 & 0 \\ 0 & \sigma_{u_2}^2 \end{bmatrix}$$

and

$$\text{Var}(\epsilon_{ij}) = \sigma^2(\widehat{\text{weight}}_{ij})^{2\delta}$$

We fit this model and store its results as S2:

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3})),
> define(phi1: U1[plot], xb) define(phi2: U2[plot], xb)
> resvariance(power _yhat, noconstant)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = -402.761818
Iteration 2: linearization log likelihood = -372.425639
Iteration 3: linearization log likelihood = -364.009461
Iteration 4: linearization log likelihood = -364.371515
Iteration 5: linearization log likelihood = -364.38263
Iteration 6: linearization log likelihood = -364.38966
Iteration 7: linearization log likelihood = -364.389167
Iteration 8: linearization log likelihood = -364.389203
```

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	412
Group variable: plot	Number of groups	=	48
	Obs per group:		
	min	=	8
	avg	=	8.6
	max	=	10

Linearization log likelihood = -364.3892

```
phi1: U1[plot], xb
phi2: U2[plot], xb
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	16.82289	.6030523	27.90	0.000	15.64093	18.00485
phi2						
_cons	51.74669	.4579626	112.99	0.000	50.8491	52.64428
/phi3	7.545369	.085632	88.11	0.000	7.377534	7.713205

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
plot: Independent					
	var(U1)	11.32134	2.831139	6.934847	18.48241
	var(U2)	2.689111	.9344037	1.36093	5.31351
Residual variance:					
Power _yhat					
	sigma2	.0509223	.004422	.0429527	.0603706
	delta	.9339856	.0244477	.886069	.9819023

```
. estimates store S2
```

Because (16) is not nested in (14), we assess the adequacy of the heteroskedastic model by using information criteria. We use `estimates stats` to display the AIC and BIC values for the three models.


```
. estimates stats S1 S2
Akaike's information criterion and Bayesian information criterion
```

Model	Obs	ll(null)	ll(model)	df	AIC	BIC
S1	412	.	-739.8344	10	1499.669	1539.879
S2	412	.	-364.3892	7	742.7784	770.9256

Note: N=Obs used in calculating BIC; see **[R] BIC note**.

The heteroskedastic model defined by (16) has smaller AIC and BIC values and thus provides a much better representation of the data than (14).



► Example 12: Heteroskedastic model with interactions

The main goal of the soybean study was to compare growth patterns of the two genotypes of soybean over the three growing seasons, represented by calendar years 1988 through 1990. More specifically, we would like to compare the limiting growth, the half-life, and the growth rate of soybeans across growing seasons and genotypes.

Let $P_j = I(\text{variety}_j = P)$ be the indicator for genotype variety P, $S_{89,j} = I(\text{year}_j = 1989)$ be the indicator for growing season 1989, and $S_{90,j} = I(\text{year}_j = 1990)$ be the indicator for growing season 1990. Genotype variety F and growing season 1988 are baselines.

Consider an extension of the model defined by (13) and (16), where, in addition to random effects, ϕ_{1j} includes main and interaction effects of growing seasons and genotype variety, ϕ_{2j} includes main effects of growing seasons and genotype variety, and ϕ_{3j} contains main effects of growing seasons only.

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_{11} + \beta_{12}S_{89,j} + \beta_{13}S_{90,j} + \beta_{14}P_j + \beta_{15}S_{89,j} \times P_j + \beta_{16}S_{90,j} \times P_j + u_{1j} \\ \beta_{21} + \beta_{22}S_{89,j} + \beta_{23}S_{90,j} + \beta_{24}P_j + u_{2j} \\ \beta_{31} + \beta_{32}S_{89,j} + \beta_{33}S_{90,j} \end{bmatrix} \tag{17}$$

To fit the model defined by (13) and (17) by using **menl**, we extend **menl**'s specification from **example 11** by including the full-factorial interaction **i.year##i.variety** in the expression **{phi1:}**, main effects **i.year** and **i.variety** in the expression **{phi2:}**, and main effects **i.year** in the expression **{phi3:}**.

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3:})),
> define(phi1: i.year##i.variety U1[plot])
> define(phi2: i.year i.variety U2[plot])
> define(phi3: i.year, xb)
> resvariance(power _yhat, noconstant)
Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1: linearization log likelihood = -292.62615
Iteration 2: linearization log likelihood = -290.243889
(iteration log omitted)
Iteration 10: linearization log likelihood = -290.90729
Iteration 11: linearization log likelihood = -290.907297
```

Computing standard errors:

Mixed-effects ML nonlinear regression
Group variable: plot

```

Number of obs    =      412
Number of groups =       48
Obs per group:
    min =         8
    avg  =        8.6
    max  =        10

```

```

Wald chi2(10)    =     413.88
Prob > chi2      =      0.0000
Linearization log likelihood = -290.9073

```

```

phi1:  i.year i.variety i.year#i.variety U1[plot]
phi2:  i.year i.variety U2[plot]
phi3:  i.year

```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
year						
1989	-8.837933	1.056113	-8.37	0.000	-10.90788	-6.76799
1990	-3.666206	1.165969	-3.14	0.002	-5.951463	-1.380949
variety						
P	1.648139	1.033433	1.59	0.111	-.3773533	3.673631
year#variety						
1989#P	5.563008	1.167783	4.76	0.000	3.274196	7.851819
1990#P	.0974816	1.178054	0.08	0.934	-2.211462	2.406425
_cons	19.42734	.9445749	20.57	0.000	17.57601	21.27867
phi2						
year						
1989	-2.253227	.9746496	-2.31	0.021	-4.163505	-.3429493
1990	-4.970736	.9778318	-5.08	0.000	-6.887251	-3.054221
variety						
P	-1.294058	.4255317	-3.04	0.002	-2.128085	-.4600314
_cons	54.81257	.758724	72.24	0.000	53.3255	56.29964
phi3						
year						
1989	-.9023768	.1992358	-4.53	0.000	-1.292872	-.5118818
1990	-.6805314	.2100799	-3.24	0.001	-1.092281	-.2687823
_cons	8.060677	.1459662	55.22	0.000	7.774588	8.346765

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
plot: Independent					
	var(U1)	.8643051	.5271731	.2615089	2.856589
	var(U2)	.1341755	.2306995	.0046146	3.901367
Residual variance:					
	Power _yhat				
	sigma2	.0467091	.0039176	.0396286	.0550546
	delta	.9451193	.0227608	.9005089	.9897297

. estimates store S3

By including more fixed effects in the model, which explain some of the variability in the average leaf weight, we substantially reduced the estimates of variance components. Compared with [example 11](#), `var(U1)` decreased from 11.32 to 0.86, and `var(U2)` decreased from 2.69 to 0.13. It often happens that specifying a better-fitting model for the fixed effects reduces the need for random effects in the model.

We can compare model S3 or the model defined by [\(17\)](#) with model S2 or the one defined by [\(16\)](#) by using, for example, information criteria.

```
. estimates stats S2 S3
Akaike's information criterion and Bayesian information criterion
```

Model	Obs	ll(null)	ll(model)	df	AIC	BIC
S2	412	.	-364.3892	7	742.7784	770.9256
S3	412	.	-290.9073	17	615.8146	684.172

Note: N=Obs used in calculating BIC; see **[R] BIC note**.

Even though S3 has many more parameters, it fits the soybean data better than S2.

By inspecting the fixed-effects estimates from the output of model S3, we see that both the type of year and genotype variety affect all three parameters: the expected maximum leaf weight, half-life, and scale. For example, all three parameters achieve their highest values in the dry year, baseline year 1988, because coefficient estimates for the other years are negative. Also, the genotype variety P reaches its half-life roughly a day later ($\beta_{24} = -1.29$) than genotype variety F.



➤ **Example 13: Obtaining predictions**

After estimation, we may want to obtain predicted values for the outcome or for the parameters of interest. Continuing with [example 12](#), we want to predict the asymptotic average leaf weight per soybean plant in each plot, $\widehat{\phi}_{1j}$. The ϕ_{1j} parameter is not constant but varies for each plot, growing season, and genotype variety. We can use `predict` after `menl` to obtain predicted values for ϕ_{1j} ; see [\[ME\] menl postestimation](#).

First, we create a new grouping variable for growing seasons, genotype variety, and plot types. We also create the `tolist` variable to mark the first observation in each group.

```
. egen group = group(year variety plot)
. by group, sort: generate byte tolist=(_n==1)
```

Next, we use `predict` to compute predicted values for the expression `{phi1:}` and store them in the new variable `phi1`. We store only unique values in `phi1`, one for each group; the remaining observations are replaced with missing values.

```
. predict double (phi1 = {phi1:})
. qui replace phi1 = . if tolist!=1
```

We now list the five smallest and the five largest values of the asymptotic average leaf weight.

```
. sort phi1
. list plot year variety phi1 if (_n<=5 | _n>43) & phi1<., sep(5)
```

	plot	year	variety	phi1
1.	1989F6	1989	F	8.8421451
2.	1989F4	1989	F	10.449521
3.	1989F5	1989	F	10.473849
4.	1989F1	1989	F	10.721364
5.	1989F7	1989	F	10.810197
44.	1988P8	1988	P	20.86739
45.	1988P2	1988	P	21.237691
46.	1988P4	1988	P	21.310511
47.	1988P3	1988	P	21.506007
48.	1988P6	1988	P	21.581873

Soybean plants with genotype variety P have substantially larger asymptotic average leaf weight in the dry year, 1988, than soybean plants with genotype variety F in the wet year, 1989.

◀

► Example 14: Within-group error correlation structure

Pinheiro and Bates (2000, chap. 8) analyzed data from a study of the estrus cycles of mares. Originally analyzed in Pierson and Ginther (1987), the data contain daily records of the number of ovarian follicles larger than 10 mm over a period ranging from 3 days before ovulation to 3 days after the subsequent ovulation. The measurement times for each mare are scaled so that the ovulations for each mare occur at times 0 and 1 and are recorded in `stime`.

The considered model is

$$\text{follicles}_{ij} = \phi_{1j} + \phi_{2j} \sin(2\pi\phi_{3j}\text{stime}_{ij}) + \phi_{4j} \cos(2\pi\phi_{3j}\text{stime}_{ij}) + \epsilon_{ij}$$

where ϕ_{1j} is an intercept, ϕ_{3j} is the frequency of the sine wave for the j th mare, and ϕ_{2j} and ϕ_{4j} are terms determining the amplitude and phase of the sine wave for the j th mare. If a_j and p_j are the amplitude and phase for mare j , then $\phi_{2j} = a_j \cos(p_j)$ and $\phi_{4j} = a_j \sin(p_j)$.

This model was fit in example 8 of [ME] **mixed** in the context of a linear mixed-effects model, where the number of ovarian follicles was a periodic function of time with known frequency ϕ_{3j} equal to 1. If we want to estimate frequency, we cannot use the `mixed` command, because ϕ_{3j} enters the model nonlinearly.

Pinheiro and Bates (2000) suggested an AR(1) correlation structure for modeling the within-group error correlation. This structure can be specified by using the `rescorrelation()` option as `rescorrelation(ar 1, t(time))`, where `time` is an integer-valued time variable used to order the observations within mares and to determine the lags between successive observations.

We also considered several random-effects structures and found that we need only one random intercept to model ϕ_{1j} .

The full specification for the stage 2 model is

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \\ \phi_{4j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix}$$

where

$$\mathbf{u}_j = u_{1j} \sim N\left(0, \sigma_u^2\right), \epsilon_j \sim N(\mathbf{0}, \sigma_\epsilon^2 \mathbf{\Lambda}_j)$$

and

$$\sigma_\epsilon^2 \mathbf{\Lambda}_j = \sigma_\epsilon^2 \begin{bmatrix} 1 & \rho & \rho^2 & \dots & \rho^{n_j-1} \\ \rho & 1 & \rho & \dots & \rho^{n_j-2} \\ \rho^2 & \rho & 1 & \dots & \rho^{n_j-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{n_j-1} & \rho^{n_j-2} & \rho^{n_j-3} & \dots & 1 \end{bmatrix}$$

We fit this model by using **menl** as follows:

```
. use http://www.stata-press.com/data/r15/ovary, clear
(0varian follicles in mares)

. menl follicles = {phi1: U1[mare], xb} + {phi2}*sin(2*_pi*stime*{phi3}) +
> {phi4}*cos(2*_pi*stime*{phi3}), rescorrelation(ar 1, t(time))

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:   linearization log likelihood = -789.434153
Iteration 2:   linearization log likelihood = -789.434391
Iteration 3:   linearization log likelihood = -789.434391

Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      308
Group variable: mare                      Number of groups   =       11
                                           Obs per group:
                                           min =      25
                                           avg  =     28.0
                                           max  =      31

Linearization log likelihood = -789.43439
      phi1:  U1[mare], xb
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	11.98929	.9055946	13.24	0.000	10.21436	13.76422
/phi2	.2226033	.3290159	0.68	0.499	-.4222559	.8674626
/phi3	4.18747	.2746499	15.25	0.000	3.649166	4.725774
/phi4	.279653	.3223277	0.87	0.386	-.3520977	.9114036

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
mare: Identity	var(U1)	4.935352	3.967836	1.020903	23.85898
Residual: AR(1), time time	var(e)	20.14587	3.492937	14.34177	28.29888
	corr	.7332304	.0463231	.6287332	.8117157

By using estimates of ϕ_{2j} and ϕ_{4j} , we can compute the amplitude and phase for the sine wave for mare j . The amplitude and the phase are the same for all the mares because ϕ_{2j} and ϕ_{4j} are constant and not mare specific.

For example, the amplitude a_j can be computed as $\sqrt{\phi_{2j}^2 + \phi_{4j}^2}$ by using the relationship $\phi_{2j}^2 + \phi_{4j}^2 = a_j^2 \{\sin^2(p_j) + \cos^2(p_j)\} = a_j^2$. The phase p_j can be computed as $p_j = \text{atan}(\phi_{4j}/\phi_{2j})$ by using the relationship $\phi_{4j}/\phi_{2j} = \{a_j \sin(p_j)\} / \{a_j \cos(p_j)\} = \tan(p_j)$.

We can use `nlcom` to compute the amplitude and the phase.

```
. nlcom (amplitude: sqrt(_b[/phi2]^2 + _b[/phi4]^2))
> (phase: atan(_b[/phi4]/_b[/phi2]))
      amplitude:  sqrt(_b[/phi2]^2 + _b[/phi4]^2)
      phase:      atan(_b[/phi4]/_b[/phi2])
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
amplitude	.3574325	.2451183	1.46	0.145	-.1229904	.8378555
phase	.8985001	1.090985	0.82	0.410	-1.23979	3.03679

As we mentioned in [example 1](#), it is important to try different initial values when fitting NLME models to investigate potential convergence to a local maximum, especially for models containing periodic functions, as in our example. We explore different initial values for this model in [Linearization approach to finding initial values](#) by considering the functional form of the mean function and arrive at a different solution with a larger log likelihood.

◀

Restricted maximum likelihood

Like [mixed](#), `menl` provides estimation by using ML or REML. The difference between the two approaches is described in detail in [Likelihood versus restricted likelihood](#) in [\[ME\] mixed](#). Briefly, REML is preferable when you have a small number of groups because it produces unbiased, at least for balanced data, estimates of variance components. In large samples, there is little difference between ML and REML. One disadvantage of REML, however, is that LR tests based on REML are inappropriate for comparing models with different fixed-effects specifications.


► Example 15: Pharmacokinetics modeling

Pharmacokinetics (PKs) is the study of drug absorption, distribution, metabolism, and excretion. It is often referred to as the study of “what the body does with a drug”. The goal of PK modeling is to summarize the concentration-time measurements using a model that relates drug input to drug

response, to relate the parameters of this model to patient characteristics, and to provide individual dose-response predictions to optimize individual doses. In other words, by understanding between-subject variation in drug disposition, we can individualize the dosage regimen for a particular patient based on relevant physiological information identified by our PK model.

Consider a PK study of the antiasthmatic agent theophylline that was reported by [Boeckmann, Sheiner, and Beal \(\[1994\] 2011\)](#) and analyzed by [Davidian and Giltinan \(1995\)](#). The drug was administered orally to 12 subjects, where dosage dose (mg/kg) was given on a per weight basis. Serum concentrations (in mg/L) were obtained at 11 time points per subject over 25 hours following administration. The graph below shows the resulting concentration-time profiles for four subjects.

```
. use http://www.stata-press.com/data/r15/theoph
(Theophylline kinetics (Boeckmann et al., 1994))
. twoway connected conc time if subject<=4, connect(L) by(subject)
```



In PKs, the pattern of rapid rise to a peak concentration followed by an apparent exponential decay may be described by a so-called one-compartment open model with first-order absorption and elimination. The model corresponds roughly to viewing the body as one “blood compartment” and is particularly useful for the PK analysis of drugs that distribute relatively rapidly throughout the body, which makes it a reasonable model for the kinetics of theophylline after oral administration. Further details about compartmental modeling may be found in [Gibaldi and Perrier \(1982\)](#). The one-compartment open model for theophylline kinetics may be expressed as

$$\text{conc}_{ij} = \frac{\text{dose}_j k_{e_j} k_{a_j}}{\text{Cl}_j (k_{a_j} - k_{e_j})} \{ \exp(-k_{e_j} \text{time}_{ij}) - \exp(-k_{a_j} \text{time}_{ij}) \} + \epsilon_{ij} \quad (18)$$

for $i = 1, \dots, 11$ and $j = 1, \dots, 12$. Model parameters are the elimination rate constant k_{e_j} , the absorption rate constant k_{a_j} , and the clearance Cl_j for each subject j .

Because each of the model parameters must be positive to be meaningful, we write

$$\begin{aligned} \text{Cl}_j &= \exp(\beta_0 + u_{0j}) \\ k_{a_j} &= \exp(\beta_1 + u_{1j}) \\ k_{e_j} &= \exp(\beta_2) \end{aligned}$$

where u_{0j} and u_{1j} are assumed independent and normally distributed with means zero and variance $\sigma_{u_0}^2$ and $\sigma_{u_1}^2$, respectively.

The model defined by (18) implies that the predicted value for the concentration at time $\text{time}_{ij} = 0$ is $\widehat{\text{conc}}_{ij} = 0$. Therefore, a power variance function, a natural candidate for this type of heteroskedastic pattern, cannot be used in this example because error variance will be 0 at $\text{time}_{ij} = 0$. So the constant plus power variance function, which adds a constant to the power term, is used instead to model the within-group error variance:

$$\text{Var}(\epsilon_{ij}) = \sigma^2 \{ (\widehat{\text{conc}}_{ij})^\delta + c \}^2$$

In `menl`, we use the `resvariance(power _yhat)` option to specify the constant plus power variance function and the following model specification:

```
. menl conc = (dose*{ke:}*{ka:}/({cl:}*({ka:}-{ke:}))) *  
> (exp(-{ke:}*time)-exp(-{ka:}*time)), define(cl: exp({b0}+{U0[subject]}))  
> define(ka: exp({b1}+{U1[subject]})) define(ke: exp({b2}))  
> resvariance(power _yhat)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -167.519533

Iteration 2: linearization log likelihood = -167.65729

(iteration log omitted)

Iteration 26: linearization log likelihood = -167.679657

Iteration 27: linearization log likelihood = -167.679641

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	132
Group variable: subject	Number of groups	=	12
	Obs per group:		
	min	=	11
	avg	=	11.0
	max	=	11

Linearization log likelihood = -167.67964

```
cl: exp({b0}+{U0[subject]})  
ka: exp({b1}+{U1[subject]})  
ke: exp({b2})
```

conc	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b0	-3.227479	.0598389	-53.94	0.000	-3.344761	-3.110197
/b1	.432931	.1980835	2.19	0.029	.0446945	.8211674
/b2	-2.453742	.0514567	-47.69	0.000	-2.554595	-2.352889

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
subject: Independent					
	var(U0)	.0288787	.0127763	.0121337	.0687323
	var(U1)	.4075667	.1948713	.1596654	1.040367
Residual variance:					
Power _yhat	sigma2	.0976905	.0833027	.018366	.519624
	delta	.3187133	.2469511	-.1653019	.8027285
	_cons	.7288982	.3822952	.2607507	2.03755

The number of groups, 12, is fairly small in these data, so we now refit the model by using REML estimation.

```
. menl conc = (dose*{ke:}*{ka:}/({cl:}*({ka:}-{ke:}))) *
> (exp(-{ke:}*time)-exp(-{ka:}*time)), define(cl: exp({b0}+{U0[subject]}))
> define(ka: exp({b1}+{U1[subject]})) define(ke: exp({b2}))
> resvariance(power _yhat) reml

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1:   linearization log restricted-likelihood = -172.31734
Iteration 2:   linearization log restricted-likelihood = -172.423246
               (iteration log omitted)
Iteration 23:  linearization log restricted-likelihood = -172.443827
Iteration 24:  linearization log restricted-likelihood = -172.443845

Computing standard errors:
Mixed-effects REML nonlinear regression      Number of obs      =      132
Group variable: subject                     Number of groups   =       12
                                           Obs per group:
                                           min =           11
                                           avg  =          11.0
                                           max  =           11

Linear. log restricted-likelihood = -172.44384
      cl: exp({b0}+{U0[subject]})
      ka: exp({b1}+{U1[subject]})
      ke: exp({b2})
```

conc	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/b0	-3.227295	.0619113	-52.13	0.000	-3.348639	-3.105951
/b1	.4354519	.2072387	2.10	0.036	.0292716	.8416322
/b2	-2.453743	.0517991	-47.37	0.000	-2.555267	-2.352218

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
subject: Independent					
	var(U0)	.0316416	.014531	.0128634	.0778326
	var(U1)	.4500585	.2228206	.1705476	1.187661
Residual variance:					
Power _yhat	sigma2	.1015759	.086535	.0191263	.5394491
	delta	.3106636	.2466547	-.1727707	.7940979
	_cons	.7150935	.3745256	.2561837	1.996063

As expected, the estimates of the random-effects variances are slightly larger than the corresponding ML estimates, but we arrive at similar inferential conclusions based on our REML estimates.

► **Example 16: Nonlinear functions of parameters**

A distinctive feature of [example 15](#) is that parameters of interest are nonlinear functions of the estimated parameters and random effects. To interpret parameters that depend on random effects, we can either integrate random effects out of the parameter expression or condition on them. The former parameter estimates are often referred to as population-based estimates. The latter parameter estimates

are referred to as conditional estimates and, when conditioning on zero random effects, $\mathbf{u}_j = 0$, as estimates for an “average” or typical subject. For linear functions, the population-based estimates coincide with the conditional estimates. This is no longer true for nonlinear functions.

In PK modeling, the parameters of interest are clearance, elimination rate, and absorption rate. These are nonlinear functions of the estimated parameters β_0 , β_1 , β_2 and subject-specific random effects. Depending on the context, we may be interested in their population-based estimates or in their conditional estimates.

In general, obtaining population-based estimates would require numerical integration to integrate the subject-specific random effects out of the expression. In our example, we can compute population-based estimates directly by using the fact that $\exp(u_{0j})$ ’s and $\exp(u_{1j})$ ’s are lognormally distributed.

Thus the population-based clearance Cl^P can be computed as $E(\text{Cl}_j) = E\{\exp(\beta_0 + u_{0j})\} = \exp(\beta_0 + \sigma_{u_0}^2/2)$ and the population-based absorption rate k_a^P as $E\{\exp(\beta_1 + u_{1j})\} = \exp(\beta_1 + \sigma_{u_1}^2/2)$. The elimination rate k_e does not depend on subject-specific effects and can thus be computed simply as $k_e^P = k_e = \exp(\beta_2)$.

Alternatively, if we want parameters to represent a typical subject, we can simply set $u_{0j} = 0$ and $u_{1j} = 0$ in their expressions. Thus we can compute clearance and absorption rate for a typical subject simply as $\text{Cl} = \exp(\beta_0)$ and $k_a = \exp(\beta_1)$. These formulas can also be viewed as a result of exponentiating population-based log-clearance and log-absorption rate; that is, $\text{Cl} = \exp[E\{\log(\text{Cl}_j)\}] = \exp(\beta_0)$ and $k_a = \exp[E\{\log(k_{a_j})\}] = \exp(\beta_1)$.

If we compare the formulas for, say, Cl^P and Cl , the former considers variation in clearances across subjects, whereas the latter ignores such variation and instead reflects what the clearance would be for a typical subject with $u_{0j} = 0$.

Both approaches have merit, and here we will compute, for example, $\text{Cl}^P = \exp(\hat{\beta}_0 + \hat{\sigma}_{u_0}^2/2) = \exp(-3.23 + 0.032/2) = 0.04$. That is, 0.04 liters of serum concentration are cleared of the theophylline drug per hour per kg body weight in the considered population. In other words, for the population of subjects that weigh 75 kg, an average of $75 \times 0.04 \approx 3$ liters of serum concentration are cleared of theophylline every hour.

We can also use `nlcom` to compute the estimates of Cl^P and Cl . To use `nlcom`, we need to know how parameters are labeled by `menl` for postestimation. We can use `menl`’s `coeflegend` option to display parameter names. We also specify `noheader` to suppress the table header.

```
. menl, coeflegend noheader
```

conc	Coef.	Legend
/b0	-3.227295	_b[/b0]
/b1	.4354519	_b[/b1]
/b2	-2.453743	_b[/b2]
/subject		
lnsd(U0)	-1.726641	_b[/subject:lnsd(U0)]
lnsd(U1)	-.3991888	_b[/subject:lnsd(U1)]
/Residual		
lnsigma	-1.143475	_b[/Residual:lnsigma]
delta	.3106636	_b[/Residual:delta]
ln_cons	-.335342	_b[/Residual:ln_cons]

If we examine the output carefully, we will notice that `menl`, `coeflegend` displayed results in the estimation metric—as log-standard deviations instead of variances. Although by default `menl` displays

parameters in their original metric, it stores them in the estimation metric, the metric that was used during optimization; see [Examples of specifying initial values](#) and [Methods and formulas](#) for more details about the estimation metric.

The parameters we need to compute CI^P and CI are coefficient `_b[/b0]` and the variance of `U0`, which can be obtained as `exp(2*_b[/subject:lnsd(U0)])` based on the stored estimate of the log-standard deviation of `U0`. We now use `nlcom` to compute our nonlinear estimates.

```
. nlcom (C1_P: exp(_b[/b0]+0.5*exp(2*_b[/subject:lnsd(U0)]))) (C1: exp(_b[/b0]))
      C1_P: exp(_b[/b0]+0.5*exp(2*_b[/subject:lnsd(U0)]))
      C1: exp(_b[/b0])
```

conc	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
C1_P	.0402972	.002512	16.04	0.000	.0353738	.0452205
C1	.0396646	.0024557	16.15	0.000	.0348516	.0444777

Working with parameters in the estimation metric can be tedious, especially when nonlinear expressions contain multiple variance components. In that case, you may consider using `estat sd` after `menl` to obtain results in the standard deviation metric or, if you also specify the `variance` option, in the variance metric; see [\[ME\] menl postestimation](#). If you specify the `post` option with `estat sd`, the results will also be stored in the standard deviation or variance metrics, which you can use for further postestimation analysis.

```
. estat sd, post variance coeflegend
```

conc	Coef.	Legend
/b0	-3.227295	_b[/b0]
/b1	.4354519	_b[/b1]
/b2	-2.453743	_b[/b2]

Random-effects Parameters		Estimate	Legend
subject: Independent			
	var(U0)	.0316416	_b[/subject:var(U0)]
	var(U1)	.4500585	_b[/subject:var(U1)]
Residual variance:			
Power _yhat			
	sigma2	.1015759	_b[/Residual:sigma2]
	delta	.3106636	_b[/Residual:delta]
	_cons	.7150935	_b[/Residual:_cons]

In addition to results being displayed in the variance metric, because of the `post` option, they are stored in that metric. We also specified the `coeflegend` option with `estat sd` to see how parameters are labeled so that we could refer to them in other postestimation commands such as `nlcom`.

Now, we can simply refer to the variance of `U0` as `_b[/subject:var(U0)]` in our `nlcom` command.

```
. nlcom (Cl_P: exp(_b[/b0]+0.5*_b[/subject:var(U0)]))
      Cl_P: exp(_b[/b0]+0.5*_b[/subject:var(U0)])
```

	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
Cl_P	.0402972	.002512	16.04	0.000	.0353738	.0452205

`estat sd`'s `post` option should be used with caution because it clears all estimation results except the parameter estimates in `e(b)` and their VCE in `e(V)`. Thus the only postestimation features that will work after `estat sd`, `post` are those that need only `e(b)` and `e(V)`, such as `lincom` and `nlcom`. Other postestimation features will not be available, and you will need to refit your model to run them. To avoid refitting your model, you may consider storing your estimation results in memory (see [R] [estimates store](#)) or saving them on disk (see [R] [estimates save](#)) before using `estat sd`, `post`. We no longer needed the estimation results from `menl`, so we did not mind clearing them.

◀

Three-level models

Representation of (1) can be extended to, for example, two-nested levels of clustering, to form the following three-level model, with observations composing the first level,

$$\mathbf{y}_{jk} = \boldsymbol{\mu} \left(\mathbf{X}_{jk}, \boldsymbol{\beta}, \mathbf{u}_k^{(3)}, \mathbf{u}_{jk}^{(2)} \right) + \boldsymbol{\epsilon}_{jk}$$

where the first-level observations $i = 1, \dots, n_{jk}$ are nested within the second-level groups $j = 1, \dots, M_k$, which are nested within the third-level groups $k = 1, \dots, M$. Group j nested within group k consists of n_{jk} observations, so $\mathbf{y}_{jk}$, $\mathbf{X}_{jk}$, and $\boldsymbol{\epsilon}_{jk}$ each have row dimension n_{jk} .

Also, assume that

$$\mathbf{u}_k^{(3)} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_3) \quad \mathbf{u}_{jk}^{(2)} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_2) \quad \boldsymbol{\epsilon}_{jk} \sim N(\mathbf{0}, \sigma^2 \mathbf{\Lambda}_{jk})$$

and that $\mathbf{u}_k^{(3)}$, $\mathbf{u}_{jk}^{(2)}$, and $\boldsymbol{\epsilon}_{jk}$ are independent.

► Example 17: Three-level model

[Hand and Crowder \(1996, 118–120\)](#) analyzed a study where the blood glucose levels `glucose` of 7 volunteers, `subject`, who took alcohol at time 0 were measured 14 times, `time`, over a period of 5 hours after alcohol consumption. The same experiment was repeated at a later date with the same subjects but with a dietary additive, `guar`, used for all subjects. Variable `guar` is a binary variable that identifies whether a subject received a dietary additive. It also identifies each experiment, with 0 corresponding to the experiment without `guar` and 1 corresponding to the experiment with `guar`. Thus we will use the `guar` variable both as the level indicator and, later, as a fixed-effects variable.

Here is a plot of the whole dataset.

```
. use http://www.stata-press.com/data/r15/glucose
(Glucose levels following alcohol ingestion (Hand and Crowder, 1996))
. twoway connected glucose time if guar==0 ||
> connected glucose time if guar==1 ||, by(subject, rows(2))
> legend(order(1 "without guar" 2 "with guar"))
```



Our preliminary assessment based on the above graph is that, except for subject 6, the effect of the dietary additive guar on the temporal trajectory of the blood glucose levels does not seem to be important. The effect of guar will be formally tested in [example 18](#).

[Hand and Crowder \(1996\)](#) proposed the following empirical model relating the expected glucose level to time,

$$\text{glucose}_{ijk} = \phi_{1jk} + \phi_{2jk} \text{time}^3 \exp(-\phi_{3jk} \text{time}) + \epsilon_{ijk} \quad (19)$$

where $k = 1, \dots, 7$, $j = 1, 2$, and $i = 1, \dots, 14$. The blood glucose level is ϕ_1 at $\text{time} = 0$ and as $\text{time} \rightarrow \infty$. This is intentional, so that ϕ_1 can be interpreted as both the blood glucose level before ingesting alcohol and the blood glucose level after the effect of alcohol ingestion has washed out.

[Pinheiro and Bates \(2000, exercise 3, 412\)](#) analyzed this dataset in the context of a three-level NLME model. They initially proposed the following stage 2 specification,

$$\begin{aligned} \phi_{1jk} &= \beta_1 + u_{1k}^{(3)} + u_{1j,k}^{(2)} \\ \phi_{2jk} &= \beta_2 + u_{2k}^{(3)} + u_{2j,k}^{(2)} \\ \phi_{3jk} &= \beta_3 \end{aligned} \quad (20)$$

$$\mathbf{u}_k^{(3)} = \begin{bmatrix} u_{1k}^{(3)} \\ u_{2k}^{(3)} \end{bmatrix} \sim N(\mathbf{0}, \Sigma_3) \quad \mathbf{u}_{j,k}^{(2)} = \begin{bmatrix} u_{1j,k}^{(2)} \\ u_{2j,k}^{(2)} \end{bmatrix} \sim N(\mathbf{0}, \Sigma_2) \quad \epsilon_{ijk} \sim N(0, \sigma_\epsilon^2)$$

where Σ_2 and Σ_3 are general symmetric covariance matrices. $u_{1j,k}^{(2)}$ and $u_{2j,k}^{(2)}$ are random intercepts at the guar-within-subject level and can be specified in **menl** as `UU1[subject>guar]` and `UU2[subject>guar]`.

The full model defined by (19) and (20) contains many parameters. We will follow our own advice from [example 11](#) and specify the `iterate()` option to check how reasonable our model is for the data we have.

```
. menl glucose = {phi1:} + {phi2:}*c.time#c.time#c.time*exp(-{phi3}*time),
> define(phi1: U1[subject] UU1[subject>guar])
> define(phi2: U2[subject] UU2[subject>guar])
> covariance(U1 U2, unstructured) covariance(UU*, unstructured)
> stddeviations iterate(3)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -189.447112

Iteration 2: linearization log likelihood = -189.441164

Iteration 3: linearization log likelihood = -189.441123

Computing standard errors:

Mixed-effects ML nonlinear regression Number of obs = 196

Path	No. of Groups	Observations per Group Minimum Average Maximum
subject	7	28 28.0 28
subject>guar	14	14 14.0 14

Linearization log likelihood = -189.44112

phi1: U1[subject] UU1[subject>guar]

phi2: U2[subject] UU2[subject>guar]

glucose	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	3.661565	.1160345	31.56	0.000	3.434142	3.888989
phi2						
_cons	.4283298	.0530028	8.08	0.000	.3244462	.5322134
/phi3	.5896813	.013861	42.54	0.000	.5625144	.6168483

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
subject: Unstructured				
sd(U1)	.2624559	.0926845	.1313592	.5243873
sd(U2)	.0598426	.0724586	.0055765	.6421878
corr(U1,U2)	-.1489676	.9201413	-.9636339	.9346912
subject>guar: Unstructured				
sd(UU1)	.0919531	.0764232	.0180353	.468823
sd(UU2)	.1227072	.0412876	.0634547	.2372882
corr(UU1,UU2)	.99999	.0047062	-1	1
sd(Residual)	.5712261	.0305339	.514409	.6343187

Warning: convergence not achieved

The estimated correlation `corr(UU1,UU2)` is near one with the confidence interval spanning the entire range for the correlation parameter, which indicates that the random-effects structure is overparameterized. The confidence interval for `corr(U1,U2)` contains zero, which suggests that this term does not contribute much to explaining between-subject variability. If we try to fit this model without the `iterate()` option, it will continue iterating without convergence.

We simplify our model by assuming independence between random effects; that is, we assume that random-effects covariance matrices Σ_2 and Σ_3 are diagonal.

Recall that `covariance(, independent)` is assumed by default, so we do not need to explicitly specify the `covariance()` option:

```
. menl glucose = {phi1:} + {phi2:}*c.time#c.time#c.time*exp(-{phi3}*time),
> define(phi1: U1[subject] UU1[subject>guar])
> define(phi2: U2[subject] UU2[subject>guar]) stddeviations
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = -190.355295
Iteration 2: linearization log likelihood = -190.36034
Iteration 3: linearization log likelihood = -190.363301
Iteration 4: linearization log likelihood = -190.364176
Iteration 5: linearization log likelihood = -190.363753
Iteration 6: linearization log likelihood = -190.363967
Iteration 7: linearization log likelihood = -190.363859
Iteration 8: linearization log likelihood = -190.363914
Iteration 9: linearization log likelihood = -190.363886
```

Computing standard errors:

Mixed-effects ML nonlinear regression Number of obs = 196

Path	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
subject	7	28	28.0	28
subject>guar	14	14	14.0	14

Linearization log likelihood = -190.36389

```
phi1: U1[subject] UU1[subject>guar]
phi2: U2[subject] UU2[subject>guar]
```

glucose	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	3.658712	.1168642	31.31	0.000	3.429662	3.887762
phi2						
_cons	.4239173	.0526333	8.05	0.000	.320758	.5270766
/phi3	.5876636	.0137214	42.83	0.000	.5607701	.6145571

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
subject: Independent					
	sd(U1)	.2685609	.092104	.137126	.5259757
	sd(U2)	.0422075	.1078501	.0002821	6.315554
subject>guar: Independent					
	sd(UU1)	.0666034	.1527526	.0007435	5.96621
	sd(UU2)	.1362263	.0433548	.0730065	.2541912
	sd(Residual)	.5732488	.0309928	.5156118	.6373288

The random-effects structure may still be overparameterized, given small estimates for `sd(U2)` and `sd(UU1)`. If we were to perform an LR test of the corresponding variance components being zero,

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
subject: Identity					
	sd(U1)	.2453634	.1013232	.1092206	.5512071
subject>guar: Identity					
	sd(UU2)	.1011852	.0276419	.0592358	.1728421
Residual: CTAR1, time time					
	sd(e)	.6208598	.0412948	.5449771	.7073086
	corr	.6547722	.0564848	.5440641	.7654804

The dietary additive guar does not seem to affect the blood-glucose-level profiles over time. This actually conforms with the plot of the data from [example 17](#), where, except for subject 6, the profiles with and without guar are similar.



➤ **Example 19: Three-level model with block-diagonal covariance matrix**

[Pinheiro and Bates \(2000\)](#) report the data from the experiment conducted by Microelectronics Division of Lucent Technologies to study the variability in the manufacturing of analog MOS circuits. The intensities of the **current** (in mA) were collected on n -channel devices at five ascending **voltages**: 0.8, 1.2, 1.6, 2.0, and 2.4 V. Measurements were made on 8 **sites** of each of 10 **wafers**. The main objective of the study was to build an empirical model to simulate the behavior of similar circuits.

The intensity of the current at the i th level of **voltage** in the j th **site** within the k th **wafer** is expressed as

$$\text{current}_{ijk} = \phi_{1jk} + \phi_{2jk} \cos(\phi_{3jk} \text{voltage}_i + \pi/4) + \epsilon_{ijk}$$

where

$$\begin{aligned} \phi_{1jk} &= \beta_0 + u_{0k}^{(3)} + u_{0j,k}^{(2)} + \left(\beta_1 + u_{1k}^{(3)} + u_{1j,k}^{(2)} \right) \text{voltage}_i + \left(\beta_2 + u_{2k}^{(3)} + u_{2j,k}^{(2)} \right) \text{voltage}_i^2 \\ \phi_{2jk} &= \beta_3 + u_{3k}^{(3)} + u_{3j,k}^{(2)} \\ \phi_{3jk} &= \beta_4 + u_{4k}^{(3)} \end{aligned}$$
$$\mathbf{u}_k^{(3)} = \begin{bmatrix} u_{0k}^{(3)} \\ u_{1k}^{(3)} \\ u_{2k}^{(3)} \\ u_{3k}^{(3)} \\ u_{4k}^{(3)} \end{bmatrix} \sim N(\mathbf{0}, \Sigma_3) \quad \mathbf{u}_{j,k}^{(2)} = \begin{bmatrix} u_{0j,k}^{(2)} \\ u_{1j,k}^{(2)} \\ u_{2j,k}^{(2)} \\ u_{3j,k}^{(2)} \end{bmatrix} \sim N(\mathbf{0}, \Sigma_2) \quad \epsilon_{ijk} \sim N(0, \sigma_\epsilon^2)$$

Parameters β_0 , β_1 , and β_2 characterize the quadratic component of the model, and amplitude β_3 and frequency β_4 characterize the periodic component represented by the cosine wave.

For illustration, consider the following random-effects covariance structures:

$$\Sigma_3 = \begin{bmatrix} \sigma_{11}^{(3)} & & & & \\ & \sigma_{22}^{(3)} & & & \\ & & \sigma_{33}^{(3)} & & \\ & & & \sigma_{44}^{(3)} & \\ & & & & \sigma_{55}^{(3)} \end{bmatrix} \quad \Sigma_2 = \begin{bmatrix} \sigma_{11}^{(2)} & \sigma_{12}^{(2)} & 0 & 0 \\ \sigma_{12}^{(2)} & \sigma_{22}^{(2)} & 0 & 0 \\ 0 & 0 & \sigma_{33}^{(2)} & \sigma_{34}^{(2)} \\ 0 & 0 & \sigma_{34}^{(2)} & \sigma_{44}^{(2)} \end{bmatrix}$$

If we were to fit this model by using `menl`, we would type

```
. use http://www.stata-press.com/data/r15/wafer
(Modeling of analog MOS circuits)

. menl current = {phi1:}+{phi2:}*cos({phi3:}*voltage + _pi/4),
> define(phi1: voltage c.voltage#c.voltage W0[wafer] S0[wafer>site]
> c.voltage#(W1[wafer] S1[wafer>site])
> c.voltage#c.voltage#(W2[wafer] S2[wafer>site]))
> define(phi2: W3[wafer] S3[wafer>site]) define(phi3: W4[wafer], xb)
> covariance(S0 S1, unstructured) covariance(S2 S3, unstructured)
> covariance(W*, independent) stddeviations
```

In the specification above, Σ_3 is specified as `covariance(W*, independent)`, although this specification could have been omitted because `independent` is `menl`'s default random-effects covariance structure. The block-diagonal matrix Σ_2 is specified by using repeated `covariance()` options: `covariance(S0 S1, unstructured)` and `covariance(S2 S3, unstructured)`. If we tried to run this model, we would find out that it is overparameterized.

Because of the large number of random effects at each grouping level, to avoid numerically unstable estimates, we will further simplify our model by assuming independence between $u_{2j,k}^{(2)}$ and $u_{3j,k}^{(2)}$, which implies that $\sigma_{34}^{(2)} = 0$:

$$\Sigma_3 = \begin{bmatrix} \sigma_{11}^{(3)} & & & & \\ & \sigma_{22}^{(3)} & & & \\ & & \sigma_{33}^{(3)} & & \\ & & & \sigma_{44}^{(3)} & \\ & & & & \sigma_{55}^{(3)} \end{bmatrix} \quad \Sigma_2 = \begin{bmatrix} \sigma_{11}^{(2)} & \sigma_{12}^{(2)} & 0 & 0 \\ \sigma_{12}^{(2)} & \sigma_{22}^{(2)} & 0 & 0 \\ 0 & 0 & \sigma_{33}^{(2)} & 0 \\ 0 & 0 & 0 & \sigma_{44}^{(2)} \end{bmatrix}$$

We now try to fit the above simpler model. Note that given the complexity of this model, it takes some time to execute.

```
. use http://www.stata-press.com/data/r15/wafer
(Modeling of analog MOS circuits)

. menl current = {phi1:}+{phi2:}*cos({phi3:}*voltage + _pi/4),
> define(phi1: voltage c.voltage#c.voltage W0[wafer] S0[wafer>site]
> c.voltage#(W1[wafer] S1[wafer>site])
> c.voltage#c.voltage#(W2[wafer] S2[wafer>site]))
> define(phi2: W3[wafer] S3[wafer>site]) define(phi3: W4[wafer], xb)
> covariance(S0 S1, unstructured) covariance(S2 S3, independent)
> covariance(W*, independent) stddeviations
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = 737.063827
Iteration 2: linearization log likelihood = 754.072576
```

Iteration 4: linearization log likelihood = 825.917104

Computing standard errors:

Mixed-effects MI nonlinear

Path	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
wafer	10	40	40.0	40
wafer>site	80	5	5.0	5

	Wald chi2(2)	=	8763.93
Linearization log likelihood = 825.91711	Prob > chi2	=	0.0000

```
phi1:  voltage c.voltage#c.voltage W0[wafer] S0[wafer>site]
       c.voltage#W1[wafer] c.voltage#S1[wafer>site]
       c.voltage#c.voltage#W2[wafer]
       c.voltage#c.voltage#S2[wafer>site]
phi2:  W3[wafer] S3[wafer>site]
phi3:  W4[wafer], xb
```

current	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
voltage	6.046937	.1022632	59.13	0.000	5.846504	6.247369
c.voltage#						
c.voltage	1.158782	.0159669	72.57	0.000	1.127487	1.190076
_cons	-4.658034	.0361763	-128.76	0.000	-4.728938	-4.58713
phi2						
_cons	.1684428	.002054	82.01	0.000	.1644171	.1724686
phi3						
_cons	6.449391	.0019631	3285.32	0.000	6.445543	6.453238

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
wafer: Independent				
sd(W0)	.1107108	.0262518	.0695589	.1762087
sd(W1)	.3041975	.0764653	.1858624	.4978744
sd(W2)	.0449994	.0125441	.026057	.0777122
sd(W3)	.0057862	.0016144	.0033489	.0099974
sd(W4)	.0061349	.0013878	.0039377	.0095579
wafer>site: Unstructured				
sd(S0)	.0729495	.006297	.0615952	.0863969
sd(S1)	.2930062	.0252425	.2474831	.346903
corr(S0,S1)	-.8113227	.0413367	-.8782248	-.7132762
wafer>site: Independent				
sd(S2)	.0627587	.0053067	.0531738	.0740712
sd(S3)	.0080611	.0006861	.0068227	.0095244
sd(Residual)	.0008407	.0000711	.0007122	.0009922

In this example, our primary focus was to demonstrate how to use `menl` to fit a block-diagonal random-effects covariance structure. But if we were to interpret our fixed-effects estimates, the average frequency of the cosine wave, $\beta_4 = E(\phi_{3jk})$, for example, is estimated to be $6.45V^{-1}$, with a corresponding estimated period of $2\pi/\hat{\beta}_4 \approx 0.97V$. Also, some of the estimates of standard deviations such as `sd(W2)`, `sd(W3)`, and `sd(W4)` are very small, which suggests that this model may still be too rich for the observed data. If we proceeded to further analyze these data, we would consider simpler models. For example, at the very least, we would have omitted the term `W3[wافر]` from this model.

◀

Obtaining initial values

Obtaining good starting or initial values is important for the estimation of many statistical models, but it is often crucial for the estimation of NLME models. NLME models are known to be sensitive to the initial values and to have difficulty converging. Highly nonlinear mean specification or complicated variance–covariance structures for random effects and errors can often lead to multiple solutions, which requires considering different sets of initial values.

By default, `menl` uses the EM algorithm to obtain initial values. This default routine works well in many cases but cannot be guaranteed to provide good initial values in all situations. Sometimes, you may need to specify your own initial values. Trying different initial values can also be useful to investigate the existence of multiple solutions and to verify convergence to a global maximum.

So far we have been “lucky” that all the examples worked without us having to specify initial estimates. You may not be that lucky with your data and model. So, in this section, we provide some guidance on how to find good initial values when the default initial values do not work well.

We present three approaches that you may choose to explore to find good initial estimates for the fixed effects. In some cases, you may also be able to obtain initial estimates for covariance parameters; see [Linearization approach to finding initial values](#).

Linearization approach to finding initial values

Sometimes, we can use an LME model to obtain initial values of the NLME model by holding some of the parameters fixed at specific values. We can then fit the resulting LME model by using the `mixed` command and use the corresponding estimates as initial values for the NLME model. We refer to this initialization method as the linearization method.

We could have used this method in [example 14](#) and [example 19](#), if the default EM method did not provide reasonable initial estimates. In any case, it is good practice to specify different initial values to investigate potential convergence of the algorithm to a local maximum.

For instance, in [example 14](#), we fit

$$\text{follicles}_{ij} = \phi_{1j} + \phi_{2j} \sin(2\pi\phi_{3j}\text{stime}_{ij}) + \phi_{4j} \cos(2\pi\phi_{3j}\text{stime}_{ij}) + \epsilon_{ij}$$

where

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \\ \phi_{4j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix}$$

This model is nonlinear because of the parameter ϕ_{3j} . To obtain initial values, we can hold ϕ_{3j} (or β_3) fixed at a specific value, say, $\beta_3 = 1$, thus making the above model linear,

$$\text{follicles}_{ij} = \phi_{1j} + \phi_{2j} \sin(2\pi\phi_{3j}\text{stime}_{ij}) + \phi_{4j} \cos(2\pi\phi_{3j}\text{stime}_{ij}) + \epsilon_{ij}$$

where

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \\ \phi_{4j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 \\ 1 \\ \beta_4 \end{bmatrix}$$

Or, more compactly,

$$\text{follicles}_{ij} = \beta_1 + u_{1j} + \beta_2 \sin(2\pi\text{stime}_{ij}) + \beta_4 \cos(2\pi\text{stime}_{ij}) + \epsilon_{ij}$$

Now that the model is linear, we can use the mixed command to obtain initial values for β_1 , β_2 , and β_4 to be used in `menl`. In the code below, variables `sin1` and `cos1` are $\sin(2\pi\text{stime}_{ij})$ and $\cos(2\pi\text{stime}_{ij})$, respectively, and `|| mare:` specifies a random intercept at the `mare` level (see [\[ME\] mixed](#)). Also, for consistency with example 13, we assume an AR(1) within-group error correlation structure:

```
. mixed follicles sin1 cos1 || mare:, residuals(ar 1, t(time)) nolog
Mixed-effects ML regression      Number of obs      =      308
Group variable: mare             Number of groups   =       11
                                Obs per group:
                                min =          25
                                avg =         28.0
                                max =          31
                                Wald chi2(2)       =      39.00
                                Prob > chi2        =      0.0000
Log likelihood = -776.51731
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
sin1	-2.958619	.4935054	-6.00	0.000	-3.925872	-1.991366
cos1	-.8798847	.5031764	-1.75	0.080	-1.866092	.1063228
_cons	12.18963	.9017435	13.52	0.000	10.42224	13.95701

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
mare: Identity					
	var(_cons)	7.095503	3.764874	2.508047	20.07385
Residual: AR(1)					
	rho	.5974664	.0547217	.4795551	.6941854
	var(e)	13.08097	1.765326	10.04078	17.0417

LR test vs. linear model: chi2(2) = 242.63 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

We will now use the estimates of the fixed effects shown in the output table as initial values for `menl` by specifying the `initial()` option. We use 1 as the initial value for `/phi3`. There are three ways to specify initial values in the `initial()` option; see [Specifying initial values](#). Here we will use the specification where we repeatedly list a parameter name followed by its initial value; also see [Examples of specifying initial values](#).

```
. local xb phi1:_cons 12.2 /phi2 -3.0 /phi3 1 /phi4 -.88
. menl follicles = {phi1: U1[mare], xb} + {phi2}*sin(2*_pi*stime*{phi3}) +
> {phi4}*cos(2*_pi*stime*{phi3}), rescorrelation(ar 1, t(time)) init('xb')
Alternating PNLS/LME algorithm:
Iteration 1:    linearization log likelihood = -775.627064
Iteration 2:    linearization log likelihood = -775.624331
Iteration 3:    linearization log likelihood = -775.624331
Computing standard errors:
Mixed-effects ML nonlinear regression      Number of obs      =      308
Group variable: mare                      Number of groups   =       11
                                           Obs per group:
                                           min =           25
                                           avg =          28.0
                                           max =           31
```

Linearization log likelihood = -775.62433

phi1: U1[mare], xb

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	12.18125	.9055136	13.45	0.000	10.40647	13.95602
/phi2	-2.874435	.5389255	-5.33	0.000	-3.930709	-1.81816
/phi3	.919119	.0512361	17.94	0.000	.8186981	1.01954
/phi4	-1.67526	.6766388	-2.48	0.013	-3.001447	-.3490722

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
mare: Identity					
	var(U1)	7.207072	3.755605	2.595361	20.01336
Residual: AR(1), time time					
	var(e)	12.63377	1.646898	9.785276	16.31146
	corr	.5823733	.0544508	.4656902	.679153

In the above, we initialized only fixed-effects parameters and used naïve initial estimates of 1 for random-intercept and error variances and 0 for the correlation. We could have specified `initial()`'s `fixed` suboption to use the EM algorithm to compute initial estimates for the random-effects parameters; see [Examples of specifying initial values](#) for details.

With the linearization approach, we can also use estimates of the random-effects parameters from the `mixed` command to initialize the corresponding parameters of `menl`. This is an advantage of the linearization approach over the other two approaches we discuss in subsequent sections. One complication with the initialization of random-effects parameters is that the initial values must be supplied in the estimation metric, the metric used during estimation, instead of the parameter original metric. For example, instead of variances, we must supply estimates of log-standard deviations, and instead of covariances or correlations, we must supply inverse hyperbolic tangents of correlation parameters. Luckily for us, `mixed` uses the same estimation metric as `menl` and provides the `estmetric` option to display parameters in the estimation metric.

In our example, the random-effects parameters are the random-intercept variance, the within-group error variance, and the correlation between error terms. We refit the earlier `mixed` command but now with the `estmetric` option to obtain the estimates of the random-effects parameters in the estimation metric.

```
. mixed follicles sin1 cos1 || mare:, residuals(ar 1, t(time)) nolog estmetric
Mixed-effects ML regression          Number of obs    =      308
Group variable: mare                 Number of groups =       11
                                     Obs per group:
                                     min =          25
                                     avg =         28.0
                                     max =          31
                                     Wald chi2(2)       =      39.00
                                     Prob > chi2        =      0.0000
Log likelihood = -776.51731
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
follicles						
sin1	-2.958619	.4935054	-6.00	0.000	-3.925872	-1.991366
cos1	-.8798847	.5031764	-1.75	0.080	-1.866092	.1063228
_cons	12.18963	.9017435	13.52	0.000	10.42224	13.95701
lns1_1_1						
_cons	.9797306	.2653			.5762499	1.665722
lnsig_e						
_cons	1.285579	.0674768	19.05	0.000	1.153327	1.417832
r_atr1						
_cons	.6891977	.0850992	8.10	0.000	.5224063	.8559892

menl uses the same ordering of the parameters as mixed does, so we can simply list all the estimates directly in the initial() option. When we list the values without parameter names, we must specify initial()'s copy suboption and specify the values for all parameters. In our example, we specify four fixed-effects coefficients and three random-effects parameters.

```
. menl follicles = {phi1: U1[mare], xb} + {phi2}*sin(2*_pi*time*{phi3}) +
> {phi4}*cos(2*_pi*time*{phi3}), rescorrelation(ar 1, t(time))
> initial(12.2 -3.0 1 -.88 .98 1.29 .69, copy)
Alternating PNLS/LME algorithm:
Iteration 1: linearization log likelihood = -775.624332
Iteration 2: linearization log likelihood = -775.624331
Computing standard errors:
Mixed-effects ML nonlinear regression          Number of obs    =      308
Group variable: mare                 Number of groups =       11
                                     Obs per group:
                                     min =          25
                                     avg =         28.0
                                     max =          31
Linearization log likelihood = -775.62433
phi1: U1[mare], xb
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	12.18125	.9055135	13.45	0.000	10.40647	13.95602
/phi2	-2.874434	.5389241	-5.33	0.000	-3.930706	-1.818162
/phi3	.919119	.0512356	17.94	0.000	.818699	1.019539
/phi4	-1.675261	.6766409	-2.48	0.013	-3.001452	-.3490689

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
mare: Identity	var(U1)	7.207072	3.755605	2.595361	20.01336
Residual: AR(1), time time					
	var(e)	12.63377	1.646898	9.785276	16.31146
	corr	.5823733	.0544508	.4656903	.679153

The results are different from those in [example 14](#). The value of the linearization log likelihood in this example, -775.62 , is larger than that from [example 14](#), -789.43 . So it appears that we have converged to a local maximum of the linearization log likelihood in [example 14](#).

Our initial values based on `mixed` turned out to be better than those computed by default by `menl`. This is not surprising. In general, `menl`'s EM algorithm should produce reasonable initial values for many nonlinear models, but the initial values may not necessarily be optimal for all of those models. In this example, our initial values were tailored to the ovary data and the model.

In general, sensitivity to initial values is one of the key issues in NLME models, especially for models that involve periodic functions. Therefore, it is important to try different sets of initial values to verify global convergence before reporting your final results. Sometimes, you may even have to rely on your knowledge of the science behind the problem to decide which set of results is more reasonable.

Graphical approach to finding initial values

If your model has parameters that have natural physical interpretations, you may be able to obtain starting values from a graph of the data.

[Draper and Smith \(1998\)](#) presented a dataset in which the trunk circumference `circumf` (in mm) of five different orange trees was measured over seven different time points, stored in `age`. [Pinheiro and Bates \(2000\)](#) suggested the following model for these data:

$$\text{circumf}_{ij} = \frac{\phi_{1j}}{1 + \exp \left\{ -(\text{age}_{ij} - \phi_{2j}) / \phi_{3j} \right\}} + \epsilon_{ij} \quad (21)$$

In this model, ϕ_{1j} is the asymptotic trunk circumference for the j th tree as $\text{age}_{ij} \rightarrow \infty$, ϕ_{2j} is the age at which the j th tree attains half of its asymptotic trunk circumference ϕ_{1j} , and ϕ_{3j} is a scale parameter; see the graph below.

The stage 2 specification of this model is

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 + u_{1j} \\ \beta_2 \\ \beta_3 \end{bmatrix}$$


where

$$u_{1j} \sim N(0, \sigma_{u_1}^2), \quad \epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$$

Because the model parameters have graphical interpretations, we can plot our data and obtain initial values from the graph.

```
. use http://www.stata-press.com/data/r15/orange
(Growth of orange trees (Draper and Smith, 1998))

. twoway connected circumf age, connect(L) yline(175) xline(1582)
> yline(87.5, lpattern(dash)) xline(700, lpattern(dash))
> yline(131.25, lpattern("-...")) xline(1000, lpattern("-..."))
> xlabel(0 118 484 700 1000 1372 1582) ylabel(#5 87.5 131.25 175)
```



From the above graph, the mean asymptotic trunk circumference can be estimated as 175 mm, which is roughly the mean of the circumference values at age 1,582 (in days). The trees attain half of their asymptotic trunk circumference, $175/2 = 87.5$, at about age 700 (in days). Therefore, we use the initial estimates $\beta_1 = 175$ for the asymptotic trunk circumference and $\beta_2 = 700$ for the location of the inflection point. To obtain an initial estimate for β_3 , we note that when $\text{age} = \beta_2 + \beta_3$ in (21), $E(\text{circumf}_{ij}) = \beta_1 / \{1 + \exp(-1)\} = 0.73\beta_1$, which we will approximate as $0.75\beta_1$ for the purpose of the graph. That is, the logistic curve reaches approximately 3/4 of its asymptotic value, $0.75 \times 175 = 131.25$, at $\text{age} = \beta_2 + \beta_3$. The above graph suggests that the trees attain 3/4 of their final trunk circumference at about 1,000 days ($= \beta_2 + \beta_3$), giving an initial estimate of $\beta_3 = 1000 - 700 = 300$. We can now supply these values to **menl** in the `initial()` option.

Unfortunately, the graph does not provide us with the estimates for variance components. In this case, we can use `initial()`'s `fixed` suboption to specify that the EM algorithm still be used to initialize variance components, while the supplied values be used to initialize fixed effects. If we do not specify `fixed`, **menl** will use naïve initial estimates for variance components such as ones for variances and zeros for covariances.

We now fit the model using our own initial estimates for fixed effects:

```
. menl circumf = {phi1: U1[tree], xb}/(1+exp(-(age-{phi2})/{phi3})),
> initial(phi1:_cons 175 /phi2 700 /phi3 300, fixed)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

Iteration 1: linearization log likelihood = -131.584945

Iteration 2: linearization log likelihood = -131.584579

Iteration 3: linearization log likelihood = -131.584579

Computing standard errors:

Mixed-effects ML nonlinear regression
Group variable: tree

Number of obs = 35
Number of groups = 5
Obs per group:

min = 7
avg = 7.0
max = 7

Linearization log likelihood = -131.58458
phi1: U1[tree], xb

circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	191.049	16.15403	11.83	0.000	159.3877	222.7103
/phi2	722.556	35.15082	20.56	0.000	653.6616	791.4503
/phi3	344.1624	27.14739	12.68	0.000	290.9545	397.3703

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Identity	var(U1)	991.1514	639.4636	279.8776	3510.038
	var(Residual)	61.56371	15.89568	37.11466	102.1184

For comparison, we fit the same model but now using the default initial values for fixed effects:

```
. menl circumf = {phi1: U1[tree], xb}/(1+exp(-(age-{phi2})/{phi3}))
```

Obtaining starting values by EM:
Alternating PNLS/LME algorithm:
Iteration 1: linearization log likelihood = -131.584579
Computing standard errors:

Mixed-effects ML nonlinear regression
Group variable: tree

Number of obs = 35
Number of groups = 5
Obs per group:

min = 7
avg = 7.0
max = 7

Linearization log likelihood = -131.58458
phi1: U1[tree], xb

circumf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
_cons	191.049	16.15403	11.83	0.000	159.3877	222.7103
/phi2	722.556	35.15082	20.56	0.000	653.6616	791.4503
/phi3	344.1624	27.14739	12.68	0.000	290.9545	397.3703

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
tree: Identity	var(U1)	991.1514	639.4636	279.8776	3510.038
	var(Residual)	61.56371	15.89568	37.11466	102.1184

The results are identical except for the iteration log.

Smart regressions approach to finding initial values

Consider the following NLME model,

$$y_{ij} = \phi_{1j} + (\phi_{2j} - \phi_{1j}) \exp \{-\exp(\phi_{3j}) x_{ij}\} + \epsilon_{ij}$$

where

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 + u_{1j} \\ \beta_3 \end{bmatrix}$$

Here ϕ_{1j} is the asymptote as $x_{ij} \rightarrow \infty$ and ϕ_{2j} is the value of y_{ij} at $x_{ij} = 0$. Thus initial estimates, $\beta_1^{(0)}$ and $\beta_2^{(0)}$, may be obtained by using the graphical approach as described in [Graphical approach to finding initial values](#). To obtain an initial estimate for β_3 , notice that, ignoring the error term ϵ_{ij} and setting $u_{1j} = 0$,

$$\log(|y_{ij} - \beta_1|) = \log(\beta_2 - \beta_1) + \{-\exp(\beta_3)\} x_{ij}$$

Therefore, we can regress $\log(|y - \beta_1^{(0)}|)$ on x and use the estimated slope, $\hat{\beta}_x = -\exp(\beta_3^{(0)})$, to obtain the initial value for $\beta_3^{(0)} = \log(-\hat{\beta}_x)$.

Examples of specifying initial values

When you want to assign initial values for a subset of the model parameters, for example, fixed effects or random-effects covariance parameters, you will often need to know their estimation names or, in other words, how **menl** labels them in **e(b)**. To learn the names, you can fit the model with the **iterate(0)** and **coeflegend** options first.

```
. menl ... , ... iterate(0) coeflegend
```

The **iterate(0)** option specifies to bypass maximization and only report the initial values and the likelihood evaluated at those values. The **coeflegend** option specifies that the legend of the parameters and how to specify them in an expression be displayed rather than displaying the statistics for the parameters.

Keep in mind, however, that **menl** does not perform estimation in the original parameter metric. For computational stability, the estimation is performed, loosely speaking, in a metric that transforms all parameters to be defined on a real line. For example, a log transformation is used for standard deviations, and an inverse hyperbolic tangent transformation is used for correlations. When you specify initial values, you must specify them for parameters in the estimation metric and not the original metric.

coeflegend displays parameter names as they are stored in **e(b)**, which, for **menl**, are the names of estimation parameters. If you also want to see parameters in the original metric, you can specify **coeflegend** on replay.

```
. menl ... , ... iterate(0)
. menl, coeflegend
```

For example, recall the NLME model for the soybean data from [example 9](#). Suppose that we want to supply our own initial values.

We fit the model with `iterate(0)` and `coeflegend`:

```
. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3:})),
> define(phi1: U1[plot], xb)
> define(phi2: U2[plot], xb)
> define(phi3: U3[plot], xb)
> covariance(U*, unstructured) iterate(0) coeflegend
```

Obtaining starting values by EM:

Computing standard errors:

Mixed-effects ML nonlinear regression	Number of obs	=	412
Group variable: plot	Number of groups	=	48
	Obs per group:		
	min	=	8
	avg	=	8.6
	max	=	10

Linearization log likelihood = -740.06177

```
phi1:  U1[plot], xb
phi2:  U2[plot], xb
phi3:  U3[plot], xb
```

weight	Coef.	Legend
phi1		
_cons	19.26527	_b[phi1:_cons]
phi2		
_cons	55.05299	_b[phi2:_cons]
phi3		
_cons	8.385531	_b[phi3:_cons]
/plot		
lnsd(U1)	1.650846	_b[/plot:lnsd(U1)]
lnsd(U2)	1.436634	_b[/plot:lnsd(U2)]
lnsd(U3)	.4081525	_b[/plot:lnsd(U3)]
athcorr(U2,		
U1)	.9055785	_b[/plot:athcorr(U2,U1)]
athcorr(U3,		
U1)	.8482105	_b[/plot:athcorr(U3,U1)]
athcorr(U3,		
U2)	1.537798	_b[/plot:athcorr(U3,U2)]
/Residual		
lnsigma	.1069986	_b[/Residual:lnsigma]

Warning: convergence not achieved

Parameter names are listed within the `_b[]` specifier.

In what follows, we will outline only the syntax of the specifications. If you actually want to run all the examples to see the initialization in action, we suggest that you specify `iterate(0)` for speed.

Let's first specify initial values for fixed effects only. The fixed-effects parameters are `phi1:_cons`, `phi2:_cons`, and `phi3:_cons`. Suppose that we want to initialize them with 19, 55, and 8.

We can type

```
. menl ..., ... initial(phi1:_cons 19 phi2:_cons 55 phi3:_cons 8)
```

Or, more compactly, we can type

```
. local fe phi1:_cons 19 phi2:_cons 55 phi3:_cons 8
. menl ..., ... initial('fe')
```

When you specify the `initial()` option, `menl` does not perform the EM algorithm to initialize the parameters but instead uses the values you supplied. If you specify values for only a subset of parameters, the remaining parameters will be initialized with naïve initial values such as zeros for fixed effects and correlations and ones for variances. Often, you may have good initial values for fixed effects but not for variance components. In this situation, `menl` provides `initial()`'s `fixed` suboption. This option specifies that the supplied values be used for fixed effects but that the EM algorithm still be used to obtain initial values for variance components. If you specify only a subset of values for fixed effects, the remaining fixed effects will still be initialized with zeros even if `fixed` is specified. We recommend that you specify `fixed` when you intend to supply initial values only for the fixed effects.

```
. local fe phi1:_cons 19 phi2:_cons 55 phi3:_cons 8
. menl ..., ... initial('fe', fixed)
```

Now suppose that we also want to assign initial values for random-effects parameters. As we mentioned earlier, remember that we assign initial values for standard deviations in the log metric and for correlation in the inverse hyperbolic tangent or `atanh` metric. For example, if you want to assign an initial value of 2 to σ_ϵ , then you should supply `log(2)` to the `initial()` option. Similarly, if you want to assign a value of 0.7 to the correlation of two random effects, then you should provide `atanh(0.7)` to the `initial()` option.

Continuing with [example 9](#), suppose that we want to specify the following initial values for the random-effects covariance parameters:

$$\begin{pmatrix} \sigma_1 = 5 & & \\ \rho_{21} = 0.72 & \sigma_2 = 4 & \\ \rho_{31} = 0.71 & \rho_{32} = 0.94 & \sigma_3 = 1.4 \end{pmatrix}$$

The names of the parameters in the estimation metric that correspond to σ_1 , σ_2 , and σ_3 are `/plot:lnsd(U1)`, `/plot:lnsd(U2)`, and `/plot:lnsd(U3)` and that correspond to ρ_{21} , ρ_{31} , and ρ_{32} are `/plot:athcorr(U2,U1)`, `/plot:athcorr(U3,U1)`, and `/plot:athcorr(U3,U2)`.

When specifying initial values for [free parameters](#) such as random-effects covariance parameters, you can omit the forward slash (/) at the beginning of their names. Keeping in mind that initial values for covariance parameters are supplied in the log and `atanh` metrics, we can type

```
. local re_cov      plot:lnsd(U1) log(5)           // log(5)
. local re_cov 're_cov' plot:lnsd(U2) 1.4          // log(4)
. local re_cov 're_cov' plot:lnsd(U3) 0.34         // log(1.4)
. local re_cov 're_cov' plot:athcorr(U2,U1) atanh(0.72) // atanh(0.72)
. local re_cov 're_cov' plot:athcorr(U3,U1) 0.89    // atanh(0.71)
. local re_cov 're_cov' plot:athcorr(U3,U2) 1.7     // atanh(0.94)
. menl ... , ... initial('fe' 're_cov' Residual:lnsigma 0.5)
```

In the above, we also specified an initial value of 0.5 for the log of the error standard deviation. For parameters `/plot:lnsd(U1)` and `/plot:athcorr(U2,U1)`, instead of specifying the values, we specified the corresponding expression. This is allowed, as long as your expression is simple and does not contain spaces.

Instead of using parameter names, we can specify a list of values directly in the `initial()` option, in which case we must also specify `initial()`'s `copy` suboption.

```
. menl ... , ... initial(19 55 8 1.6 1.4 0.34 0.9 0.89 1.7 0.5, copy)
```

Or we can provide these values as a matrix:

```
. matrix initvals = (19, 55, 8, 1.6, 1.4, 0.34, 0.9, 0.89, 1.7, 0.5)
. matrix list initvals
initvals[1,10]
      c1  c2  c3  c4  c5  c6  c7  c8  c9  c10
r1    19  55   8  1.6  1.4  .34  .9  .89  1.7  .5
. menl ... , ... initial(initvals, copy)
```

If we label the columns of the `initvals` matrix properly, we do not need to specify `copy`:

```
. local fullcolnames : colfullnames e(b)
. matrix colnames initvals = 'fullcolnames'
. matrix list initvals
initvals[1,10]
      phi1:      phi2:      phi3:      /plot:      /plot:
      _cons      _cons      _cons      lnsd(U1)      lnsd(U2)
r1      19      55      8      1.6      1.4
      /plot:      /plot:      /plot:      /plot:      /Residual:
      lnsd(U3)      athcorr(U2,      athcorr(U3,      athcorr(U3,
      U1)      U1)      U2)      lnsigma
r1      .34      .9      .89      1.7      .5
. menl ... , ... initial(initvals)
```

Using a properly labeled initial-value matrix, we can also specify initial values for a subset of parameters. For example, we can specify initial values for fixed effects only as follows:

```
. matrix initvals = initvals[1,1..3]
. matrix list initvals
initvals[1,3]
      phi1:  phi2:  phi3:
      _cons  _cons  _cons
r1      19      55      8
. menl ... , ... initial(initvals)
```

Stored results

`menl` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(k_res)</code>	number of within-group error parameters
<code>e(k_eq)</code>	number of equations
<code>e(k_feq)</code>	number of fixed-effects equations
<code>e(k_req)</code>	number of random-effects equations
<code>e(k_reseq)</code>	number of within-group error equations
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	linearization log (restricted) likelihood
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>menl</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(ivars)</code>	grouping variables
<code>e(title)</code>	title in estimation output
<code>e(varlist)</code>	variables used in the specified equation
<code>e(eq_depvar)</code>	user-specified equation
<code>e(expressions)</code>	names of defined expressions, <i>expr_1</i> , <i>expr_2</i> , ..., <i>expr_k</i>
<code>e(ex_expr_i)</code>	defined expression <i>expr_i</i> , <i>i</i> =1,..., <i>k</i>
<code>e(hierarchy)</code>	random-effects hierarchy structure, (<i>path:covtype:REs</i>) (...)
<code>e(revars)</code>	names of random effects
<code>e(rstructlab)</code>	within-group error covariance output label
<code>e(timevar)</code>	within-group error covariance <code>t()</code> variable, if specified
<code>e(indexvar)</code>	within-group error covariance <code>index()</code> variable, if specified
<code>e(covbyvar)</code>	within-group error covariance <code>by()</code> variable, if specified
<code>e(stratavar)</code>	within-group error variance <code>strata()</code> variable, if specified
<code>e(corrbyvar)</code>	within-group error correlation <code>by()</code> variable, if specified
<code>e(rescovopt)</code>	within-group error covariance option, if <code>rescovariance()</code> specified
<code>e(resvaropt)</code>	within-group error variance option, if <code>resvariance()</code> specified
<code>e(rescorropt)</code>	within-group error correlation option, if <code>rescorrelation()</code> specified
<code>e(chi2type)</code>	Wald; type of model χ^2 test
<code>e(vce)</code>	conventional
<code>e(method)</code>	ML or REML
<code>e(opt)</code>	type of optimization, <code>lbates</code>
<code>e(crittype)</code>	optimization criterion, <code>linearization</code> log likelihood
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsok)</code>	predictions allowed by <code>margins</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginsdefault)</code>	default <code>predict()</code> specification for <code>margins</code>
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	factor-variable constraint matrix
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance
<code>e(b_sd)</code>	random-effects and within-group error estimates in the standard deviation metric
<code>e(V_sd)</code>	VCE for parameters in the standard deviation metric
<code>e(b_var)</code>	random-effects and within-group error estimates in the variance metric
<code>e(V_var)</code>	VCE for parameters in the variance metric
<code>e(cov_#)</code>	random-effects covariance structure at the hierarchical level $k-\#+1$ in a k -level model
<code>e(hierstats)</code>	group-size statistics for each hierarchy

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

Methods and formulas are presented under the following headings:

Introduction
Variance-components parameters
Inference based on linearization
Initial values

Introduction

Recall (1), a two-level NLME model, from the [Introduction](#),

$$y_{ij} = \mu(\mathbf{x}'_{ij}, \boldsymbol{\beta}, \mathbf{u}_j) + \epsilon_{ij} \quad i = 1, \dots, n_j; \quad j = 1, \dots, M$$

where M is the number of clusters and, for each cluster j , n_j is the number of observations in that cluster; $\mathbf{y}_j = (y_{1j}, y_{2j}, \dots, y_{n_j j})'$ is the $n_j \times 1$ response vector; $\mathbf{X}_j = (\mathbf{x}_{1j}, \mathbf{x}_{2j}, \dots, \mathbf{x}_{n_j j})'$ is the $n_j \times l$ matrix of covariates, including within-subject and between-subjects covariates; $\boldsymbol{\beta}$ is the $p \times 1$ vector of unknown parameters; $\mathbf{u}_j$ is the $q \times 1$ vector of random effects; and $\boldsymbol{\epsilon}_j = (\epsilon_{1j}, \epsilon_{2j}, \dots, \epsilon_{n_j j})'$ is the $n_j \times 1$ vector of within-group or within-cluster errors. $\mathbf{u}_j$'s follow a multivariate normal distribution with mean 0 and $q \times q$ variance–covariance matrix $\boldsymbol{\Sigma}$, and $\boldsymbol{\epsilon}_j$'s follow a multivariate normal distribution with mean 0 and $n_j \times n_j$ variance–covariance matrix $\sigma^2 \boldsymbol{\Lambda}_j$; $\mathbf{u}_j$'s are assumed to be independent of $\boldsymbol{\epsilon}_j$'s. Depending on the form of $\boldsymbol{\Lambda}_j$, σ^2 is either a within-group error variance σ_ϵ^2 or a squared scale parameter σ^2 . For example, when errors are i.i.d., that is, when $\boldsymbol{\Lambda}_j$ is the identity matrix, $\sigma^2 = \sigma_\epsilon^2$ is the within-group error variance. When $\boldsymbol{\Lambda}_j$ corresponds to the heteroskedastic power structure, σ^2 is a multiplier or a scale parameter.

Positive-definite matrices $\boldsymbol{\Sigma}/\sigma^2$ and $\boldsymbol{\Lambda}_j$ are expressed as functions of unconstrained parameter vectors $\boldsymbol{\alpha}_u$ and $\boldsymbol{\alpha}_w$, respectively, to recast a constrained optimization problem into an unconstrained one. Thus $\boldsymbol{\alpha}_u$ contains unconstrained random-effects covariance parameters and $\boldsymbol{\alpha}_w$ contains unconstrained within-group error covariance parameters. $\boldsymbol{\Lambda}_j$ may also depend on the random effects $\mathbf{u}_j$ and the fixed effects $\boldsymbol{\beta}$. For more details about $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}_j$ and about functional forms of parameter vectors $\boldsymbol{\alpha}_u$ and $\boldsymbol{\alpha}_w$ given different covariance structures, see [Variance-components parameters](#).

Based on (1), the marginal, with respect to $\mathbf{u}_j$'s, log likelihood for $(\boldsymbol{\beta}, \boldsymbol{\alpha}, \sigma^2)$ is

$$L(\boldsymbol{\beta}, \boldsymbol{\alpha}, \sigma^2) = \log \left\{ \prod_{j=1}^M \int f(\mathbf{y}_j | \mathbf{X}_j, \mathbf{u}_j; \boldsymbol{\beta}, \boldsymbol{\alpha}_w, \sigma^2) f(\mathbf{u}_j; \boldsymbol{\alpha}_u) d\mathbf{u}_j \right\} \quad (22)$$

where $\boldsymbol{\alpha} = (\boldsymbol{\alpha}'_u, \boldsymbol{\alpha}'_w)'$, $f(\mathbf{y}_j | \mathbf{X}_j, \mathbf{u}_j; \boldsymbol{\beta}, \boldsymbol{\alpha}_w, \sigma^2)$ is the conditional density of $\mathbf{y}_j$ given $\mathbf{X}_j$ and $\mathbf{u}_j$, and $f(\mathbf{u}_j; \boldsymbol{\alpha}_u)$ is the density of $\mathbf{u}_j$.

In general, there are no closed-form expressions for (22) or the marginal moments of an NLME model. This is because the random effects $\mathbf{u}_j$ enter the model nonlinearly, making the q -dimensional integral in (22) analytically intractable in all but simpler cases. Several estimation techniques have been proposed for estimating parameters $\boldsymbol{\beta}$, $\boldsymbol{\alpha}$, and σ^2 , including numerical integration of the integral in (22) by using an adaptive Gaussian quadrature and a linearization of the mean function in (1) by using a Taylor-series expansion.

menl implements the linearization method of [Lindstrom and Bates \(1990\)](#), with extensions from [Pinheiro and Bates \(1995\)](#), which is described in [Inference based on linearization](#).

Variance-components parameters

For numerical stability, maximization of (22) is performed with respect to the unique elements of the matrix $\mathbf{G} = \boldsymbol{\Sigma}/\sigma^2$ expressed as logarithms of standard deviations for the diagonal elements and hyperbolic arctangents of the correlations for off-diagonal elements. Let $\boldsymbol{\alpha}_u$ be the vector containing these elements. For example, if we assume that the elements of the random-effects vector $\mathbf{u}_j$ are independent, then $\boldsymbol{\Sigma}$ is diagonal and $\boldsymbol{\alpha}_u$ will contain q distinct parameters— q logarithms of standard deviations. [Table 1](#) lists the vectors of parameters $\boldsymbol{\alpha}_u$ for all random-effects covariance structures supported by menl in the `covariance(vartype)` option.

Table 1. Variance-components parameters

<i>vartype</i>	α'_u
independent	$(g_1, g_2, \dots, g_q)$
exchangeable	(g_1, g_{12})
identity	g_1
unstructured	$(g_1, g_2, \dots, g_q, g_{12}, g_{13}, \dots, g_{q-1q})$
Notes: $g_u = \log(\sqrt{[\mathbf{G}]_{uu}})$, $g_{uv} = \text{atanh}([\mathbf{G}]_{uv})$. unstructured has $q(q+1)/2$ parameters.	

The within-group error covariance matrix is parameterized as follows,

$$\text{Var}(\epsilon_j | \mathbf{u}_j) = \sigma^2 \mathbf{\Lambda}_j(\mathbf{X}_j, \beta, \mathbf{u}_j, \alpha_w) = \sigma^2 \mathbf{S}_j(\delta, \mathbf{v}_j) \mathbf{C}_j(\rho) \mathbf{S}_j(\delta, \mathbf{v}_j)$$

where $\alpha_w = (\delta^{*'}, \rho^{*'})'$ and δ^* and ρ^* are unconstrained versions of δ and ρ defined in [table 2](#) and [table 3](#), respectively. For example, for a positive δ_1 , $\delta_1^* = \log(\delta_1)$. $\mathbf{S}_j = \mathbf{S}_j(\delta, \mathbf{v}_j)$ is an $n_j \times n_j$ diagonal matrix with nonnegative diagonal elements $g(\delta, v_{1j}), g(\delta, v_{2j}), \dots, g(\delta, v_{n_j j})$ such that $\text{Var}(\epsilon_{ij}) = \sigma^2 [\mathbf{S}_j]_{ii}^2 = \sigma^2 g^2(\delta, v_{ij})$, where v_{ij} 's are the values of a variance covariate or the values of a mean function $\mu(\mathbf{x}'_{ij}, \beta, \mathbf{u}_j)$, in which case $\mathbf{\Lambda}_j$ will depend on $\mathbf{X}_j$, β , and $\mathbf{u}_j$. $\mathbf{C}_j = \mathbf{C}_j(\rho)$ is a correlation matrix such that $\text{corr}(\epsilon_{ij}, \epsilon_{kj}) = [\mathbf{C}_j]_{ik} = h(|t_{ij} - t_{kj}|, \rho)$, where t_{ij} is a value of a time variable for time-dependent correlation structures such as AR and MA structures. A list of the supported $g(\cdot)$ and $h(\cdot)$ functions is given in [table 2](#) and [table 3](#), respectively.

[Carroll and Ruppert \(1988\)](#) introduced various variance functions $g(\delta, v_{ij})$ to model heteroskedasticity, which were further studied in the context of NLME models by [Davidian and Giltinan \(1995\)](#). [Table 2](#) lists variance functions supported by the `resvariance(resvarfunc ...)` option.

Table 2. Supported variance functions $g(\cdot)$

<i>resvarfunc</i>	$g(\delta, v_{ij})$	δ'
identity	1	—
linear	$\sqrt{v_{ij}}$	—
power	$c + v_{ij} ^\delta$	$(c, \delta), c \geq 0$
power, noconstant	$ v_{ij} ^\delta$	δ
exponential	$\exp(\delta v_{ij})$	δ

The variance function $g(\cdot)$ and thus the within-group error covariance may depend on β and $\mathbf{u}_j$ through $\mu(\cdot)$, when $v_{ij} = \mu_{ij} = \mu(\mathbf{x}'_{ij}, \mathbf{u}_j, \beta)$ in [table 2](#). This is particularly appealing in PK applications, where there is often considerable intraindividual heterogeneity that is modeled, for example, as a power function of the mean.

The within-group error correlation structure is governed by the $h(\cdot)$ function. [Table 3](#) lists correlation structures that are supported by the `rescorrelation(rescorr ...)` option and also have a closed-form expression. In addition, the AR and MA correlation structures are defined below.

The **ar** p structure assumes that the errors have an AR structure of order p . That is,

$$\epsilon_{ij} = \phi_1 \epsilon_{i-1,j} + \cdots + \phi_p \epsilon_{i-p,j} + z_{ij}$$

where z_{ij} are i.i.d. Gaussian with mean 0 and variance σ_z^2 . **menl** reports estimates of $\phi_1, \dots, \phi_p$ and the overall error variance σ_ϵ^2 , which can be derived from the above expression. This structure has a closed-form expression only for $p = 1$, in which case $\phi_1 = \rho$ is the correlation between error terms.

The **ma** q structure assumes that the errors are an MA process of order q . That is,

$$\epsilon_{ij} = Z_i + \theta_1 Z_{i-1} + \cdots + \theta_q Z_{i-q}$$

where Z_l are i.i.d. Gaussian with mean 0 and variance σ_Z^2 . **menl** reports estimates of $\theta_1, \dots, \theta_q$ and the overall error variance σ_ϵ^2 , which can be derived from the above expression.

Table 3. Within-group error correlation functions $h(\cdot)$

<i>rescorr</i>	$h(t_{ij} - t_{kj} , \rho)$	Expression	ρ
identity	$h(k)$	$I(k = 0)$	—
ar 1	$h(k, \rho)$	$\rho^k, k = 0, 1, \dots$	$\rho, \rho < 1$
ar $p, p > 1$	$h(k, \phi)$	no closed form	$(\phi_1, \phi_2, \dots, \phi_p)$
ctar1	$h(s, \rho)$	$\rho^s, s \geq 0$	$\rho, \rho < 1$
ma q	$h(k, \theta)$	$\begin{cases} \frac{\sum_{j=0}^{q- k } \theta_j \theta_{j+ k }}{\sum_{j=0}^q \theta_j^2} & k \leq q \\ 0 & k > q \end{cases}$	$(\theta_0 = 1, \theta_1, \dots, \theta_q)$

You can build many flexible within-group error covariance structures by combining different functions $g(\cdot)$ and $h(\cdot)$, that is, by combining the **resvariance()** and **rescorrelation()** options. For example, you can combine an AR(1) correlation structure with a heteroskedastic structure that is expressed as a power function of the mean by specifying **rescorrelation(ar 1, t(timevar))** and **resvariance(power _yhat)**.

Inference based on linearization

Let's write (1), equivalently, in matrix form as

$$\mathbf{y}_j = \boldsymbol{\mu}(\mathbf{X}_j, \boldsymbol{\beta}, \mathbf{u}_j) + \boldsymbol{\Lambda}_j^{\frac{1}{2}}(\mathbf{X}_j, \boldsymbol{\beta}, \mathbf{u}_j, \boldsymbol{\alpha}_w) \mathbf{e}_j$$

Here $\boldsymbol{\mu}(\mathbf{X}_j, \boldsymbol{\beta}, \mathbf{u}_j)$ depends on $\boldsymbol{\beta}$ and $\mathbf{u}_j$ through the function $\mathbf{d}(\cdot)$ in (2), and $\mathbf{e}_j$'s $\sim N(\mathbf{0}, \sigma^2 I_{n_j})$, where I_{n_j} is the identity matrix of dimension n_j . In what follows, for brevity, we suppress the dependence of $\boldsymbol{\mu}$ and $\boldsymbol{\Lambda}_j$ on $\mathbf{X}_j$.

Following [Lindstrom and Bates \(1990\)](#), we will initially assume that $\boldsymbol{\Lambda}_j$ does not depend on $\mathbf{X}_j$, $\boldsymbol{\beta}$, and $\mathbf{u}_j$ or, equivalently, on ϕ_j but rather on j only through its dimension; that is, $\boldsymbol{\Lambda}_j = \boldsymbol{\Lambda}_j(\boldsymbol{\alpha}_w)$. Therefore, heteroskedastic structures that depend on the mean are not yet allowed in this context. Toward the end of this section, we will present a modified version of the algorithm that accounts for the dependence of $\boldsymbol{\Lambda}_j$ on ϕ_j .

Lindstrom and Bates discuss a natural extension of the methods for the LME models to NLME models. For a known α (and thus known Σ and Λ_j) and σ^2 , the estimates of β and $\mathbf{u}_j$ jointly minimize

$$\sum_{j=1}^M \left[\log |\Sigma(\alpha_u)| + \mathbf{u}_j' \{\Sigma(\alpha_u)\}^{-1} \mathbf{u}_j + \log \left| \sigma^2 \Lambda_j(\alpha_w) \right| \right. \\ \left. + \sigma^{-2} \{\mathbf{y}_j - \mu(\beta, \mathbf{u}_j)\}' \Lambda_j^{-1}(\alpha_w) \{\mathbf{y}_j - \mu(\beta, \mathbf{u}_j)\} \right]$$

which is twice the negative log likelihood for β when $\mathbf{u}_j$ is fixed or twice the negative log of the posterior density of $\mathbf{u}_j$ when β is fixed. Consequently, one strategy for estimating β and (predicting) $\mathbf{u}_j$ is to minimize the above objective function with respect to β and $\mathbf{u}_j$ given suitable estimates of α and σ^2 . Estimation of α and σ^2 can be accomplished by using MLE with respect to the marginal density of $\mathbf{y}_j$, in which $\mathbf{u}_j$'s are integrated out. But because no closed-form expression for this density is available, we approximate the conditional distribution of $\mathbf{y}_j$ given $\mathbf{u}_j$ by a multivariate normal distribution with an expectation that is linear in $\mathbf{u}_j$ and β . This is illustrated in step 2 of the algorithm below.

Lindstrom and Bates (1990) propose the following two-step estimation method or alternating algorithm.

Step 1 (PNLS step). Given current estimates $\hat{\alpha}$ (and thus $\hat{\alpha}_u$ and $\hat{\alpha}_w$) of α and $\hat{\sigma}^2$ of σ^2 , minimize with respect to β and $\mathbf{u}_j$

$$\sum_{j=1}^M \left[\log |\Sigma(\hat{\alpha}_u)| + \mathbf{u}_j' \{\Sigma(\hat{\alpha}_u)\}^{-1} \mathbf{u}_j + \log \left| \hat{\sigma}^2 \Lambda_j(\hat{\alpha}_w) \right| \right. \\ \left. + \hat{\sigma}^{-2} \{\mathbf{y}_j - \mu(\beta, \mathbf{u}_j)\}' \Lambda_j^{-1}(\hat{\alpha}_w) \{\mathbf{y}_j - \mu(\beta, \mathbf{u}_j)\} \right] \quad (23)$$

Define Δ such that $\sigma^2 \Sigma^{-1} = \Delta' \Delta$. Note that $\Delta = \Delta(\alpha_u)$, but for notational convenience, this dependency is suppressed throughout the rest of this section. Equation (23) is equivalent to minimizing the penalized least-square objective function

$$\text{PNLS step:} \quad \sum_{j=1}^M \left[\left\| \{\Lambda_j'(\alpha_w)\}^{-1/2} \{\mathbf{y}_j - \mu(\beta, \mathbf{u}_j)\} \right\|^2 + \|\Delta \mathbf{u}_j\|^2 \right]$$

with respect to β and $\mathbf{u}_j$ while holding the current estimates of α (and, consequently, of Δ and of Λ_j) fixed. `pnlsopts(iterate(#))` iterations are performed at this step, unless the convergence criterion (CC) is met. The CC for PNLS optimization is controlled by `pnlsopts(nrtolerance(#))` and one of `pnlsopts(ltolerance(#))` or `pnlsopts(tolerance(#))`; see `menlmaxopts` for details.

Denote the resulting estimates as $\hat{\mathbf{u}}_j$ and $\hat{\beta}$.

Step 2 (LME step). Perform a first-order Taylor-series expansion of the model mean function around the current estimates of β and of the conditional modes of the random effects $\mathbf{u}_j$, yielding

$$\mathbf{y}_j = \mu(\hat{\beta}, \hat{\mathbf{u}}_j) + \hat{\mathbf{X}}_j (\beta - \hat{\beta}) + \hat{\mathbf{Z}}_j (\mathbf{u}_j - \hat{\mathbf{u}}_j) + \Lambda_j^{\frac{1}{2}}(\alpha_w) \mathbf{e}_j \quad (24)$$

where

$$\begin{aligned}\widehat{\mathbf{X}}_j &= \left. \frac{\partial \boldsymbol{\mu}(\boldsymbol{\beta}, \mathbf{u}_j)}{\partial \boldsymbol{\beta}'} \right|_{\boldsymbol{\beta}=\widehat{\boldsymbol{\beta}}, \mathbf{u}_j=\widehat{\mathbf{u}}_j} \\ \widehat{\mathbf{Z}}_j &= \left. \frac{\partial \boldsymbol{\mu}(\boldsymbol{\beta}, \mathbf{u}_j)}{\partial \mathbf{u}_j'} \right|_{\boldsymbol{\beta}=\widehat{\boldsymbol{\beta}}, \mathbf{u}_j=\widehat{\mathbf{u}}_j}\end{aligned}$$

Model (24) is essentially an LME model, and we use notations $\widehat{\mathbf{X}}_j$ and $\widehat{\mathbf{Z}}_j$ for the derivatives to emphasize this. That is, $\widehat{\mathbf{X}}_j$ and $\widehat{\mathbf{Z}}_j$ represent the corresponding fixed-effects and random-effects design matrices of an LME model.

Thus the approximate conditional distribution of $\mathbf{y}_j$ is

$$\mathbf{y}_j | \mathbf{u}_j \sim N \left\{ \boldsymbol{\mu}(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{u}}_j) + \widehat{\mathbf{X}}_j (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}) + \widehat{\mathbf{Z}}_j (\mathbf{u}_j - \widehat{\mathbf{u}}_j), \sigma^2 \boldsymbol{\Lambda}_j \right\}$$

Because the expectation is now linear in random effects $\mathbf{u}_j$, the approximate conditional distribution of $\mathbf{y}_j$, along with distribution of $\mathbf{u}_j$, allows us to approximate the marginal distribution of $\mathbf{y}_j$ as

$$\mathbf{y}_j \sim N \left\{ \boldsymbol{\mu}(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{u}}_j) + \widehat{\mathbf{X}}_j (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}) - \widehat{\mathbf{Z}}_j \widehat{\mathbf{u}}_j, \sigma^2 \mathbf{V}_j(\boldsymbol{\alpha}) \right\} \quad (25)$$

where $\mathbf{V}_j(\boldsymbol{\alpha}) = \widehat{\mathbf{Z}}_j \boldsymbol{\Delta}^{-1} (\boldsymbol{\Delta}^{-1})' \widehat{\mathbf{Z}}_j' + \boldsymbol{\Lambda}_j(\boldsymbol{\alpha}_w)$.

Let $\widehat{\mathbf{w}}_j = \mathbf{y}_j - \boldsymbol{\mu}(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{u}}_j) + \widehat{\mathbf{X}}_j \widehat{\boldsymbol{\beta}} + \widehat{\mathbf{Z}}_j \widehat{\mathbf{u}}_j$. Estimation of $\boldsymbol{\alpha}$ and σ^2 can now be accomplished by maximizing the log likelihood corresponding to the approximate marginal distribution in (25),

$$\begin{aligned} \text{LME step:} \quad l_{\text{LB}}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \sigma^2) &= -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2} \sum_{j=1}^M \left\{ \log |\mathbf{V}_j(\boldsymbol{\alpha})| \right. \\ &\quad \left. + \sigma^{-2} (\widehat{\mathbf{w}}_j - \widehat{\mathbf{X}}_j \boldsymbol{\beta})' \mathbf{V}_j^{-1}(\boldsymbol{\alpha}) (\widehat{\mathbf{w}}_j - \widehat{\mathbf{X}}_j \boldsymbol{\beta}) \right\} \end{aligned} \quad (26)$$

where $n = \sum_{j=1}^M n_j$.

Alternatively, when the `rem1` option is specified, we take an REML approach and maximize

$$l_{\text{LB},R}(\boldsymbol{\alpha}, \sigma^2) = l_{\text{LB}}(\boldsymbol{\alpha}, \widehat{\boldsymbol{\beta}}(\boldsymbol{\alpha}), \sigma^2) - \frac{1}{2} \sum_{j=1}^M \log \left| \sigma^{-2} \widehat{\mathbf{X}}_j' \mathbf{V}_j^{-1}(\boldsymbol{\alpha}) \widehat{\mathbf{X}}_j \right| \quad (27)$$

The LME step (step 2) of the alternating algorithm consists of optimizing an LME log likelihood, in which the response vector is given by $\widehat{\mathbf{w}}_j$ and the fixed- and random-effects design matrices are given by $\widehat{\mathbf{X}}_j$ and $\widehat{\mathbf{Z}}_j$, respectively. `lmeopts(iterate(#))` iterations are performed at this step, unless the CC is met. The CC for LME optimization is controlled by `lmeopts(nrtolerance(#))` and one of `lmeopts(ltolerance(#))` or `lmeopts(tolerance(#))`; see [menlmaxopts](#) for details.

The LME step produces estimates $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\sigma}^2$. (The estimates $\widehat{\boldsymbol{\beta}}$ can also be obtained at this step, but it is generally more computationally efficient to compute them at the PNLS step.) These estimates will now be used in step 1, the PNLS step.

Stopping rules. One PNLS step and one LME step correspond to one iteration of the alternating algorithm. The log likelihood reported by `menl` at each iteration is the log likelihood (26) or, if the `rem1` option is specified, (27) from the last iteration of the LME step. `menl` refers to this log likelihood as “linearization log likelihood” because it corresponds to the log likelihood of the LME model, which was the result of the linearization of the NLME model. The algorithm stops when the linearization likelihoods from successive iterations satisfy `ltolerance(#)`, when the parameter estimates from successive iterations satisfy `tolerance(#)`, or if the model does not converge, when the maximum number of iterations in `iterate()` is reached; see `menlmaxopts` for details about maximization options. Because the alternating algorithm does not provide a joint Hessian matrix for all parameters, there is no check for the tolerance of the scaled gradient; thus the convergence cannot be established in its strict sense. The convergence is declared based on the stopping rules described above.

When $\Lambda_j = \Lambda_j(\beta, \mathbf{u}_j, \alpha_w)$ depends on $\mathbf{u}_j$ and β , which is the case, for example, with `resvariance(power _yhat)` and `resvariance(exponential _yhat)`, an intermediate step between the PNLS and the LME step is performed to replace the fixed effects and random effects in Λ_j , or more precisely in the variance function $g(\cdot)$, by their current estimates from the PNLS step. After that, $\Lambda_j(\alpha_w; \hat{\beta}, \hat{\mathbf{u}}_j) = \Lambda_j(\alpha_w)$ depends only on α_w because both $\mathbf{u}_j$ and β are held fixed at their current estimates throughout the LME step.

Efficient methods for computing (26) or (27) are given in chapters 2 and 5 of [Pinheiro and Bates \(2000\)](#). Namely, to simplify the optimization problem, one can express the optimal values of β and σ^2 as functions of α (and thus of Δ and α_w) and work with the profiled log likelihood of α .

For the PNLS step, the objective function to be minimized is the penalized sum of squares

$$\sum_{j=1}^M \left[\|(\Lambda'_j)^{-1/2} \{\mathbf{y}_j - \boldsymbol{\mu}(\beta, \mathbf{u}_j)\}\|^2 + \|\Delta \mathbf{u}_j\|^2 \right]$$

By adding “pseudo”-observations to the data, the PNLS problem can be reexpressed as a standard nonlinear least-squares problem. Thus step 1 of the alternating algorithm is sometimes called the “pseudodata step”. Define pseudoobservations $\tilde{\mathbf{y}}_j$ as follows:

$$\tilde{\mathbf{y}}_j = \begin{bmatrix} (\Lambda'_j)^{-1/2} \mathbf{y}_j \\ \mathbf{0} \end{bmatrix} \quad \tilde{\boldsymbol{\mu}}(\beta, \mathbf{u}_j) = \begin{bmatrix} (\Lambda'_j)^{-1/2} \boldsymbol{\mu}(\beta, \mathbf{u}_j) \\ \Delta \mathbf{u}_j \end{bmatrix}$$

Then, the PNLS step can be rewritten as

$$\sum_{j=1}^M \|\tilde{\mathbf{y}}_j - \tilde{\boldsymbol{\mu}}(\beta, \mathbf{u}_j)\|^2$$

Hence, for values of α and σ^2 fixed at the current estimates, the estimation of β and $\mathbf{u}_j$ in the PNLS step can be regarded as a standard nonlinear least-squares problem. A popular iterative estimation technique for standard nonlinear least-squares is the Gauss–Newton method (see [Pinheiro and Bates \[2000, chap. 7\]](#) for more details).

After the completion of the alternating algorithm, an extra LME iteration is performed, with fixed effects profiled-out of the likelihood, to reparameterize $[\alpha, \log(\sigma)]$ to their natural metric and to compute their standard errors with the delta method. This step is labeled **Computing standard errors**: in the output of `menl`. If you are interested only in standard errors for fixed effects, you can skip this step by specifying the `nstderr` option, in which case standard errors for the random-effects

and within-group error covariance parameters will not be computed and will be shown as missing in the output table. The standard errors for the fixed effects are obtained from the PNLS step, and the standard errors for random-effects parameters are obtained from the LME step.

Inference on the parameters of the NLME model is based on the approximating LME model with log likelihood and restricted log likelihood functions defined in (26) and (27). Therefore, all the inferential machinery available within the context of LME models can be used. For example, under the LME approximation, the distribution of the (restricted) MLE $\hat{\beta}$ of the fixed effects is

$$\hat{\beta} \sim N \left\{ \beta, \sigma^2 \left(\sum_{j=1}^M \hat{\mathbf{X}}_j' \mathbf{V}_j^{-1}(\alpha) \hat{\mathbf{X}}_j \right)^{-1} \right\}$$

and for random-effects and within-group error parameters is

$$\begin{bmatrix} \hat{\alpha} \\ \log \hat{\sigma} \end{bmatrix} \sim N \left\{ \begin{bmatrix} \alpha \\ \log \sigma \end{bmatrix}, I^{-1}(\alpha, \sigma) \right\}$$

where

$$I(\alpha, \sigma) = - \begin{bmatrix} \partial^2 l_{\text{LB}_p} / \partial \alpha \partial \alpha' & \partial^2 l_{\text{LB}_p} / \partial \log \sigma \partial \alpha' \\ \partial^2 l_{\text{LB}_p} / \partial \alpha \partial \log \sigma & \partial^2 l_{\text{LB}_p} / \partial^2 \log \sigma \end{bmatrix}$$

and $l_{\text{LB}_p} = l_{\text{LB}_p}(\alpha, \sigma)$ is the approximated log likelihood from the LME step with fixed effects profiled out. Because inference is based on the LME approximation of the original NLME model, asymptotic results are technically “approximately asymptotic” and are thus less accurate than the asymptotic inferential results for LME models as described in [ME] **mixed**.

Initial values

The PNLS step requires starting values for β and $\mathbf{u}_j$. These are obtained from the EM algorithm; see, for example, Bates and Pinheiro (1998) for details. You can control optimization within the EM algorithm by specifying the `emtolerance()` and `emiterate()` options. You can also supply your own initial values; see [Examples of specifying initial values](#). NLME models are often sensitive to initial values, so it is good practice to try different sets of initial values to verify that your results are robust to them.

References

- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies.
<http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Boeckmann, A. J., L. B. Sheiner, and S. L. Beal. [1994] 2011. *NONMEM Users Guide, Part V: Introductory Guide*. San Francisco: Regents of the University of California. <https://nonmem.iconplc.com/nonmem720/guides/v.pdf>.
- Carroll, R. J., and D. Ruppert. 1988. *Transformation and Weighting in Regression*. New York: Chapman & Hall.
- Davidian, M., and D. M. Giltinan. 1995. *Nonlinear Models for Repeated Measurement Data*. Boca Raton, FL: Chapman & Hall/CRC.
- . 2003. Nonlinear models for repeated measurement data: An overview and update. *Journal of Agricultural, Biological, and Environmental Statistics* 8: 387–419.
- Demidenko, E. 2013. *Mixed Models: Theory and Applications with R*. 2nd ed. Hoboken, NJ: Wiley.

- Draper, N., and H. Smith. 1998. *Applied Regression Analysis*. 3rd ed. New York: Wiley.
- Fitzmaurice, G. M., M. Davidian, G. Verbeke, and G. Molenberghs, ed. 2009. *Longitudinal Data Analysis*. Boca Raton, FL: Chapman & Hall/CRC.
- Gibaldi, M., and D. Perrier. 1982. *Pharmacokinetics*. 2nd ed, revised and expanded. New York: Dekker.
- Hand, D. J., and M. J. Crowder. 1996. *Practical Longitudinal Data Analysis*. Boca Raton, FL: Chapman & Hall.
- Harring, J. R., and J. Liu. 2016. A comparison of estimation methods for nonlinear mixed-effects models under model misspecification and data sparseness: A simulation study. *Journal of Modern Applied Statistical Methods* 15(1): Article 27.
- Joyner, W. B., and D. M. Boore. 1981. Peak horizontal acceleration and velocity from strong-motion records including records from the 1979 imperial valley, California, earthquake. *Bulletin of the Seismological Society of America* 71: 2011–2038.
- Lindstrom, M. J., and D. M. Bates. 1990. Nonlinear mixed effects models for repeated measures data. *Biometrics* 46: 673–687.
- Pierson, R. A., and O. J. Ginther. 1987. Follicular population dynamics during the estrous cycle of the mare. *Animal Reproduction Science* 14: 219–231.
- Pinheiro, J. C., and D. M. Bates. 1995. Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics* 4: 12–35.
- . 2000. *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- Plan, E. L., A. Maloney, F. Mentré, M. O. Karlsson, and J. Bertrand. 2012. Performance comparison of various maximum likelihood nonlinear mixed-effects estimation methods for dose–response models. *AAPS Journal* 14: 420–432.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.
- Thompson, W. A., Jr. 1962. The problem of negative estimates of variance components. *Annals of Mathematical Statistics* 33: 273–289.
- Vonesh, E., and V. M. Chinchilli. 1997. *Linear and Nonlinear Models for the Analysis of Repeated Measurements*. New York: Dekker.
- Wolfinger, R. D., and X. Lin. 1997. Two Taylor-series approximation methods for nonlinear mixed models. *Computational Statistics & Data Analysis* 25: 465–490.

Also see

- [ME] **menl postestimation** — Postestimation tools for menl
- [ME] **meglm** — Multilevel mixed-effects generalized linear model
- [ME] **mixed** — Multilevel mixed-effects linear regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [R] **nl** — Nonlinear least-squares estimation
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples
Also see

[predict](#)
Methods and formulas

[margins](#)
Reference

Postestimation commands

The following postestimation commands are of special interest after `menl`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat recovariance</code>	display the estimated random-effects covariance matrices
<code>estat sd</code>	display variance components as standard deviations and correlations
<code>estat wcorrelation</code>	display model-implied within-cluster correlations and standard deviations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estimates</code>	cataloging estimation results
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
<code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

predict

Description for predict

`predict` creates a new variable containing predictions of mean values, residuals, or standardized residuals. It can also create multiple new variables containing estimates of random effects and their standard errors or containing predicted [named substitutable expressions](#).

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [reses(newvarsspec)
relevel(levelvar) options]
```

Syntax for predicting named substitutable expressions (parameters)

Predict specific parameters

```
predict [type] (newvar = {param:}) [(newvar = {param:})] [...] [if] [in]
[, fixedonly relevel(levelvar) options]
```

```
predict [type] newvarsspec [if] [in], parameters(paramnames) [ fixedonly
relevel(levelvar) options]
```

Predict all parameters

```
predict [type] newvarsspec [if] [in], parameters [fixedonly relevel(levelvar)
options]
```

Syntax for obtaining other predictions

```
predict [type] newvar [if] [in] [, statistic fixedonly relevel(levelvar) options]
```

newvarsspec is *stub** or *newvarlist*.

paramnames is *param* [*param* [...]] and *param* is a name of a substitutable expression as specified in one of `menl`'s `define()` options.

<i>statistic</i>	Description
Main	
<code>yhat</code>	prediction for the expected response conditional on the random effects
<code>mu</code>	synonym for <code>yhat</code>
<code>residuals</code>	residuals, response minus predicted values
* <code>rstandard</code>	standardized residuals

Unstarred statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample. Starred statistics are calculated only for the estimation sample, even when `if e(sample)` is not specified.

<i>options</i>	Description
Main	
<code>iterate(#)</code>	maximum number of iterations when computing random effects; default is <code>iterate(10)</code>
<code>tolerance(#)</code>	convergence tolerance when computing random effects; default is <code>tolerance(1e-6)</code>

Options for predict

Main

`yhat` calculates the predicted values, which are the mean-response values conditional on the random effects, $\mu(\mathbf{x}'_{ij}, \hat{\beta}, \hat{\mathbf{u}}_j)$. By default, the predicted values account for random effects from all levels in the model; however, if the `relevel(levelvar)` option is specified, then the predicted values are fit beginning with the topmost level down to and including level `levelvar`. For example, if `classes` are nested within `schools`, then typing

```
. predict yhat_school, yhat relevel(school)
```

would produce school-level predictions. That is, the predictions would incorporate school-specific random effects but not those for each class nested within each school. If the `fixedonly` option is specified, predicted values conditional on zero random effects, $\mu(\mathbf{x}'_{ij}, \hat{\beta}, \mathbf{0})$, are calculated based on the estimated fixed effects (coefficients) in the model when the random effects are fixed at their theoretical mean value of `0`.

`mu` is a synonym for `yhat`.

`parameters` and `parameters(paramnames)` calculate predictions for all or a subset of the [named substitutable expressions](#) in the model. By default, the predictions account for random effects from all levels in the model; however, if the `relevel(levelvar)` option is specified, then the predictions would incorporate random effects from the topmost level down to and including level `levelvar`. Option `parameters(param)` is useful with `margins`. `parameters()` does not appear in the dialog box.

`reffects` calculates predictions of the random effects. For the Lindstrom–Bates estimation method of [menl](#), these are essentially the best linear unbiased predictions (BLUPs) of the random effects in the LME approximated log likelihood; see [Inference based on linearization](#) in [ME] [menl](#). By default, estimates of all random effects in the model are calculated. However, if the `relevel(levelvar)` option is specified, then estimates of random effects for only level `levelvar` in the model are calculated. For example, if `classes` are nested within `schools`, then typing

```
. predict b*, reffects relevel(school)
```

would produce estimates at the school level. You must specify q new variables, where q is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you.

reses(*newvarsspec*) calculates the standard errors of the estimates of the random effects. By default, standard errors for all random effects in the model are calculated. However, if the **relevel**(*levelvar*) option is specified, then standard errors of the estimates of the random effects for only level *levelvar* in the model are calculated; see the **reffects** option.

You must specify q new variables, where q is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you. The new variables will have the same storage type as the corresponding random-effects variables.

The **reffects** and **reses**() options often generate multiple new variables at once. When this occurs, the random effects (or standard errors) contained in the generated variables correspond to the order in which the variance components are listed in the output of **men1**. Still, examining the variable labels of the generated variables (with the **describe** command, for instance) can be useful in deciphering which variables correspond to which terms in the model.

residuals calculates residuals, equal to the responses minus the predicted values **yhat**. By default, the predicted values account for random effects from all levels in the model; however, if the **relevel**(*levelvar*) option is specified, then the predicted values are fit beginning at the topmost level down to and including level *levelvar*.

rstandard calculates standardized residuals, equal to the residuals multiplied by the inverse square root of the estimated error covariance matrix.

fixedonly specifies that all random effects be set to zero, equivalent to using only the fixed portion of the model.

relevel(*levelvar*) specifies the level in the model at which predictions involving random effects are to be obtained; see the options above for the specifics. *levelvar* is the name of the model level; it is the name of the variable describing the grouping at that level.

iterate(#) specifies the maximum number of iterations when computing estimates of the random effects. The default is **iterate**(10). This option is relevant only to predictions that depend on random effects. This option is not allowed if the **fixedonly** option is specified.

tolerance(#) specifies a convergence tolerance when computing estimates of the random effects. The default is **tolerance**(1e-6). This option is relevant only to predictions that depend on random effects. This option is not allowed if the **fixedonly** option is specified.

margins

Description for margins

`margins` estimates margins of response for predicted mean values or [named substitutable expressions](#).

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [options]
```

statistic	Description
yhat	predicted values conditional on zero random effects; implies <code>fixedonly</code> ; the default
<u>parameters</u> (<i>param</i>)	predicted named substitutable expression <i>param</i> conditional on zero random effects; implies <code>fixedonly</code>
<u>re</u> ffects	not allowed with <code>margins</code>
<u>re</u> siduals	not allowed with <code>margins</code>
<u>r</u> standard	not allowed with <code>margins</code>

The `fixedonly` option is assumed for the predictions used with `margins`.

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.

For the full syntax, see [\[R\] margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting an NLME model using `menl`. For the most part, calculation centers on obtaining estimates of the random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation. The estimates of the random effects are in turn used to obtain predicted values and residuals at different nesting levels. These are useful for checking model assumptions and may be used in general as model-building tools.

► Example 1: Testing variance components

In [example 9](#) and [example 12](#) of [\[ME\] menl](#), we modeled the average leaf weight of two genotypes of soybean plants over three growing seasons as

$$\text{weight}_{ij} = \frac{\phi_{1j}}{1 + \exp\{-(\text{time}_{ij} - \phi_{2j})/\phi_{3j}\}} + \epsilon_{ij}$$

for $j = 1, \dots, 48$ and $i = 1, \dots, n_j$, with $8 \leq n_j \leq 10$. Here we consider a simplified version of the stage 2 model specification from [example 12](#),

$$\phi_j = \begin{bmatrix} \phi_{1j} \\ \phi_{2j} \\ \phi_{3j} \end{bmatrix} = \begin{bmatrix} \beta_{11} + \beta_{12}S_{89,j} + \beta_{13}S_{90,j} + u_{1j} \\ \beta_{21} + \beta_{22}S_{89,j} + \beta_{23}S_{90,j} + \beta_{24}P_j \\ \beta_{31} + \beta_{32}S_{89,j} + \beta_{33}S_{90,j} \end{bmatrix}$$

where $P_j = I(\text{variety}_j = \text{P})$, $S_{89,j} = I(\text{year}_j = 1989)$, and $S_{90,j} = I(\text{year}_j = 1990)$. The random effects u_{1j} 's are normally distributed with mean 0 and variance σ_{u1}^2 and errors ϵ_{ij} 's are normally distributed with mean 0 and error variance

$$\text{Var}(\epsilon_{ij}) = \sigma^2(\widehat{\text{weight}_{ij}})^{2\delta}$$

Let's fit this model using menl.

```
. use http://www.stata-press.com/data/r15/soybean
(Growth of soybean plants (Davidian and Giltinan, 1995))

. menl weight = {phi1:}/(1+exp(-(time-{phi2:})/{phi3:})),
> define(phi1: i.year U1[plot])
> define(phi2: i.year i.variety)
> define(phi3: i.year, xb) resvariance(power _yhat, noconstant)
```

Obtaining starting values by EM:

Alternating PNLS/LME algorithm:

```
Iteration 1: linearization log likelihood = -324.215786
Iteration 2: linearization log likelihood = -313.897333
Iteration 3: linearization log likelihood = -314.762871
Iteration 4: linearization log likelihood = -314.431695
Iteration 5: linearization log likelihood = -314.513097
Iteration 6: linearization log likelihood = -314.493994
Iteration 7: linearization log likelihood = -314.499224
Iteration 8: linearization log likelihood = -314.498385
Iteration 9: linearization log likelihood = -314.498535
Iteration 10: linearization log likelihood = -314.498506
```

Computing standard errors:

```
Mixed-effects ML nonlinear regression      Number of obs      =      412
Group variable: plot                      Number of groups   =      48
                                           Obs per group:
                                           min =      8
                                           avg =      8.6
                                           max =     10
                                           Wald chi2(7)       =     193.98
                                           Prob > chi2        =     0.0000

Linearization log likelihood = -314.49851
      phi1: i.year U1[plot]
      phi2: i.year i.variety
      phi3: i.year
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
phi1						
year						
1989	-6.614797	1.195633	-5.53	0.000	-8.958195	-4.271399
1990	-3.749016	1.264716	-2.96	0.003	-6.227814	-1.270218
_cons	20.28009	.953959	21.26	0.000	18.41036	22.14981
phi2						
year						
1989	-2.623514	.9752055	-2.69	0.007	-4.534882	-.7121467
1990	-5.142726	.9879783	-5.21	0.000	-7.079128	-3.206324
variety						
P	-2.265482	.3274228	-6.92	0.000	-2.907219	-1.623745
_cons	55.25935	.7554917	73.14	0.000	53.77861	56.74008
phi3						
year						
1989	-.9538782	.1963606	-4.86	0.000	-1.338738	-.5690186
1990	-.7220007	.2081227	-3.47	0.001	-1.129914	-.3140877
_cons	8.042677	.1452638	55.37	0.000	7.757965	8.327389

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
plot: Identity	var(U1)	4.260962	1.114483	2.551939	7.11451
Residual variance:					
Power _yhat					
	sigma2	.046573	.0038271	.0396449	.0547118
	delta	.9667451	.0229472	.9217694	1.011721

`menl` does not report tests against zeros for parameters in the random-effects table because they are not appropriate for all types of parameters such as variances. For some parameters such as power parameter δ in our example, labeled as `delta` in the output, the test of $H_0: \delta = 0$ is sensible. In fact, it corresponds to the test of homoskedastic within-plot errors because under the null hypothesis the error variance $\text{Var}(\epsilon_{ij}) = \sigma^2(\widehat{\text{weight}}_{ij})^{2\delta}$ reduces to σ^2 .

We can use the `test` command to perform this test.

```
. test _b[/Residual:delta] = 0
( 1)  [/Residual]delta = 0
      chi2( 1) = 1774.87
      Prob > chi2 =    0.0000
```

The Wald test strongly rejects the null hypothesis of homoskedastic errors.



➤ Example 2: Obtaining predictions

Continuing with [example 1](#), we can also obtain the estimates of the plot-level random effects u_{1j} 's. Because `menl` used the Lindstrom–Bates linearization method, the estimated random effects are essentially BLUPs; see [Inference based on linearization](#) in [ME] `menl`.


We need to specify the name of the variable to be created and then use `predict, reffects`. For example, below we obtain the predictions of random effects for the first 10 plots.

```
. predict u1, reffects
. by plot, sort: generate tolist = (_n==1)
. list plot u1 if plot <=10 & tolist
```

	plot	u1
1.	1988F4	-1.716238
11.	1988F2	-.1668753
20.	1988F1	-.2153712
30.	1988F7	-.3337681
39.	1988F5	1.331566
47.	1988F8	-.7153563
57.	1988F6	.0699628
67.	1988F3	1.353845
77.	1988P1	-.6681811
87.	1988P5	-.9152615

Next, we obtain the predicted mean values and plot them. By default, the mean response conditional on the estimated random effects is computed. Predicted values based on the fixed-effects estimates alone, that is, conditional on zero random effects, may be obtained by specifying the `fixedonly` option.

```
. predict fitweight, yhat
. twoway connected weight fitweight time if plot<=9, sort by(plot)
> ytitle("Average leaf weight per plant (g)")
> legend(order(1 "observed" 2 "predicted"))
```



The predicted values closely match the observed average leaf weights, confirming the adequacy of the model.

Also see [example 13](#) in [\[ME\] menl](#) for how to predict parameters defined as functions of other parameters with substitutable expressions.

➤ **Example 3: Checking model assumptions based on residuals**


The raw residuals are useful to check for heterogeneity of the within-group error variance; see [example 10](#) in [\[ME\] menl](#). They are less recommended, however, for checking normality assumptions and for detecting outlying observations. This is because raw residuals are usually correlated and have different variances. Instead, we can use standardized residuals to check for normality and outlying observations. If the normality assumption is reasonable and the model fits data well, standardized residuals should follow a standard normal distribution.

Let’s check whether the standardized residuals from our model are approximately normally distributed with mean zero and variance one.

```
. predict rs, rstandard
. summarize rs
```

Variable	Obs	Mean	Std. Dev.	Min	Max
rs	412	.0150192	.9564895	-2.547287	3.661603

```
. qnorm rs
```



The plot does not indicate serious departures from normality, and the estimated mean and standard deviation are close to zero and one, respectively. It appears that the power of the mean function is a reasonable choice for modeling heteroskedasticity of the within-group errors in this example.

➤ **Example 4: estat group and level-specific predictions**

In [example 19](#) of [\[ME\] menl](#), we modeled the intensity of current at the i th level of **voltage** in the j th site within the k th wafer as

$$\text{current}_{ijk} = \phi_{1jk} + \phi_{2jk} \cos(\phi_{3jk} \text{voltage}_i + \pi/4) + \epsilon_{ijk}$$

for $k = 1, \dots, 10$, $j = 1, \dots, 8$, and $i = 1, \dots, 5$. In that example, we considered fairly complicated specifications for ϕ_j ’s in stage 2 with many random effects at different levels, which lead to slow execution of the command. To illustrate some of the commands available after **menl**, we will substantially simplify the stage 2 specification to speed up the estimation of the model.

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
wafer: Independent					
	var(W0)	.0048116	.0027516	.0015686	.014759
	var(W1)	.0396212	.0184213	.0159285	.0985555
wafer>site: Independent					
	var(S0)	.0086449	.0017812	.0057726	.0129463
	var(S1)	.0118703	.0021115	.0083763	.0168217
	var(Residual)	.001069	.0000952	.0008978	.0012729

We can use estat group to see how the data are broken down by wafer and site:

```
. estat group
```

Path	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
wafer	10	40	40.0	40
wafer>site	80	5	5.0	5

We are reminded that we have balanced data for each site (all sites were measured at 5 ascending voltages).

Suppose that we want to predict random effects at the wafer level only; that is, we want to compute $\hat{\mathbf{u}}_k^{(3)}$. This can be done by specifying the relevel(wafer) option:

```
. predict u_wafer*, reffects relevel(wafer)
```

Notice how predict labels the generated variables for you to avoid confusion.

```
. describe u_wafer*
```

variable name	storage type	display format	value label	variable label
u_wafer1	float	%9.0g		BLUP r.e. for W0[wafer]
u_wafer2	float	%9.0g		BLUP r.e. for W1[wafer]

We can use predict, yhat to get the predicted values $\mu(\text{voltage}_i, \hat{\beta}, \hat{\mathbf{u}}_k^{(3)}, \hat{\mathbf{u}}_{j,k}^{(2)})$. If instead we want to predict values at the wafer level, $\mu(\text{voltage}_i, \hat{\beta}, \hat{\mathbf{u}}_k^{(3)}, \mathbf{0})$, we again need to specify the relevel() option:

```
. predict curr_wafer, yhat relevel(wafer)
. list wafer site current curr_wafer in 1/10
```

	wafer	site	current	curr_wafer
1.	1	1	.90088	.8898317
2.	1	1	3.8682	3.920231
3.	1	1	7.6406	7.65254
4.	1	1	11.736	11.76189
5.	1	1	15.934	15.91457
6.	1	2	1.032	.8898317
7.	1	2	4.1022	3.920231
8.	1	2	7.9316	7.65254
9.	1	2	12.064	11.76189
10.	1	2	16.294	15.91457

The predicted values `curr_wafer` do not vary across sites, because $\mu(\text{voltage}_i, \hat{\beta}, \hat{\mathbf{u}}_k^{(3)}, \mathbf{0})$ does not depend on j .

◀

Methods and formulas

Following the notation defined throughout [ME] **menl**, estimates of random effects $\mathbf{u}_j$ are obtained by using PNLS iterations with parameters β , α , and σ^2 held fixed at their values obtained at convergence. Starting with $\hat{\mathbf{u}}_j^{(0)} = \mathbf{0}$, at the k th iteration, we have

$$\hat{\mathbf{u}}_j^{(k)} = \hat{\Sigma} \hat{\mathbf{Z}}_j'^{(k-1)} \left(\hat{\mathbf{Z}}_j^{(k-1)} \hat{\Sigma} \hat{\mathbf{Z}}_j'^{(k-1)} + \hat{\sigma}^2 \hat{\Lambda}_j \right)^{-1} \left(\hat{\mathbf{w}}_j^{(k-1)} - \hat{\mathbf{X}}_j^{(k-1)} \hat{\beta} \right)$$

where $\hat{\Sigma}$ and $\hat{\Lambda}$ are Σ and Λ with maximum likelihood (ML) or restricted maximum likelihood (REML) estimates of the variance components plugged in and $\hat{\mathbf{X}}_j^{(k-1)} = \hat{\mathbf{X}}_j(\hat{\mathbf{u}}_j^{(k-1)})$, $\hat{\mathbf{Z}}_j^{(k-1)} = \hat{\mathbf{Z}}_j(\hat{\mathbf{u}}_j^{(k-1)})$, and $\hat{\mathbf{w}}_j^{(k-1)} = \hat{\mathbf{w}}_j(\hat{\mathbf{u}}_j^{(k-1)})$ are defined in the LME step of *Inference based on linearization* in *Methods and formulas* of [ME] **menl**. When the variance structure depends on $\mathbf{u}_j$, such as when the `resvariance(power _yhat)` option is specified during estimation, $\hat{\Lambda}_j$ will also be updated at each iteration; that is, $\hat{\Lambda}_j = \hat{\Lambda}_j(\hat{\mathbf{u}}_j^{(k-1)})$. The iterative process stops when the relative difference between $\hat{\mathbf{u}}_j^{(k-1)}$ and $\hat{\mathbf{u}}_j^{(k)}$ is less than `tolerance(#)` or, if the stopping rule is not met, when the maximum number of iterations in `iterate(#)` is reached.

Standard errors for the estimates of the random effects are calculated based on [Bates and Pinheiro \(1998, sec. 3.3\)](#). If estimation is done by REML, these standard errors account for uncertainty in the estimate of β , whereas for ML, the standard errors treat β as known. As such, standard errors of REML-based estimates will usually be larger.

Predicted mean values are computed as $\mu_j(\mathbf{X}_j, \hat{\beta}, \hat{\mathbf{u}}_j)$, predicted parameters as $\hat{\phi}_j = d(\mathbf{x}_j^b, \hat{\beta}, \hat{\mathbf{u}}_j)$, residuals as $\hat{\epsilon}_j = \mathbf{y}_j - \mu_j(\mathbf{X}_j, \hat{\beta}, \hat{\mathbf{u}}_j)$, and standardized residuals as

$$\hat{\epsilon}_j^* = \hat{\sigma}^{-1} \hat{\Lambda}_j^{-1/2} \hat{\epsilon}_j$$

If the **relevel** (*levelvar*) option is specified, predicted values, residuals, and standardized residuals consider only those random-effects terms up to and including level *levelvar* in the model. If the **fixedonly** option is specified, all statistics and named substitutable expressions are computed conditional on zero random effects; that is, their computation is based on the estimated fixed effects only.

Reference

Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies.
<http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.

Also see

[ME] **menl** — Nonlinear mixed-effects regression

[U] **20 Estimation and postestimation commands**

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

meologit fits mixed-effects logistic models for ordered responses. The actual values taken on by the response are irrelevant except that larger values are assumed to correspond to “higher” outcomes. The conditional distribution of the response given the random effects is assumed to be multinomial, with success probability determined by the logistic cumulative distribution function.

Quick start

Two-level ordered logit regression of y on [indicators](#) for levels of a and random intercepts by lev2

```
meologit y i.a || lev2:
```

Two-level model including fixed and random coefficients for x

```
meologit y i.a x || lev2: x
```

As above, but report odds ratios instead of coefficients

```
meologit y i.a x || lev2: x, or
```

Three-level model of y on a, x, and their interaction using [factor variable](#) notation and random intercepts by lev2 and lev3 with lev2 nested within lev3

```
meologit y a##c.x || lev3: || lev2:
```

Menu

Statistics > Multilevel mixed-effects models > Ordered logistic regression

Syntax

```
meologit depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
or	report fixed-effects coefficients as odds ratios
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<u>mv</u> aghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mc</u> aghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: meologit.

vce() and *weights* are not allowed with the *svy* prefix; see [SVY] svy.

*fw*weights, *iw*weights, and *pw*weights are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, *unstructured*, *fixed*(*matname*), or *pattern*(*matname*).

covariance(*independent*) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is *covariance*(*independent*) unless a crossed random-effects model is fit, in which case the default is *covariance*(*identity*).

covariance(*exchangeable*) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(*identity*) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(*unstructured*) allows for all variances and covariances to be distinct. If an equation consists of *p* random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(*fixed*(*matname*)) and *covariance*(*pattern*(*matname*)) covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a *fixed*(*matname*) covariance structure, (co)variance (*i*, *j*) is constrained to equal the value specified in the *i*, *j*th entry of *matname*. In a *pattern*(*matname*) covariance structure, (co)variances (*i*, *j*) and (*k*, *l*) are constrained to be equal if *matname*[*i*, *j*] = *matname*[*k*, *l*].

noconstant suppresses the constant (intercept) term; may be specified for any of or all the random-effects equations.

fweight(*varname*) specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, [*fw=fwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **fweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) frequency weights, and **wt2** would hold the second-level (the school-level) frequency weights.

iweight(*varname*) specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, [*iw=iwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **iweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) importance weights, and **wt2** would hold the second-level (the school-level) importance weights.

pweight(*varname*) specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, [*pw=pwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **pweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable **wt1** would hold the first-level (the observation-level) sampling weights, and **wt2** would hold the second-level (the school-level) sampling weights.

constraints(*constraints*), **collinear**; see [R] [estimation options](#).

SE/Robust

vce(*vcetype*) specifies the type of standard error reported, which includes types that are derived from asymptotic theory (**oim**), that are robust to some kinds of misspecification (**robust**), and that allow for intragroup correlation (**cluster** *clustvar*); see [R] [vce-option](#). If **vce**(**robust**) is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

level(#); see [R] [estimation options](#).

or reports estimated fixed-effects coefficients transformed to odds ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated. **or** may be specified either at estimation or upon replay.

nocnsreport; see [R] [estimation options](#).

notable suppresses the estimation table, either at estimation or upon replay.

noheader suppresses the output header, either at estimation or upon replay.

nogroup suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: **nocl**, **nopvalues**, **noomitted**, **vsquish**, **noemptycells**, **baselevels**, **allbaselevels**, **nofvlabel**, **fvwrap**(#), **fvwrapon**(*style*), **cformat**(%*fnt*), **pformat**(%*fnt*), **sformat**(%*fnt*), and **nolstretch**; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

`maximize_options`: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meologit` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `meologit` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[ (gridspec) ]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects ordered logistic regression is ordered logistic regression containing both fixed effects and random effects. An ordered response is a variable that is categorical and ordered, for instance, “poor”, “good”, and “excellent”, which might indicate a person’s current health status or the repair record of a car.

`meologit` allows for many levels of random effects. However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of fixed effects $\mathbf{x}_{ij}$, a set of cutpoints κ , and a set of random effects $\mathbf{u}_j$, the cumulative probability of the response being in a category higher than k is

$$\Pr(y_{ij} > k | \mathbf{x}_{ij}, \kappa, \mathbf{u}_j) = H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j - \kappa_k) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The cutpoints κ are labeled $\kappa_1, \kappa_2, \dots, \kappa_{K-1}$, where K is the number of possible outcomes. $H(\cdot)$ is the logistic cumulative distribution function that represents cumulative probability.

The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard logistic regression model, with regression coefficients (fixed effects) β . In our parameterization, $\mathbf{x}_{ij}$ does not contain a constant term because its effect is absorbed into the cutpoints. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of Σ , known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$, so that all covariate effects are essentially random and distributed as multivariate normal with mean β and variance Σ .

From (1), we can derive the probability of observing outcome k as

$$\begin{aligned} \Pr(y_{ij} = k | \kappa, \mathbf{u}_j) &= \Pr(\kappa_{k-1} < \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij} \leq \kappa_k) \\ &= \Pr(\kappa_{k-1} - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j < \epsilon_{ij} \leq \kappa_k - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) \\ &= H(\kappa_k - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) - H(\kappa_{k-1} - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) \end{aligned}$$

where κ_0 is taken as $-\infty$ and κ_K is taken as $+\infty$.

From the above, we may also write the model in terms of a latent linear response, where observed ordinal responses y_{ij} are generated from the latent continuous responses, such that

$$y_{ij}^* = \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

and

$$y_{ij} = \begin{cases} 1 & \text{if } y_{ij}^* \leq \kappa_1 \\ 2 & \text{if } \kappa_1 < y_{ij}^* \leq \kappa_2 \\ \vdots & \\ K & \text{if } \kappa_{K-1} < y_{ij}^* \end{cases}$$

The errors ϵ_{ij} are distributed as logistic with mean 0 and variance $\pi^2/3$ and are independent of $\mathbf{u}_j$.

Below we present two short examples of mixed-effects ordered logistic regression; refer to [ME] **me** and [ME] **meglm** for examples of other random-effects models. A two-level ordered logistic model can also be fit using **xtologit** with the **re** option; see [XT] **xtologit**. In the absence of random effects, mixed-effects ordered logistic regression reduces to standard ordered logistic regression; see [R] **ologit**.

► Example 1: Two-level random-intercept model

We use the data from the Television, School, and Family Smoking Prevention and Cessation Project (Flay et al. 1988; Rabe-Hesketh and Skrondal 2012, chap. 11), where schools were randomly assigned into one of four groups defined by two treatment variables. Students within each school are nested in classes, and classes are nested in schools. In this example, we ignore the variability of classes within schools and fit a two-level model; we incorporate classes in a three-level model in [example 2](#). The dependent variable is the tobacco and health knowledge (THK) scale score collapsed into four ordered categories. We regress the outcome on the treatment variables and their interaction and control for the pretreatment score.

```
. use http://www.stata-press.com/data/r15/tvsfpors
. meologit thk prethk cc##tv || school:
Fitting fixed-effects model:
Iteration 0:   log likelihood =  -2212.775
Iteration 1:   log likelihood =  -2125.509
Iteration 2:   log likelihood =  -2125.1034
Iteration 3:   log likelihood =  -2125.1032
Refining starting values:
Grid node 0:   log likelihood =  -2136.2426
Fitting full model:
Iteration 0:   log likelihood =  -2136.2426   (not concave)
Iteration 1:   log likelihood =  -2120.2577
Iteration 2:   log likelihood =  -2119.7574
Iteration 3:   log likelihood =  -2119.7428
Iteration 4:   log likelihood =  -2119.7428

Mixed-effects ologit regression
Group variable:      school

Number of obs      =      1,600
Number of groups   =      28
Obs per group:
    min =      18
    avg =     57.1
    max =     137

Integration method: mvaghermite
Integration pts.   =      7
Wald chi2(4)      =     128.06
Prob > chi2       =     0.0000

Log likelihood = -2119.7428
```

thk	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
prethk	.4032892	.03886	10.38	0.000	.327125	.4794534
1.cc	.9237904	.204074	4.53	0.000	.5238127	1.323768
1.tv	.2749937	.1977424	1.39	0.164	-.1125744	.6625618
cc#tv 1 1	-.4659256	.2845963	-1.64	0.102	-1.023724	.0918728
/cut1	-.0884493	.1641062			-.4100916	.233193
/cut2	1.153364	.165616			.8287625	1.477965
/cut3	2.33195	.1734199			1.992053	2.671846
school var(_cons)	.0735112	.0383106			.0264695	.2041551

```
LR test vs. ologit model: chibar2(01) = 10.72      Prob >= chibar2 = 0.0005
```

The estimation table reports the fixed effects, the estimated cutpoints ($\kappa_1, \kappa_2, \kappa_3$), and the estimated variance components. The fixed effects can be interpreted just as you would the output from `ologit`. We find that students with higher preintervention scores tend to have higher postintervention scores. Because of their interaction, the impact of the treatment variables on the knowledge score is not straightforward; we defer this discussion to [example 1](#) of [\[ME\] meologit postestimation](#).

Underneath the fixed effects and the cutpoints, the table shows the estimated variance components. The random-effects equation is labeled `school`, meaning that these are random effects at the `school` level. Because we have only one random effect at this level, the table shows only one variance component. The estimate of σ_u^2 is 0.07 with standard error 0.04. The reported likelihood-ratio test shows that there is enough variability between schools to favor a mixed-effects ordered logistic regression over a standard ordered logistic regression; see [Distribution theory for likelihood-ratio test](#) in [ME] [me](#) for a discussion of likelihood-ratio testing of variance components.

We now store our estimates for later use.

```
. estimates store r_2
```

◀

Two-level models extend naturally to models with three or more levels with nested random effects. Below we continue with [example 1](#).

► Example 2: Three-level random-intercept model

In this example, we fit a three-level model incorporating classes nested within schools. The fixed-effects part remains the same. Our model now has two random-effects equations, separated by `||`. The first is a random intercept (constant only) at the `school` level (level three), and the second is a random intercept at the `class` level (level two). The order in which these are specified (from left to right) is significant—`meologit` assumes that `class` is nested within `school`.

```
. meologit thk prethk cc##tv || school: || class:
Fitting fixed-effects model:
Iteration 0:   log likelihood = -2212.775
Iteration 1:   log likelihood = -2125.509
Iteration 2:   log likelihood = -2125.1034
Iteration 3:   log likelihood = -2125.1032
Refining starting values:
Grid node 0:   log likelihood = -2152.1514
Fitting full model:
Iteration 0:   log likelihood = -2152.1514 (not concave)
Iteration 1:   log likelihood = -2125.9213 (not concave)
Iteration 2:   log likelihood = -2120.1861
Iteration 3:   log likelihood = -2115.6177
Iteration 4:   log likelihood = -2114.5896
Iteration 5:   log likelihood = -2114.5881
Iteration 6:   log likelihood = -2114.5881
Mixed-effects ologit regression                                Number of obs      =      1,600
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
school	28	18	57.1	137
class	135	1	11.9	28

```
Integration method: mvaghermite                                Integration pts. =      7
                                                                Wald chi2(4)      =    124.39
Log likelihood = -2114.5881                                    Prob > chi2       =    0.0000
```

thk	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
prethk	.4085273	.039616	10.31	0.000	.3308814	.4861731
1.cc	.8844369	.2099124	4.21	0.000	.4730161	1.295858
1.tv	.236448	.2049065	1.15	0.249	-.1651614	.6380575
cc##tv 1 1	-.3717699	.2958887	-1.26	0.209	-.951701	.2081612
/cut1	-.0959459	.1688988			-.4269815	.2350896
/cut2	1.177478	.1704946			.8433151	1.511642
/cut3	2.383672	.1786736			2.033478	2.733865
school var(_cons)	.0448735	.0425387			.0069997	.2876749
school>class var(_cons)	.1482157	.0637521			.063792	.3443674

```
LR test vs. ologit model: chi2(2) = 21.03                      Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

We see that we have 135 classes from 28 schools. The variance-component estimates are now organized and labeled according to level. The variance component for `class` is labeled `school>class` to emphasize that classes are nested within schools.

Compared with the two-level model from [example 1](#), the estimate of the variance of the random intercept at the school level dropped from 0.07 to 0.04. This is not surprising because we now use two random components versus one random component to account for unobserved heterogeneity among students. We can use `lrtest` and our stored estimation result from [example 1](#) to see which model provides a better fit:

```
. lrtest r_2 .
Likelihood-ratio test                                LR chi2(1) =      10.31
(Assumption: r_2 nested in .)                       Prob > chi2 =      0.0013

Note: The reported degrees of freedom assumes the null hypothesis is not on
      the boundary of the parameter space.  If this is not true, then the
      reported test is conservative.
```

The likelihood-ratio test favors the three-level model. For more information about the likelihood-ratio test in the context of mixed-effects models, see [Distribution theory for likelihood-ratio test](#) in [ME] [me](#).

◀

The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by `||`.

Stored results

`meologit` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_cat)</code>	number of categories
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>meologit</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified

<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	<code>ologit</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>logit</code>
<code>e(family)</code>	<code>ordinal</code>
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	<code>offset</code>
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	<code>max</code> or <code>min</code> ; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of <code>ml</code> method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>
<code>e(marginswexp)</code>	weight expression for <code>margins</code>
<code>e(marginsdefault)</code>	default <code>predict()</code> specification for <code>margins</code>
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

`meologit` is a convenience command for `meglm` with a `logit` link and an `ordinal` family; see [ME] `meglm`.

Without a loss of generality, consider a two-level ordered logistic model. The probability of observing outcome k for response y_{ij} is then

$$\begin{aligned} p_{ij} &= \Pr(y_{ij} = k | \boldsymbol{\kappa}, \mathbf{u}_j) = \Pr(\kappa_{k-1} < \boldsymbol{\eta}_{ij} + \epsilon_{it} \leq \kappa_k) \\ &= \frac{1}{1 + \exp(-\kappa_k + \boldsymbol{\eta}_{ij})} - \frac{1}{1 + \exp(-\kappa_{k-1} + \boldsymbol{\eta}_{ij})} \end{aligned}$$

where $\boldsymbol{\eta}_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \text{offset}_{ij}$, κ_0 is taken as $-\infty$, and κ_K is taken as $+\infty$. Here $\mathbf{x}_{ij}$ does not contain a constant term because its effect is absorbed into the cutpoints.

For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$ given a set of cluster-level random effects $\mathbf{u}_j$ is

$$\begin{aligned} f(\mathbf{y}_j | \boldsymbol{\kappa}, \mathbf{u}_j) &= \prod_{i=1}^{n_j} p_{ij}^{I_k(y_{ij})} \\ &= \exp \sum_{i=1}^{n_j} \left\{ I_k(y_{ij}) \log(p_{ij}) \right\} \end{aligned}$$

where

$$I_k(y_{ij}) = \begin{cases} 1 & \text{if } y_{ij} = k \\ 0 & \text{otherwise} \end{cases}$$

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{y}_j | \boldsymbol{\kappa}, \mathbf{u}_j) \exp(-\mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int \exp\{h(\boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \tag{2}$$

where

$$h(\boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Sigma}, \mathbf{u}_j) = \sum_{i=1}^{n_j} \left\{ I_k(y_{ij}) \log(p_{ij}) \right\} - \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see *Methods and formulas* in [ME] `meglm` for details.

`meologit` supports multilevel weights and survey data; see *Methods and formulas* in [ME] `meglm` for details.

References

- Flay, B. R., B. R. Brannon, C. A. Johnson, W. B. Hansen, A. L. Ulene, D. A. Whitney-Saltiel, L. R. Gleason, S. Sussman, M. D. Gavin, K. M. Glowacz, D. F. Sobol, and D. C. Spiegel. 1988. The television, school, and family smoking cessation and prevention project: I. Theoretical basis and program development. *Preventive Medicine* 17: 585–607.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.

Also see

- [ME] **meologit postestimation** — Postestimation tools for meologit
- [ME] **meoprobit** — Multilevel mixed-effects ordered probit regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: meologit** — Bayesian multilevel ordered logistic regression
- [SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtologit** — Random-effects ordered logistic models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation command is of special interest after `meologit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike’s and Schwarz’s Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman’s specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as probabilities, linear predictions, density and distribution functions, and standard errors.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
<code>pr</code>	predicted probabilities; the default
<code>eta</code>	fitted linear predictor
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code>density</code>	predicted density function
<code>distribution</code>	predicted distribution function

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
<u>outcome</u> (<i>outcome</i>)	outcome category for predicted probabilities
Integration	
<u>int_options</u>	integration options
You specify one or k new variables in <i>newvarlist</i> with <code>pr</code> , where k is the number of outcomes. If you do not specify <code>outcome()</code> , these options assume <code>outcome(#1)</code> .	
<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only
<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<u>int_options</u>	integration options
<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`pr`, the default, calculates the predicted probabilities.

You specify one or k new variables, where k is the number of categories of the dependent variable. If you specify the `outcome()` option, the probabilities will be predicted for the requested outcome only, in which case you specify only one new variable. If you specify one new variable and do not specify `outcome()`, `outcome(#1)` is assumed.

`eta`, `xb`, `stdp`, `density`, `distribution`, `scores`, `conditional()`, `marginal`, and `nooffset`; see [ME] **meglm postestimation**.

`outcome(outcome)` specifies the outcome for which the predicted probabilities are to be calculated. `outcome()` should contain either one value of the dependent variable or one of #1, #2, ..., with #1 meaning the first category of the dependent variable, #2 meaning the second category, etc.

`reffects`, `ebmeans`, `ebmodes`, and `reses()`, see [ME] **meglm postestimation**.

Integration

`intpoints()`, `iterate()`, `tolerance()`; see [ME] **meglm postestimation**.

margins

Description for margins

`margins` estimates margins of response for probabilities and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
default	probabilities for each outcome
pr	predicted probabilities for a specified outcome
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
density	not allowed with margins
distribution	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

`pr` defaults to the first outcome.

Options `conditional(ebmeans)` and `conditional(ebmodes)` are not allowed with `margins`.

Option `marginal` is assumed where applicable if `conditional(fixedonly)` is not specified.

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.

For the full syntax, see [R] **margins**.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting an ordered logistic mixed-effects model with `meologit`. Here we show a short example of predicted probabilities and predicted random effects; refer to [ME] `meglm postestimation` for additional examples applicable to mixed-effects generalized linear models.

► Example 1: Obtaining predicted probabilities and random effects

In [example 2](#) of [ME] `meologit`, we modeled the tobacco and health knowledge (`thk`) score—coded 1, 2, 3, 4—among students as a function of two treatments (`cc` and `tv`) by using a three-level ordered logistic model with random effects at the school and class levels.

```
. use http://www.stata-press.com/data/r15/tvsfpors
. meologit thk prethk cc##tv || school: || class:
(output omitted)
```

We obtain predicted probabilities for all four outcomes based on the contribution of both fixed effects and random effects by typing

```
. predict pr*
(predictions based on fixed effects and posterior means of random effects)
(option mu assumed)
(using 7 quadrature points)
```

As the note says, the predicted values are based on the posterior means of random effects. You can use the `modes` option to obtain predictions based on the posterior modes of random effects.

Because we specified a stub name, Stata saved the predicted random effects in variables `pr1` through `pr4`. Here we list the predicted probabilities for the first two classes for school 515:


```
. list class thk pr? if school==515 & (class==515101 | class==515102),  
> sepby(class)
```

	class	thk	pr1	pr2	pr3	pr4
1464.	515101	2	.1485538	.2354556	.2915916	.3243991
1465.	515101	2	.372757	.3070787	.1966117	.1235526
1466.	515101	1	.372757	.3070787	.1966117	.1235526
1467.	515101	4	.2831409	.3021398	.2397316	.1749877
1468.	515101	3	.2079277	.2760683	.2740791	.2419248
1469.	515101	3	.2831409	.3021398	.2397316	.1749877
1470.	515102	1	.3251654	.3074122	.2193101	.1481123
1471.	515102	2	.4202843	.3011963	.1749344	.103585
1472.	515102	2	.4202843	.3011963	.1749344	.103585
1473.	515102	2	.4202843	.3011963	.1749344	.103585
1474.	515102	2	.3251654	.3074122	.2193101	.1481123
1475.	515102	1	.4202843	.3011963	.1749344	.103585
1476.	515102	2	.3251654	.3074122	.2193101	.1481123

For each observation, our best guess for the predicted outcome is the one with the highest predicted probability. For example, for the very first observation in the table above, we would choose outcome 4 as the most likely to occur.

We obtain predictions of the posterior means themselves at the school and class levels by typing

```
. predict re_s re_c, reffects  
(calculating posterior means of random effects)  
(using 7 quadrature points)
```

Here we list the predicted random effects for the first two classes for school 515:

```
. list class re_s re_c if school==515 & (class==515101 | class==515102),  
> sepby(class)
```

	class	re_s	re_c
1464.	515101	-.0473739	.0633081
1465.	515101	-.0473739	.0633081
1466.	515101	-.0473739	.0633081
1467.	515101	-.0473739	.0633081
1468.	515101	-.0473739	.0633081
1469.	515101	-.0473739	.0633081
1470.	515102	-.0473739	-.1354929
1471.	515102	-.0473739	-.1354929
1472.	515102	-.0473739	-.1354929
1473.	515102	-.0473739	-.1354929
1474.	515102	-.0473739	-.1354929
1475.	515102	-.0473739	-.1354929
1476.	515102	-.0473739	-.1354929

We can see that the predicted random effects at the school level (**re_s**) are the same for all classes and that the predicted random effects at the class level (**re_c**) are constant within each class.

Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in *Methods and formulas* of [ME] **meglm postestimation**.

Also see

[ME] **meologit** — Multilevel mixed-effects ordered logistic regression

[ME] **meglm postestimation** — Postestimation tools for meglm

[U] **20 Estimation and postestimation commands**

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`meoprobit` fits mixed-effects probit models for ordered responses. The actual values taken on by the response are irrelevant except that larger values are assumed to correspond to “higher” outcomes. The conditional distribution of the response given the random effects is assumed to be multinomial, with success probability determined by the standard normal cumulative distribution function.

Quick start

```
Two-level ordered probit regression of y on x and random intercepts by lev2
  meoprobit y x || lev2:

Add random coefficients for x
  meoprobit y x || lev2: x

Nested three-level ordered probit model with random intercepts by lev2 and lev3 for lev2 nested
within lev3
  meoprobit y x || lev3: || lev2:
```

Menu

Statistics > Multilevel mixed-effects models > Ordered probit regression

Syntax

```
meoprobit depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <u>oim</u> , <u>robust</u> , or <u>cluster</u> <i>clustvar</i>
Reporting	
<u>level</u> (#)	set confidence level; default is <u>level</u> (95)
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <u>intpoints</u> (7)
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <u>R.</u> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <u>R.</u> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<u>mv</u> aghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mc</u> aghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] *bayes*: meoprobfit.

vce() and weights are not allowed with the *svy* prefix; see [SVY] *svy*.

*fw*eight, *iw*eight, and *pw*eight are allowed; see [U] 11.1.6 *weight*. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, *unstructured*, *fixed*(*matname*), or *pattern*(*matname*).

covariance(*independent*) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is *covariance*(*independent*) unless a crossed random-effects model is fit, in which case the default is *covariance*(*identity*).

covariance(*exchangeable*) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(*identity*) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(*unstructured*) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(*fixed*(*matname*)) and *covariance*(*pattern*(*matname*)) covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a *fixed*(*matname*) covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a *pattern*(*matname*) covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if *matname*[i, j] = *matname*[k, l].

noconstant suppresses the constant (intercept) term; may be specified for any of or all the random-effects equations.

fweight(*varname*) specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, [*fw=fwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **fweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) frequency weights, and **wt2** would hold the second-level (the school-level) frequency weights.

iweight(*varname*) specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, [*iw=iwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **iweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) importance weights, and **wt2** would hold the second-level (the school-level) importance weights.

pweight(*varname*) specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, [*pw=pwtvar1*]. *varname* can be any valid Stata variable name, and you can specify **pweight**() at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable **wt1** would hold the first-level (the observation-level) sampling weights, and **wt2** would hold the second-level (the school-level) sampling weights.

constraints(*constraints*), *collinear*; see [\[R\] estimation options](#).

SE/Robust

vce(*vcetype*) specifies the type of standard error reported, which includes types that are derived from asymptotic theory (*oim*), that are robust to some kinds of misspecification (*robust*), and that allow for intragroup correlation (*cluster clustvar*); see [\[R\] vce_option](#). If **vce**(*robust*) is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

level(#), *nocnsreport*; see [\[R\] estimation options](#).

notable suppresses the estimation table, either at estimation or upon replay.

noheader suppresses the output header, either at estimation or upon replay.

nogroup suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: *nocl*, *nopvalues*, *noomitted*, *vsquish*, *noemptycells*, *baselevels*, *allbaselevels*, *nofvlabel*, *fvwrap*(#), *fvwrapon*(*style*), *cformat*(%*fmt*), *pformat*(%*fmt*), *sformat*(%*fmt*), and *nolstretch*; see [\[R\] estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

`maximize_options`: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meoprobit` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `meoprobit` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[ (gridspec) ]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects ordered probit regression is ordered probit regression containing both fixed effects and random effects. An ordered response is a variable that is categorical and ordered, for instance, “poor”, “good”, and “excellent”, which might indicate a person’s current health status or the repair record of a car.

`meoprobit` allows for many levels of random effects. However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of fixed effects $\mathbf{x}_{ij}$, a set of cutpoints $\boldsymbol{\kappa}$, and a set of random effects $\mathbf{u}_j$, the cumulative probability of the response being in a category higher than k is

$$\Pr(y_{ij} > k | \mathbf{x}_{ij}, \boldsymbol{\kappa}, \mathbf{u}_j) = \Phi(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j - \kappa_k) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The cutpoints are labeled $\kappa_1, \kappa_2, \dots, \kappa_{K-1}$, where K is the number of possible outcomes. $\Phi(\cdot)$ is the standard normal cumulative distribution function that represents cumulative probability.

The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard probit regression model, with regression coefficients (fixed effects) β . In our parameterization, $\mathbf{x}_{ij}$ does not contain a constant term because its effect is absorbed into the cutpoints. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of Σ , known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean β and variance Σ .

From (1), we can derive the probability of observing outcome k as

$$\begin{aligned} \Pr(y_{ij} = k | \kappa, \mathbf{u}_j) &= \Pr(\kappa_{k-1} < \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij} \leq \kappa_k) \\ &= \Pr(\kappa_{k-1} - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j < \epsilon_{ij} \leq \kappa_k - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) \\ &= \Phi(\kappa_k - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) - \Phi(\kappa_{k-1} - \mathbf{x}_{ij}\beta - \mathbf{z}_{ij}\mathbf{u}_j) \end{aligned}$$

where κ_0 is taken as $-\infty$ and κ_K is taken as $+\infty$.

From the above, we may also write the model in terms of a latent linear response, where observed ordinal responses y_{ij} are generated from the latent continuous responses, such that

$$y_{ij}^* = \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

and

$$y_{ij} = \begin{cases} 1 & \text{if } y_{ij}^* \leq \kappa_1 \\ 2 & \text{if } \kappa_1 < y_{ij}^* \leq \kappa_2 \\ \vdots & \\ K & \text{if } \kappa_{K-1} < y_{ij}^* \end{cases}$$

The errors ϵ_{ij} are distributed as standard normal with mean 0 and variance 1 and are independent of $\mathbf{u}_j$.

Below we present two short examples of mixed-effects ordered probit regression; refer to [ME] **me** and [ME] **meglm** for examples of other random-effects models. A two-level ordered probit model can also be fit using **xtoprobit** with the **re** option; see [XT] **xtoprobit**. In the absence of random effects, mixed-effects ordered probit regression reduces to standard ordered probit regression; see [R] **oprobit**.

► Example 1: Two-level random-intercept model

We use the data from the Television, School, and Family Smoking Prevention and Cessation Project (Flay et al. 1988; Rabe-Hesketh and Skrondal 2012, chap. 11), where schools were randomly assigned into one of four groups defined by two treatment variables. Students within each school are nested in classes, and classes are nested in schools. In this example, we ignore the variability of classes within schools and fit a two-level model; we incorporate classes in a three-level model in [example 2](#). The dependent variable is the tobacco and health knowledge (THK) scale score collapsed into four ordered categories. We regress the outcome on the treatment variables and their interaction and control for the pretreatment score.

```

. use http://www.stata-press.com/data/r15/tvsfpor
. meoprobit thk prethk cc##tv || school:
Fitting fixed-effects model:
Iteration 0:   log likelihood =  -2212.775
Iteration 1:   log likelihood = -2127.8111
Iteration 2:   log likelihood = -2127.7612
Iteration 3:   log likelihood = -2127.7612
Refining starting values:
Grid node 0:   log likelihood = -2149.7302
Fitting full model:
Iteration 0:   log likelihood = -2149.7302   (not concave)
Iteration 1:   log likelihood = -2129.6838   (not concave)
Iteration 2:   log likelihood = -2123.5143
Iteration 3:   log likelihood = -2122.2896
Iteration 4:   log likelihood = -2121.7949
Iteration 5:   log likelihood = -2121.7716
Iteration 6:   log likelihood = -2121.7715

Mixed-effects opobit regression              Number of obs      =       1,600
Group variable:                school          Number of groups   =        28
                                   Obs per group:
                                   min =         18
                                   avg =        57.1
                                   max =        137

Integration method: mvaghermite              Integration pts.   =         7
Wald chi2(4)                      =       128.05
Prob > chi2                        =       0.0000

Log likelihood = -2121.7715

```

thk	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
prethk	.2369804	.0227739	10.41	0.000	.1923444	.2816164
1.cc	.5490957	.1255108	4.37	0.000	.303099	.7950923
1.tv	.1695405	.1215889	1.39	0.163	-.0687693	.4078504
cc#tv						
1 1	-.2951837	.1751969	-1.68	0.092	-.6385634	.0481959
/cut1	-.0682011	.1003374			-.2648587	.1284565
/cut2	.67681	.1008836			.4790817	.8745382
/cut3	1.390649	.1037494			1.187304	1.593995
school						
var(_cons)	.0288527	.0146201			.0106874	.0778937

```

LR test vs. opobit model: chibar2(01) = 11.98      Prob >= chibar2 = 0.0003

```

The estimation table reports the fixed effects, the estimated cutpoints ($\kappa_1, \kappa_2, \kappa_3$), and the estimated variance components. The fixed effects can be interpreted just as you would the output from `oprobit`. We find that students with higher preintervention scores tend to have higher postintervention scores. Because of their interaction, the impact of the treatment variables on the knowledge score is not straightforward; we defer this discussion to [example 1](#) of [\[ME\] meoprobit postestimation](#).

Underneath the fixed effects and the cutpoints, the table shows the estimated variance components. The random-effects equation is labeled `school`, meaning that these are random effects at the `school` level. Because we have only one random effect at this level, the table shows only one variance component. The estimate of σ_u^2 is 0.03 with standard error 0.01. The reported likelihood-ratio test shows that there is enough variability between schools to favor a mixed-effects ordered probit regression over a standard ordered probit regression; see *Distribution theory for likelihood-ratio test* in [ME] **me** for a discussion of likelihood-ratio testing of variance components.

We now store our estimates for later use.

```
. estimates store r_2
```

◀

Two-level models extend naturally to models with three or more levels with nested random effects. Below we continue with [example 1](#).

▷ Example 2: Three-level random-intercept model

In this example, we fit a three-level model incorporating classes nested within schools. The fixed-effects part remains the same. Our model now has two random-effects equations, separated by `||`. The first is a random intercept (constant only) at the `school` level (level three), and the second is a random intercept at the `class` level (level two). The order in which these are specified (from left to right) is significant—`meoprobit` assumes that `class` is nested within `school`.

components versus one random component to account for unobserved heterogeneity among students. We can use `lrtest` and our stored estimation result from example 1 to see which model provides a better fit:

```
. lrtest r_2 .
Likelihood-ratio test                                LR chi2(1) =      10.15
(Assumption: r_2 nested in .)                       Prob > chi2 =      0.0014
Note: The reported degrees of freedom assumes the null hypothesis is not on
      the boundary of the parameter space.  If this is not true, then the
      reported test is conservative.
```

The likelihood-ratio test favors the three-level model. For more information about the likelihood-ratio test in the context of mixed-effects models, see [Distribution theory for likelihood-ratio test in \[ME\] me](#).

◀

The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by `||`.

Stored results

`meoprobit` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_cat)</code>	number of categories
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>meoprobit</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables

e(model)	oprobit
e(title)	title in estimation output
e(link)	probit
e(family)	ordinal
e(clustvar)	name of cluster variable
e(offset)	offset
e(intmethod)	integration method
e(n_quad)	number of integration points
e(chi2type)	Wald; type of model χ^2
e(vce)	<i>vcetype</i> specified in <i>vce()</i>
e(vcetype)	title used to label Std. Err.
e(opt)	type of optimization
e(which)	max or min; whether optimizer is to perform maximization or minimization
e(ml_method)	type of ml method
e(user)	name of likelihood-evaluator program
e(technique)	maximization technique
e(datasignature)	the checksum
e(datasignaturevars)	variables used in calculation of checksum
e(properties)	b V
e(estat_cmd)	program used to implement estat
e(predict)	program used to implement predict
e(marginsnotok)	predictions disallowed by margins
e(marginswtype)	weight type for margins
e(marginswexp)	weight expression for margins
e(marginsdefault)	default predict() specification for margins
e(asbalanced)	factor variables <i>fvset</i> as <i>asbalanced</i>
e(asobserved)	factor variables <i>fvset</i> as <i>asobserved</i>

Matrices

e(b)	coefficient vector
e(Cns)	constraints matrix
e(ilog)	iteration log (up to 20 iterations)
e(gradient)	gradient vector
e(N_g)	group counts
e(g_min)	group-size minimums
e(g_avg)	group-size averages
e(g_max)	group-size maximums
e(V)	variance-covariance matrix of the estimators
e(V_modelbased)	model-based variance

Functions

e(sample)	marks estimation sample

Methods and formulas

`meoprobit` is a convenience command for `meglm` with a `probit` link and an `ordinal` family; see [ME] [meglm](#).

Without a loss of generality, consider a two-level ordered probit model. The probability of observing outcome k for response y_{ij} is then

$$\begin{aligned} p_{ij} &= \Pr(y_{ij} = k | \boldsymbol{\kappa}, \mathbf{u}_j) = \Pr(\kappa_{k-1} < \boldsymbol{\eta}_{ij} + \epsilon_{it} \leq \kappa_k) \\ &= \Phi(\kappa_k - \boldsymbol{\eta}_{ij}) - \Phi(\kappa_{k-1} - \boldsymbol{\eta}_{ij}) \end{aligned}$$

where $\boldsymbol{\eta}_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \text{offset}_{ij}$, κ_0 is taken as $-\infty$, and κ_K is taken as $+\infty$. Here $\mathbf{x}_{ij}$ does not contain a constant term because its effect is absorbed into the cutpoints.

For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$ given a set of cluster-level random effects $\mathbf{u}_j$ is

$$\begin{aligned} f(\mathbf{y}_j | \mathbf{u}_j) &= \prod_{i=1}^{n_j} p_{ij}^{I_k(y_{ij})} \\ &= \exp \sum_{i=1}^{n_j} \left\{ I_k(y_{ij}) \log(p_{ij}) \right\} \end{aligned}$$

where

$$I_k(y_{ij}) = \begin{cases} 1 & \text{if } y_{ij} = k \\ 0 & \text{otherwise} \end{cases}$$

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ , the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\beta, \kappa, \Sigma) &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int f(\mathbf{y}_j | \kappa, \mathbf{u}_j) \exp(-\mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int \exp\{h(\beta, \kappa, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \quad (2)$$

where

$$h(\beta, \kappa, \Sigma, \mathbf{u}_j) = \sum_{i=1}^{n_j} \left\{ I_k(y_{ij}) \log(p_{ij}) \right\} - \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

meoprobbit supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Flay, B. R., B. R. Brannon, C. A. Johnson, W. B. Hansen, A. L. Ulene, D. A. Whitney-Saltiel, L. R. Gleason, S. Sussman, M. D. Gavin, K. M. Glowacz, D. F. Sobol, and D. C. Spiegel. 1988. The television, school, and family smoking cessation and prevention project: I. Theoretical basis and program development. *Preventive Medicine* 17: 585–607.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Harbord, R. M., and P. Whiting. 2009. [metandi: Meta-analysis of diagnostic accuracy using hierarchical logistic regression](#). *Stata Journal* 9: 211–229.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.

- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Rabe-Hesketh, S., and A. Skrondal. 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.

Also see

- [ME] [meoprobit postestimation](#) — Postestimation tools for meoprobit
- [ME] [meologit](#) — Multilevel mixed-effects ordered logistic regression
- [ME] [me](#) — Introduction to multilevel mixed-effects models
- [BAYES] [bayes: meoprobit](#) — Bayesian multilevel ordered probit regression
- [SEM] [intro 5](#) — Tour of models ([Multilevel mixed-effects models](#))
- [SVY] [svy estimation](#) — Estimation commands for survey data
- [XT] [xtoprobit](#) — Random-effects ordered probit models
- [U] [20 Estimation and postestimation commands](#)

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation command is of special interest after `meoprobit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike’s and Schwarz’s Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman’s specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as probabilities, linear predictions, density and distribution functions, and standard errors.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
<code>pr</code>	predicted probabilities; the default
<code>eta</code>	fitted linear predictor
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code>density</code>	predicted density function
<code>distribution</code>	predicted distribution function

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

options	Description
Main	
<code>conditional(ctype)</code>	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<code>marginal</code>	compute <i>statistic</i> marginally with respect to the random effects
<code>nooffset</code>	make calculation ignoring offset or exposure
<code>outcome(outcome)</code>	outcome category for predicted probabilities
Integration	
<code>int_options</code>	integration options

You specify one or k new variables in `newvarlist` with `pr`, where k is the number of outcomes. If you do not specify `outcome()`, these options assume `outcome(#1)`.

ctype	Description
<code>ebmeans</code>	empirical Bayes means of random effects; the default
<code>ebmodes</code>	empirical Bayes modes of random effects
<code>fixedonly</code>	prediction for the fixed portion of the model only

re_options	Description
Main	
<code>ebmeans</code>	use empirical Bayes means of random effects; the default
<code>ebmodes</code>	use empirical Bayes modes of random effects
<code>reses(stub* newvarlist)</code>	calculate standard errors of empirical Bayes estimates
Integration	
<code>int_options</code>	integration options

int_options	Description
<code>intpoints(#)</code>	use # quadrature points to compute marginal predictions and empirical Bayes means
<code>iterate(#)</code>	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<code>tolerance(#)</code>	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`pr`, the default, calculates the predicted probabilities.

You specify one or k new variables, where k is the number of categories of the dependent variable. If you specify the `outcome()` option, the probabilities will be predicted for the requested outcome only, in which case you specify only one new variable. If you specify one new variable and do not specify `outcome()`, `outcome(#1)` is assumed.

`eta`, `xb`, `stdp`, `density`, `distribution`, `scores`, `conditional()`, `marginal`, and `nooffset`; see [ME] [meglm postestimation](#).

`outcome(outcome)` specifies the outcome for which the predicted probabilities are to be calculated. `outcome()` should contain either one value of the dependent variable or one of #1, #2, ..., with #1 meaning the first category of the dependent variable, #2 meaning the second category, etc.

`reffects`, `ebmeans`, `ebmodes`, and `reses()`, see [ME] [meglm postestimation](#).

Integration

`intpoints()`, `iterate()`, `tolerance()`; see [ME] [meglm postestimation](#).

margins

Description for margins

`margins` estimates margins of response for probabilities and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
default	probabilities for each outcome
pr	predicted probabilities for a specified outcome
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
density	not allowed with margins
distribution	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

`pr` defaults to the first outcome.

Options `conditional(ebmeans)` and `conditional(ebmodes)` are not allowed with `margins`.

Option `marginal` is assumed where applicable if `conditional(fixedonly)` is not specified.

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.

For the full syntax, see [R] [margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting an ordered probit mixed-effects model using **meoprobit**. Here we show a short example of predicted probabilities and predicted random effects; refer to [\[ME\] meglm postestimation](#) for additional examples applicable to mixed-effects generalized linear models.

► Example 1: Obtaining predicted probabilities and random effects

In [example 2](#) of [\[ME\] meoprobit](#), we modeled the tobacco and health knowledge (**thk**) score—coded 1, 2, 3, 4—among students as a function of two treatments (**cc** and **tv**) using a three-level ordered probit model with random effects at the school and class levels.

```
. use http://www.stata-press.com/data/r15/tvsfpors
. meoprobit thk prethk cc##tv || school: || class:
      (output omitted)
```

We obtain predicted probabilities for all four outcomes based on the contribution of both fixed effects and random effects by typing

```
. predict pr*
      (predictions based on fixed effects and posterior means of random effects)
      (option mu assumed)
      (using 7 quadrature points)
```

As the note says, the predicted values are based on the posterior means of random effects. You can use the **modes** option to obtain predictions based on the posterior modes of random effects.

Because we specified a stub name, Stata saved the predicted random effects in variables **pr1** through **pr4**. Here we list the predicted probabilities for the first two classes for school 515:

```
. list class thk pr? if school==515 & (class==515101 | class==515102),
> sepby(class)
```

	class	thk	pr1	pr2	pr3	pr4
1464.	515101	2	.1503512	.2416885	.2828209	.3251394
1465.	515101	2	.3750887	.2958534	.2080368	.121021
1466.	515101	1	.3750887	.2958534	.2080368	.121021
1467.	515101	4	.2886795	.2920168	.2433916	.1759121
1468.	515101	3	.2129906	.2729831	.2696254	.2444009
1469.	515101	3	.2886795	.2920168	.2433916	.1759121
1470.	515102	1	.3318574	.2959802	.2261095	.1460529
1471.	515102	2	.4223251	.2916287	.187929	.0981172
1472.	515102	2	.4223251	.2916287	.187929	.0981172
1473.	515102	2	.4223251	.2916287	.187929	.0981172
1474.	515102	2	.3318574	.2959802	.2261095	.1460529
1475.	515102	1	.4223251	.2916287	.187929	.0981172
1476.	515102	2	.3318574	.2959802	.2261095	.1460529

For each observation, our best guess for the predicted outcome is the one with the highest predicted probability. For example, for the very first observation in the table above, we would choose outcome 4 as the most likely to occur.

We obtain predictions of the posterior means themselves at the school and class levels by typing

```
. predict re_s re_c, reffects
(calculating posterior means of random effects)
(using 7 quadrature points)
```

Here we list the predicted random effects for the first two classes for school 515:

```
. list class re_s re_c if school==515 & (class==515101 | class==515102),
> sepby(class)
```

	class	re_s	re_c
1464.	515101	-.0340769	.0390243
1465.	515101	-.0340769	.0390243
1466.	515101	-.0340769	.0390243
1467.	515101	-.0340769	.0390243
1468.	515101	-.0340769	.0390243
1469.	515101	-.0340769	.0390243
1470.	515102	-.0340769	-.0834322
1471.	515102	-.0340769	-.0834322
1472.	515102	-.0340769	-.0834322
1473.	515102	-.0340769	-.0834322
1474.	515102	-.0340769	-.0834322
1475.	515102	-.0340769	-.0834322
1476.	515102	-.0340769	-.0834322

We can see that the predicted random effects at the school level (**re_s**) are the same for all classes and that the predicted random effects at the class level (**re_c**) are constant within each class.

Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in *[Methods and formulas](#)* of [ME] [meglm postestimation](#).

Also see

[ME] [meoprobit](#) — Multilevel mixed-effects ordered probit regression

[ME] [meglm postestimation](#) — Postestimation tools for meglm

[U] [20 Estimation and postestimation commands](#)

Description
Options
References

Quick start
Remarks and examples
Also see

Menu
Stored results

Syntax
Methods and formulas

Description

`mepoisson` fits mixed-effects models for count responses. The conditional distribution of the response given the random effects is assumed to be Poisson.

`mepoisson` performs optimization with the original metric of variance components. When variance components are near the boundary of the parameter space, you may consider using the `meqrpoisson` command, which provides alternative parameterizations of variance components; see [\[ME\]](#) `meqrpoisson`.

Quick start

Without weights

Two-level Poisson regression of *y* on *x* with random intercepts by `lev2`

```
mepoisson y x || lev2:
```

Add `evvar` measuring exposure

```
mepoisson y x, exposure(evvar) || lev2:
```

As above, but report incidence-rate ratios

```
mepoisson y x, exposure(evvar) || lev2:, irr
```

Add [indicators](#) for levels of categorical variable *a* and random coefficients on *x*

```
mepoisson y x i.a || lev2: x, irr
```

Three-level random-intercept model of *y* on *x* with `lev2` nested within `lev3`

```
mepoisson y x || lev3: || lev2:
```

With weights

Two-level Poisson regression of *y* on *x* with random intercepts by `lev2` and observation-level frequency weights `wvar1`

```
mepoisson y x [fweight=wvar1] || lev2:
```

Two-level random-intercept model from a two-stage sampling design with PSUs identified by `psu` using PSU-level and observation-level sampling weights `wvar2` and `wvar1`, respectively

```
mepoisson y x [pweight=wvar1] || psu:, pweight(wvar2)
```

Add secondary sampling stage with units identified by `ssu` having weights `wvar2` and PSU-level weights `wvar3` for a three-level random-intercept model

```
mepoisson y x [pw=wvar1] || psu:, pw(wvar3) || ssu:, pw(wvar2)
```



```
Same as above, but svyset data first
svyset psu, weight(wvar3) || ssu, weight(wvar2) || _n, weight(wvar1)
svy: mepoisson y x || psu: || ssu:
```

Menu

Statistics > Multilevel mixed-effects models > Poisson regression

Syntax

```
mepoisson depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

- for random coefficients and intercepts
levelvar: [*varlist*] [, *re_options*]
- for random effects among the values of a factor variable
levelvar: R.*varname*

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress the constant term from the fixed-effects equation
<code>exposure(varname_e)</code>	include $\ln(\text{varname}_e)$ in model with coefficient constrained to 1
<code>offset(varname_o)</code>	include <i>varname_o</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(vartype)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(varname)</code>	frequency weights at higher levels
<code>iweight(varname)</code>	importance weights at higher levels
<code>pweight(varname)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>irr</u>	report fixed-effects coefficients as incidence-rate ratios
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<u>mv</u> aghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mc</u> aghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 **Factor variables**.
depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 **Time-series varlists**.
bayes, *by*, and *svy* are allowed; see [U] 11.1.10 **Prefix commands**. For more details, see [BAYES] *bayes*; *mepoisson*.
vce() and weights are not allowed with the *svy* prefix; see [SVY] *svy*.
fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 **weight**. Only one type of weight may be specified.
Weights are not supported under the Laplacian approximation or for crossed models.
startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.
See [U] 20 **Estimation and postestimation commands** for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

exposure(varname_e) specifies a variable that reflects the amount of exposure over which the *depvar* events were observed for each observation; *ln(varname_e)* is included in the fixed-effects portion of the model with the coefficient constrained to be 1.

offset(varname_o) specifies that *varname_o* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(vartype) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, *unstructured*, *fixed(matname)*, or *pattern(matname)*.

covariance(independent) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is *covariance(independent)* unless a crossed random-effects model is fit, in which case the default is *covariance(identity)*.

covariance(exchangeable) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(identity) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(unstructured) allows for all variances and covariances to be distinct. If an equation consists of *p* random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(fixed(matname)) and *covariance(pattern(matname))* covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names

of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if `matname[i, j] = matname[k, l]`.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`constraints(constraints)`, `collinear`; see [R] [estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] [estimation options](#).

`irr` reports estimated fixed-effects coefficients transformed to incidence-rate ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated or stored. `irr` may be specified either at estimation or upon replay.

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `noci`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `mepoisson` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `mepoisson` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid([gridspec])`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

For a general introduction to `me` commands, see [\[ME\] me](#).

Remarks are presented under the following headings:

Introduction

Two-level models

Higher-level models

Introduction

Mixed-effects Poisson regression is Poisson regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

`mepoisson` allows for many levels of random effects. However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = y | \mathbf{x}_{ij}, \mathbf{u}_j) = \exp(-\mu_{ij}) \mu_{ij}^y / y! \quad (1)$$

for $\mu_{ij} = \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)$, $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are counts y_{ij} . The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard Poisson regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

As noted in chapter 13.7 of [Rabe-Hesketh and Skrondal \(2012\)](#), the inclusion of a random intercept causes the marginal variance of y_{ij} to be greater than the marginal mean, provided the variance of the random intercept is not 0. Thus the random intercept in a mixed-effects Poisson model produces overdispersion, a measure of variability above and beyond that allowed by a Poisson process; see [\[R\] nbreg](#) and [\[ME\] menbreg](#).

Below we present examples of mixed-effects Poisson regression; refer to [\[ME\] me](#) and [\[ME\] meglm](#) for additional examples including crossed random-effects models. A two-level Poisson model can also be fit using `xtpoisson` with the `re` option; see [\[XT\] xtpoisson](#). In the absence of random effects, mixed-effects Poisson regression reduces to standard Poisson regression; see [\[R\] poisson](#).

Two-level models

► Example 1: Two-level random-intercept model

[Breslow and Clayton \(1993\)](#) fit a mixed-effects Poisson model to data from a randomized trial of the drug progabide for the treatment of epilepsy.

```
. use http://www.stata-press.com/data/r15/epilepsy
(Epilepsy data; progabide drug treatment)

. describe

Contains data from http://www.stata-press.com/data/r15/epilepsy.dta
   obs:                236                Epilepsy data; progabide drug
                                         treatment
   vars:                8                31 May 2016 14:09
   size:              4,956              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
subject	byte	%9.0g		Subject ID: 1-59
seizures	int	%9.0g		No. of seizures
treat	byte	%9.0g		1: progabide; 0: placebo
visit	float	%9.0g		Dr. visit; coded as (-.3, -.1, .1, .3)
lage	float	%9.0g		log(age), mean-centered
lbas	float	%9.0g		log(0.25*baseline seizures), mean-centered
lbas_trt	float	%9.0g		lbas/treat interaction
v4	byte	%8.0g		Fourth visit indicator

Sorted by: subject

Originally from [Thall and Vail \(1990\)](#), data were collected on 59 subjects (31 on progabide, 28 on placebo). The number of epileptic seizures (`seizures`) was recorded during the two weeks prior to each of four doctor visits (`visit`). The treatment group is identified by the indicator variable `treat`. Data were also collected on the logarithm of age (`lage`) and the logarithm of one-quarter the number of seizures during the eight weeks prior to the study (`lbas`). The variable `lbas_trt` represents the interaction between `lbas` and treatment. `lage`, `lbas`, and `lbas_trt` are mean centered. Because the study originally noted a substantial decrease in seizures prior to the fourth doctor visit, an indicator, `v4`, for the fourth visit was also recorded.

[Breslow and Clayton \(1993\)](#) fit a random-effects Poisson model for the number of observed seizures

$$\log(\mu_{ij}) = \beta_0 + \beta_1 \text{treat}_{ij} + \beta_2 \text{lbas}_{ij} + \beta_3 \text{lbas\_trt}_{ij} + \beta_4 \text{lage}_{ij} + \beta_5 \text{v4}_{ij} + u_j$$

for $j = 1, \dots, 59$ subjects and $i = 1, \dots, 4$ visits. The random effects u_j are assumed to be normally distributed with mean 0 and variance σ_u^2 .

```
. mepoisson seizures treat lbas lbas_trt lage v4 || subject:
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1016.4106
Iteration 1:   log likelihood = -819.20112
Iteration 2:   log likelihood = -817.66006
Iteration 3:   log likelihood = -817.65925
Iteration 4:   log likelihood = -817.65925
Refining starting values:
Grid node 0:   log likelihood = -680.40523
Refining starting values (unscaled likelihoods):
Grid node 0:   log likelihood = -680.40523
Fitting full model:
Iteration 0:   log likelihood = -680.40523   (not concave)
Iteration 1:   log likelihood = -672.95766   (not concave)
Iteration 2:   log likelihood = -667.14039
Iteration 3:   log likelihood = -665.51823
Iteration 4:   log likelihood = -665.29165
Iteration 5:   log likelihood = -665.29067
Iteration 6:   log likelihood = -665.29067

Mixed-effects Poisson regression           Number of obs   =       236
Group variable:      subject                Number of groups =        59
                                           Obs per group:
                                           min =           4
                                           avg =          4.0
                                           max =           4

Integration method: mvaghermite             Integration pts. =         7
                                           Wald chi2(5)     =       121.70
Log likelihood = -665.29067                 Prob > chi2      =       0.0000
```

seizures	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
treat	-.9330306	.4007512	-2.33	0.020	-1.718489	-.1475727
lbas	.8844225	.1312033	6.74	0.000	.6272689	1.141576
lbas_trt	.3382561	.2033021	1.66	0.096	-.0602087	.736721
lage	.4842226	.3471905	1.39	0.163	-.1962582	1.164703
v4	-.1610871	.0545758	-2.95	0.003	-.2680536	-.0541206
_cons	2.154578	.2199928	9.79	0.000	1.7234	2.585756
subject						
var(_cons)	.2528664	.0589844			.1600801	.399434

```
LR test vs. Poisson model: chibar2(01) = 304.74      Prob >= chibar2 = 0.0000
```

The number of seizures before the fourth visit does exhibit a significant drop, and the patients on progabide demonstrate a decrease in frequency of seizures compared with the placebo group. The subject-specific random effects also appear significant: $\hat{\sigma}_u^2 = 0.25$ with standard error 0.06.

Because this is a simple random-intercept model, you can obtain equivalent results by using `xtpoisson` with the `re` and `normal` options.

Example 2: Two-level random-slope model

In their study of PQL, [Breslow and Clayton \(1993\)](#) also fit a model where they dropped the fixed effect on `v4` and replaced it with a random subject-specific linear trend over the four doctor visits. The model they fit is

$$\log(\mu_{ij}) = \beta_0 + \beta_1 \text{treat}_{ij} + \beta_2 \text{lbas}_{ij} + \beta_3 \text{lbas\_trt}_{ij} + \beta_4 \text{lage}_{ij} + \beta_5 \text{visit}_{ij} + u_j + v_j \text{visit}_{ij}$$

where (u_j, v_j) are bivariate normal with 0 mean and variance–covariance matrix

$$\Sigma = \text{Var} \begin{bmatrix} u_j \\ v_j \end{bmatrix} = \begin{bmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{bmatrix}$$

```
. mepoisson seizures treat lbas lbas_trt lage visit || subject: visit,
> covariance(unstructured) intpoints(9) nolog
```

Mixed-effects Poisson regression
Group variable: subject

Number of obs = 236
Number of groups = 59
Obs per group:
min = 4
avg = 4.0
max = 4

Integration method: mvaghermite

Integration pts. = 9
Wald chi2(5) = 115.56
Prob > chi2 = 0.0000

seizures	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
treat	-.9286592	.4021715	-2.31	0.021	-1.716901	-.1404175
lbas	.8849762	.1312535	6.74	0.000	.627724	1.142228
lbas_trt	.3379759	.2044471	1.65	0.098	-.062733	.7386849
lage	.4767192	.3536276	1.35	0.178	-.2163781	1.169817
visit	-.2664098	.1647098	-1.62	0.106	-.5892352	.0564156
_cons	2.099555	.2203749	9.53	0.000	1.667629	2.531482
subject						
var(visit)	.5314803	.229385			.2280928	1.238405
var(_cons)	.2514923	.0587902			.1590534	.3976549
subject						
cov(visit, _cons)	.0028715	.0887037	0.03	0.974	-.1709846	.1767276

LR test vs. Poisson model: chi2(3) = 324.54

Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

In the above, we specified the `covariance(unstructured)` option to allow correlation between u_j and v_j , although on the basis of the above output it probably was not necessary—the default `independent` structure would have sufficed. In the interest of getting more accurate estimates, we also increased the number of quadrature points to nine, although the estimates do not change much when compared with estimates based on the default seven quadrature points.

The essence of the above-fitted model is that after adjusting for other covariates, the log trend in seizures is modeled as a random subject-specific line, with intercept distributed as $N(\beta_0, \sigma_u^2)$ and slope distributed as $N(\beta_5, \sigma_v^2)$. From the above output, $\hat{\beta}_0 = 2.10$, $\hat{\sigma}_u^2 = 0.25$, $\hat{\beta}_5 = -0.27$, and $\hat{\sigma}_v^2 = 0.53$.

You can predict the random effects u_j and v_j by using `predict` after `mepoisson`; see [\[ME\] mepoisson postestimation](#). Better still, you can obtain a predicted number of seizures that takes these random effects into account.



Higher-level models

► Example 3: Three- and four-level random-intercept model

Rabe-Hesketh and Skrondal (2012, exercise 13.7) describe data from the *Atlas of Cancer Mortality in the European Economic Community* (EEC) (Smans, Mair, and Boyle 1993). The data were analyzed in Langford, Bentham, and McDonald (1998) and record the number of deaths among males due to malignant melanoma during 1971–1980.

```
. use http://www.stata-press.com/data/r15/melanoma
(Skin cancer (melanoma) data)

. describe
```

Contains data from http://www.stata-press.com/data/r15/melanoma.dta

obs:	354	Skin cancer (melanoma) data
vars:	6	30 May 2016 17:10
size:	4,956	(_dta has notes)

variable name	storage type	display format	value label	variable label
nation	byte	%11.0g	n	Nation ID
region	byte	%9.0g		Region ID: EEC level-I areas
county	int	%9.0g		County ID: EEC level-II/level-III areas
deaths	int	%9.0g		No. deaths during 1971-1980
expected	float	%9.0g		No. expected deaths
uv	float	%9.0g		UV dose, mean-centered

Sorted by:

Nine European nations (variable `nation`) are represented, and data were collected over geographical regions defined by EEC statistical services as level I areas (variable `region`), with deaths being recorded for each of 354 counties, which are level II or level III EEC-defined areas (variable `county`, which identifies the observations). Counties are nested within regions, and regions are nested within nations.

The variable `deaths` records the number of deaths for each county, and `expected` records the expected number of deaths (the exposure) on the basis of crude rates for the combined countries. Finally, the variable `uv` is a measure of exposure to ultraviolet (UV) radiation.

In modeling the number of deaths, one possibility is to include dummy variables for the nine nations as fixed effects. Another is to treat these as random effects and fit the three-level random-intercept Poisson model,

$$\log(\mu_{ijk}) = \log(\text{expected}_{ijk}) + \beta_0 + \beta_1 \text{uv}_{ijk} + u_k + v_{jk}$$

for nation k , region j , and county i . The model includes an exposure term for expected deaths.

```
. mepoisson deaths uv, exposure(expected) || nation: || region:

Fitting fixed-effects model:
Iteration 0:   log likelihood = -2136.5847
Iteration 1:   log likelihood = -1723.8955
Iteration 2:   log likelihood = -1723.7727
Iteration 3:   log likelihood = -1723.7727

Refining starting values:
Grid node 0:   log likelihood = -1166.6536

Refining starting values (unscaled likelihoods):
Grid node 0:   log likelihood = -1166.6536

Fitting full model:
Iteration 0:   log likelihood = -1166.6536   (not concave)
Iteration 1:   log likelihood = -1152.2741   (not concave)
Iteration 2:   log likelihood = -1146.3094   (not concave)
Iteration 3:   log likelihood = -1119.8479   (not concave)
Iteration 4:   log likelihood = -1108.0129   (not concave)
Iteration 5:   log likelihood = -1098.8067
Iteration 6:   log likelihood = -1095.7563
Iteration 7:   log likelihood = -1095.3164
Iteration 8:   log likelihood = -1095.31
Iteration 9:   log likelihood = -1095.31

Mixed-effects Poisson regression                               Number of obs   =           354

+-----+-----+-----+-----+-----+
| Group Variable | No. of | Observations per Group |
|                 | Groups | Minimum  Average  Maximum |
+-----+-----+-----+-----+-----+
| nation         |      9 |      3    39.3        95 |
| region         |     78 |      1     4.5        13 |
+-----+-----+-----+-----+-----+

Integration method: mvaghermite                               Integration pts. =           7

Log likelihood = -1095.31                                     Wald chi2(1)      =           6.12
                                                            Prob > chi2       =           0.0134

+-----+-----+-----+-----+-----+-----+
| deaths | Coef.  Std. Err.  z    P>|z|    [95% Conf. Interval] |
+-----+-----+-----+-----+-----+-----+
| uv      | -.0282041 .0113998  -2.47  0.013  -.0505473  -.0058608 |
| _cons   | -.0639672 .1335515  -0.48  0.632  -.3257234  .1977789 |
| ln(expected) |      1 (exposure) |
+-----+-----+-----+-----+-----+
| nation  |      |      |      |      |      |      |
| var(_cons) | .1371732 .0723303 |      |      | .048802  .3855676 |
+-----+-----+-----+-----+-----+
| nation>  |      |      |      |      |      |      |
| region   |      |      |      |      |      |      |
| var(_cons) | .0483483 .0109079 |      |      | .0310699  .0752353 |
+-----+-----+-----+-----+-----+

LR test vs. Poisson model: chi2(2) = 1256.93                Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

By including an exposure variable that is an expected rate, we are in effect specifying a linear model for the log of the standardized mortality ratio, the ratio of observed deaths to expected deaths that is based on a reference population. Here the reference population is all nine nations.

Looking at the estimated variance components, we can see there is more unobserved variability between nations than between regions within each nation. This may be due to, for example, country-specific informational campaigns on the risks of sun exposure.

We now add a random intercept for counties nested within regions, making this a four-level model. Because counties also identify the observations, the corresponding variance component can be interpreted as a measure of overdispersion, variability above and beyond that allowed by a Poisson process; see [R] [nbreg](#) and [ME] [menbreg](#).

```
. mepoisson deaths uv, exposure(expected) || nation: || region: || county:,
> intmethod(mcaghermite)
```

Fitting fixed-effects model:

```
Iteration 0: log likelihood = -2136.5847
Iteration 1: log likelihood = -1723.8955
Iteration 2: log likelihood = -1723.7727
Iteration 3: log likelihood = -1723.7727
```

Refining starting values:

```
Grid node 0:    log likelihood = -1379.3466
```

Refining starting values (unscaled likelihoods):

```
Grid node 0:    log likelihood = -1379.3466
```

Fitting full model:

```
Iteration 0: log likelihood = -1379.3466 (not concave)
Iteration 1: log likelihood = -1310.4947 (not concave)
Iteration 2: log likelihood = -1245.534 (not concave)
Iteration 3: log likelihood = -1218.5474 (not concave)
Iteration 4: log likelihood = -1207.881 (not concave)
Iteration 5: log likelihood = -1122.0585 (not concave)
Iteration 6: log likelihood = -1092.4049
Iteration 7: log likelihood = -1088.0486
Iteration 8: log likelihood = -1086.7175
Iteration 9: log likelihood = -1086.6756
Iteration 10: log likelihood = -1086.6754
Iteration 11: log likelihood = -1086.6754
```

Mixed-effects Poisson regression Number of obs = 354

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
nation	9	3	39.3	95
region	78	1	4.5	13
county	354	1	1.0	1

Integration method: mcaghermite

Integration pts. = 7

Log likelihood = -1086.6754

Wald chi2(1) = 8.62

Prob > chi2 = 0.0033

deaths	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
uv	-.0334702	.0113968	-2.94	0.003	-.0558075	-.0111329
_cons	-.0864583	.1299275	-0.67	0.506	-.3411115	.168195
ln(expected)	1 (exposure)					
nation var(_cons)	.1288627	.0681643			.0456949	.3634011
nation> region var(_cons)	.0406279	.0105154			.0244633	.0674735
nation> region> county var(_cons)	.0146672	.0050979			.0074215	.0289867

LR test vs. Poisson model: chi2(3) = 1274.19

Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

In the above, we used `intmethod(mcaghermite)`, which is not only faster but also produces estimates that closely agree with those obtained with the default `mvaghermite` integration method.



Stored results

`mepoisson` stores the following in `e()`:

Scalars	
<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise
Macros	
<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>mepoisson</code>

<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	poisson
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	log
<code>e(family)</code>	poisson
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	offset
<code>e(exposure)</code>	exposure variable
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	b V
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by margins
<code>e(marginswtype)</code>	weight type for margins
<code>e(marginswexp)</code>	weight expression for margins
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

`mepoisson` is a convenience command for `meglm` with a log link and an poisson family; see [ME] [meglm](#).

In a two-level Poisson model, for cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$f(\mathbf{y}_j|\mathbf{u}_j) = \prod_{i=1}^{n_j} [\{\exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\}^{y_{ij}} \exp\{-\exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\} / y_{ij}!]$$

$$= \exp \left[\sum_{i=1}^{n_j} \{y_{ij}(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) - \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) - \log(y_{ij}!)\} \right]$$

Defining $c(\mathbf{y}_j) = \sum_{i=1}^{n_j} \log(y_{ij}!)$, where $c(\mathbf{y}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j|\mathbf{u}_j) = \exp \{ \mathbf{y}_j' (\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - \mathbf{1}' \exp(\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - c(\mathbf{y}_j) \}$$

where $\mathbf{X}_j$ is formed by stacking the row vectors $\mathbf{x}_{ij}$ and $\mathbf{Z}_j$ is formed by stacking the row vectors $\mathbf{z}_{ij}$. We extend the definition of $\exp(\cdot)$ to be a vector function where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{y}_j|\mathbf{u}_j) \exp(-\mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp\{-c(\mathbf{y}_j)\} (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int \exp\{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \quad (2)$$

where

$$h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) = \mathbf{y}_j' (\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - \mathbf{1}' \exp(\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

mepoisson supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Breslow, N. E., and D. G. Clayton. 1993. Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* 88: 9–25.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- Langford, I. H., G. Bentham, and A. McDonald. 1998. Multi-level modelling of geographically aggregated health data: A case study on malignant melanoma mortality and UV exposure in the European community. *Statistics in Medicine* 17: 41–57.
- Leyland, A. H., and H. Goldstein, ed. 2001. *Multilevel Modelling of Health Statistics*. New York: Wiley.

- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Rabe-Hesketh, S., and A. Skrondal. 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Rabe-Hesketh, S., A. Skrondal, and A. Pickles. 2005. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics* 128: 301–323.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.
- Skrondal, A., and S. Rabe-Hesketh. 2004. [Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models](#). Boca Raton, FL: Chapman & Hall/CRC.
- Smans, M., C. S. Mair, and P. Boyle. 1993. *Atlas of Cancer Mortality in the European Economic Community*. Lyon, France: IARC Scientific Publications.
- Thall, P. F., and S. C. Vail. 1990. Some covariance models for longitudinal count data with overdispersion. *Biometrics* 46: 657–671.

Also see

- [ME] [mepoisson postestimation](#) — Postestimation tools for mepoisson
- [ME] [menbreg](#) — Multilevel mixed-effects negative binomial regression
- [ME] [meqrpoisson](#) — Multilevel mixed-effects Poisson regression (QR decomposition)
- [ME] [me](#) — Introduction to multilevel mixed-effects models
- [BAYES] [bayes: mepoisson](#) — Bayesian multilevel Poisson regression
- [SEM] [intro 5](#) — Tour of models ([Multilevel mixed-effects models](#))
- [SVY] [svy estimation](#) — Estimation commands for survey data
- [XT] [xtpoisson](#) — Fixed-effects, random-effects, and population-averaged Poisson models
- [U] [20 Estimation and postestimation commands](#)

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation command is of special interest after `mepoisson`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
density	predicted density function
distribution	predicted distribution function
pearson	Pearson residuals
deviance	deviance residuals
anscombe	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

options	Description
Main	
<code>conditional(ctype)</code>	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<code>marginal</code>	compute <i>statistic</i> marginally with respect to the random effects
<code>nooffset</code>	make calculation ignoring offset or exposure
Integration	
<code>int_options</code>	integration options
pearson, deviance, anscombe may not be combined with marginal.	
ctype	Description
<code>ebmeans</code>	empirical Bayes means of random effects; the default
<code>ebmodes</code>	empirical Bayes modes of random effects
<code>fixedonly</code>	prediction for the fixed portion of the model only
re_options	Description
Main	
<code>ebmeans</code>	use empirical Bayes means of random effects; the default
<code>ebmodes</code>	use empirical Bayes modes of random effects
<code>reses(stub* newvarlist)</code>	calculate standard errors of empirical Bayes estimates
Integration	
<code>int_options</code>	integration options
int_options	Description
<code>intpoints(#)</code>	use # quadrature points to compute marginal predictions and empirical Bayes means
<code>iterate(#)</code>	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<code>tolerance(#)</code>	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

mu, the default, calculates the predicted mean, that is, the predicted number of events.
eta, xb, stdp, density, distribution, pearson, deviance, anscombe, scores, conditional(), marginal, and nooffset; see [ME] [meglm postestimation](#).
reffects, ebmeans, ebmodes, and reses(); see [ME] [meglm postestimation](#).

Integration

intpoints(), iterate(), and tolerance(); see [ME] [meglm postestimation](#).

margins

Description for margins

margins estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
<u>density</u>	not allowed with margins
<u>distribution</u>	not allowed with margins
<u>pearson</u>	not allowed with margins
<u>deviance</u>	not allowed with margins
<u>anscombe</u>	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options conditional(ebmeans) and conditional(ebmodes) are not allowed with margins.
Option marginal is assumed where applicable if conditional(fixedonly) is not specified.

Statistics not allowed with margins are functions of stochastic quantities other than e(b).
For the full syntax, see [R] margins.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects Poisson model with mepoisson. For the most part, calculation centers around obtaining estimates of the subject/group-specific random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation.

Here we show a short example of predicted counts and predicted random effects; refer to [ME] meglm postestimation for additional examples applicable to mixed-effects generalized linear models.

➤ Example 1: Predicting counts and random effects

In [example 2](#) of [\[ME\] mepoisson](#), we modeled the number of observed epileptic seizures as a function of treatment with the drug progabide and other covariates,

$$\log(\mu_{ij}) = \beta_0 + \beta_1 \text{treat}_{ij} + \beta_2 \text{lbas}_{ij} + \beta_3 \text{lbas\_trt}_{ij} + \beta_4 \text{lage}_{ij} + \beta_5 \text{visit}_{ij} + u_j + v_j \text{visit}_{ij}$$

where (u_j, v_j) are bivariate normal with 0 mean and variance–covariance matrix

$$\Sigma = \text{Var} \begin{bmatrix} u_j \\ v_j \end{bmatrix} = \begin{bmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{bmatrix}$$

```
. use http://www.stata-press.com/data/r15/epilepsy
(Epilepsy data; progabide drug treatment)
. mepoisson seizures treat lbas lbas_trt lage visit || subject: visit,
> cov(unstructured) intpoints(9)
(iteration log omitted)
```

Mixed-effects Poisson regression	Number of obs	=	236
Group variable: subject	Number of groups	=	59
	Obs per group:		
	min	=	4
	avg	=	4.0
	max	=	4
Integration method: mvaghermite	Integration pts.	=	9
	Wald chi2(5)	=	115.56
Log likelihood = -655.68103	Prob > chi2	=	0.0000

seizures	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
treat	-.9286592	.4021715	-2.31	0.021	-1.716901	-.1404175
lbas	.8849762	.1312535	6.74	0.000	.627724	1.142228
lbas_trt	.3379759	.2044471	1.65	0.098	-.062733	.7386849
lage	.4767192	.3536276	1.35	0.178	-.2163781	1.169817
visit	-.2664098	.1647098	-1.62	0.106	-.5892352	.0564156
_cons	2.099555	.2203749	9.53	0.000	1.667629	2.531482

subject						
var(visit)	.5314803	.229385			.2280928	1.238405
var(_cons)	.2514923	.0587902			.1590534	.3976549

subject						
cov(visit,						
_cons)	.0028715	.0887037	0.03	0.974	-.1709846	.1767276

LR test vs. Poisson model: chi2(3) = 324.54 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

The purpose of this model was to allow subject-specific linear log trends over each subject’s four doctor visits, after adjusting for the other covariates. The intercepts of these lines are distributed $N(\beta_0, \sigma_u^2)$, and the slopes are distributed $N(\beta_5, \sigma_v^2)$, based on the fixed effects and assumed distribution of the random effects.

We can use `predict` to obtain estimates of the random effects u_j and v_j and combine these with our estimates of β_0 and β_5 to obtain the intercepts and slopes of the linear log trends.

```
. predict re_visit re_cons, reffects
(calculting posterior means of random effects)
(using 9 quadrature points)
. generate b1 = _b[visit] + re_visit
. generate b0 = _b[_cons] + re_cons
. by subject, sort: generate tolist = _n==1
. list subject treat b1 b0 if tolist & (subject <=5 | subject >=55)
```

	subject	treat	b1	b0
1.	1	0	-.428854	2.13539
5.	2	0	-.2731013	2.149744
9.	3	0	.0022089	2.417506
13.	4	0	-.3197094	2.238224
17.	5	0	.6082718	2.110739
217.	55	1	-.2308834	2.282539
221.	56	1	.2912798	3.19678
225.	57	1	-.4828764	1.423153
229.	58	1	-.2519466	1.131373
233.	59	1	-.1269573	2.171541

We list these slopes (b1) and intercepts (b0) for five control subjects and five subjects on the treatment.

```
. count if tolist & treat
31
. count if tolist & treat & b1 < 0
25
. count if tolist & !treat
28
. count if tolist & !treat & b1 < 0
20
```

We also find that 25 of the 31 subjects taking progabide were estimated to have a downward trend in seizures over their four doctor visits, compared with 20 of the 28 control subjects.

We also obtain predictions for number of seizures, and unless we specify the conditional(fixedonly) option, these predictions will incorporate the estimated subject-specific random effects.

```
. predict n
(predictions based on fixed effects and posterior means of random effects)
(option mu assumed)
(using 9 quadrature points)
```

```
. list subject treat visit seizures n if subject <= 2 | subject >= 58, sep(0)
```

	subject	treat	visit	seizures	n
1.	1	0	-.3	5	3.775774
2.	1	0	-.1	3	3.465422
3.	1	0	.1	3	3.18058
4.	1	0	.3	3	2.919151
5.	2	0	-.3	3	3.598805
6.	2	0	-.1	5	3.40751
7.	2	0	.1	3	3.226382
8.	2	0	.3	3	3.054883
229.	58	1	-.3	0	.9611137
230.	58	1	-.1	0	.9138838
231.	58	1	.1	0	.8689747
232.	58	1	.3	0	.8262726
233.	59	1	-.3	1	2.40652
234.	59	1	-.1	4	2.346184
235.	59	1	.1	3	2.287361
236.	59	1	.3	2	2.230013



Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [\[ME\] meglm postestimation](#).

Also see

- [\[ME\] mepoisson](#) — Multilevel mixed-effects Poisson regression
- [\[ME\] meglm postestimation](#) — Postestimation tools for meglm
- [\[U\] 20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

meprobit fits mixed-effects models for binary or binomial responses. The conditional distribution of the response given the random effects is assumed to be Bernoulli, with success probability determined by the standard normal cumulative distribution function.

Quick start

Two-level probit model of y and covariate x and random intercepts by lev2

```
meprobit y x || lev2:
```

Add random coefficients for x

```
meprobit y x || lev2: x
```

As above, but specify that y records the number of successes from 10 trials

```
meprobit y x || lev2: x, binomial(10)
```

As above, but with the number of trials stored in variable n

```
meprobit y x || lev2: x, binomial(n)
```

Three-level random-intercept model of y and covariate x with lev2 nested within lev3

```
meprobit y x || lev3: || lev2:
```

Two-way crossed random effects by factors a and b

```
meprobit y x || _all:R.a || b:
```

Menu

Statistics > Multilevel mixed-effects models > Probit regression

Syntax

```
meprobit depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1
<code>asis</code>	retain perfect predictor variables

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>binomial</u> (<i>varname</i> #)	set binomial trials if data are in binomial form
<u>constraints</u> (<i>constraints</i>)	apply specified linear constraints
<u>collinear</u>	keep collinear variables
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>nocnsreport</u>	do not display constraints
<u>notable</u>	suppress coefficient table
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intmethod</u> (<i>intmethod</i>)	integration method
<u>intpoints</u> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>startvalues</u> (<i>svmethod</i>)	method for obtaining starting values
<u>startgrid</u> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>noestimate</u>	do not fit the model; show starting values instead
<u>dnnumerical</u>	use numerical derivative techniques
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated
<u>fixed</u> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>pattern</u> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<u>mv</u> aghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mc</u> aghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: meprobit.

vce() and weights are not allowed with the *svy* prefix; see [SVY] svy.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 weight. Only one type of weight may be specified. Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

asis forces retention of perfect predictor variables and their associated, perfectly predicted observations and may produce instabilities in maximization; see [R] **probit**.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: **independent**, **exchangeable**, **identity**, **unstructured**, **fixed(*matname*)**, or **pattern(*matname*)**.

covariance(independent) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is **covariance(independent)** unless a crossed random-effects model is fit, in which case the default is **covariance(identity)**.

covariance(exchangeable) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(identity) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(unstructured) allows for all variances and covariances to be distinct. If an equation consists of *p* random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(fixed(*matname*)) and **covariance(pattern(*matname*))** covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a **fixed(*matname*)** covariance structure, (co)variance (*i*, *j*) is constrained to equal the

value specified in the i, j th entry of *matname*. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if `matname[i, j] = matname[k, l]`.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`binomial(varname | #)` specifies that the data are in binomial form; that is, *devar* records the number of successes from a series of binomial trials. This number of trials is given either as *varname*, which allows this number to vary over the observations, or as the constant `#`. If `binomial()` is not specified (the default), *devar* is treated as Bernoulli, with any nonzero, nonmissing values indicating positive responses.

`constraints(constraints)`, `collinear`; see [R] [estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`, `nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `noci`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meprobit` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `meprobit` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid([gridspec])`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects probit regression is probit regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

`meprobit` allows for many levels of random effects. However, for simplicity, we here consider the two-level model, where for a series of M independent clusters, and conditional on a set of fixed effects $\mathbf{x}_{ij}$ and a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = 1 | \mathbf{x}_{ij}, \mathbf{u}_j) = H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are the binary-valued y_{ij} , and we follow the standard Stata convention of treating $y_{ij} = 1$ if `depvarij` $\neq 0$ and treating $y_{ij} = 0$ otherwise. The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard probit regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$. For notational convenience here and throughout this manual entry, we suppress the dependence of y_{ij} on $\mathbf{x}_{ij}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$, so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Finally, because this is probit regression, $H(\cdot)$ is the standard normal cumulative distribution function, which maps the linear predictor to the probability of a success ($y_{ij} = 1$) with $H(v) = \Phi(v)$.

Model (1) may also be stated in terms of a latent linear response, where only $y_{ij} = I(y_{ij}^* > 0)$ is observed for the latent

$$y_{ij}^* = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

The errors ϵ_{ij} are distributed as a standard normal with mean 0 and variance 1 and are independent of $\mathbf{u}_j$.

Below we present two short examples of mixed-effects probit regression; refer to [ME] `me` and [ME] `meglm` for examples of other random-effects models. A two-level probit model can also be fit using `xtprobit` with the `re` option; see [XT] `xtprobit`. In the absence of random effects, mixed-effects probit regression reduces to standard probit regression; see [R] `probit`.

► Example 1: Two-level random-intercept model

Ng et al. (2006) analyzed a subsample of data from the 1989 Bangladesh fertility survey (Huq and Cleland 1990), which polled 1,934 Bangladeshi women on their use of contraception. The women sampled were from 60 districts, identified by the variable `district`. Each district contained either urban or rural areas (variable `urban`) or both. The variable `c_use` is the binary response, with a value of 1 indicating contraceptive use. Other covariates include mean-centered `age` and three indicator variables recording number of children.

```
. use http://www.stata-press.com/data/r15/bangladesh
(Bangladesh Fertility Survey, 1989)

. meprobit c_use urban age child* || district:

Fitting fixed-effects model:
Iteration 0:   log likelihood = -1228.8313
Iteration 1:   log likelihood = -1228.2466
Iteration 2:   log likelihood = -1228.2466

Refining starting values:
Grid node 0:   log likelihood = -1237.3973

Fitting full model:
Iteration 0:   log likelihood = -1237.3973   (not concave)
Iteration 1:   log likelihood = -1221.2111   (not concave)
Iteration 2:   log likelihood = -1207.4451
Iteration 3:   log likelihood = -1206.7002
Iteration 4:   log likelihood = -1206.5346
Iteration 5:   log likelihood = -1206.5336
Iteration 6:   log likelihood = -1206.5336

Mixed-effects probit regression
Group variable:      district

Number of obs      =      1,934
Number of groups   =         60
Obs per group:
    min =          2
    avg =        32.2
    max =        118

Integration method: mvaghermite           Integration pts. =         7
Wald chi2(5)       =        115.36
Prob > chi2        =         0.0000

Log likelihood = -1206.5336
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	.4490191	.0727176	6.17	0.000	.3064953	.5915429
age	-.0162203	.0048005	-3.38	0.001	-.0256291	-.0068114
child1	.674377	.0947829	7.11	0.000	.488606	.8601481
child2	.8281581	.1048136	7.90	0.000	.6227272	1.033589
child3	.8137876	.1073951	7.58	0.000	.6032972	1.024278
_cons	-1.02799	.0870307	-11.81	0.000	-1.198567	-.8574132
district var(_cons)	.0798719	.026886			.0412921	.1544972

LR test vs. probit model: chibar2(01) = 43.43 Prob >= chibar2 = 0.0000

Probit regression coefficients are most commonly interpreted in terms of partial effects, as we demonstrate in [example 1](#) of [\[ME\] meprobit postestimation](#). For now, we only note that urban women and women with more children are more likely to use contraceptives and that contraceptive use decreases with age. The estimated variance of the random intercept at the district level, $\hat{\sigma}^2$, is 0.08 with standard error 0.03. The reported likelihood-ratio test shows that there is enough variability between districts to favor a mixed-effects probit regression over an ordinary probit regression; see [Distribution theory for likelihood-ratio test](#) in [\[ME\] me](#) for a discussion of likelihood-ratio testing of variance components.

➤ Example 2: Three-level random-intercept model

Rabe-Hesketh, Touloupoulou, and Murray (2001) analyzed data from a study that measured the cognitive ability of patients with schizophrenia compared with their relatives and control subjects. Cognitive ability was measured as the successful completion of the “Tower of London”, a computerized task, measured at three levels of difficulty. For all but one of the 226 subjects, there were three measurements (one for each difficulty level). Because patients’ relatives were also tested, a family identifier, `family`, was also recorded.

We fit a probit model with response `dtlm`, the indicator of cognitive function, and with covariates `difficulty` and a set of indicator variables for `group`, with the controls (`group==1`) being the base category. We also allow for random effects due to families and due to subjects within families. The first is a random intercept (constant only) at the `family` level, and the second is a random intercept at the `subject` level. The order in which these are specified (from left to right) is significant—`meprobit` assumes that `subject` is nested within `family`. The equations are separated by `||`.


```
. use http://www.stata-press.com/data/r15/towerlondon
(Tower of London data)
. meprobit dtlm difficulty i.group || family: || subject:

Fitting fixed-effects model:
Iteration 0:  log likelihood = -317.11238
Iteration 1:  log likelihood = -314.50535
Iteration 2:  log likelihood = -314.50121
Iteration 3:  log likelihood = -314.50121

Refining starting values:
Grid node 0:  log likelihood = -326.18533

Fitting full model:
Iteration 0:  log likelihood = -326.18533 (not concave)
Iteration 1:  log likelihood = -313.16256 (not concave)
Iteration 2:  log likelihood = -308.47528
Iteration 3:  log likelihood = -305.02228
Iteration 4:  log likelihood = -304.88972
Iteration 5:  log likelihood = -304.88845
Iteration 6:  log likelihood = -304.88845

Mixed-effects probit regression                                Number of obs      =          677

+-----+-----+-----+-----+-----+-----+
| Group Variable | No. of | Observations per Group |
|                 | Groups | Minimum   Average   Maximum |
+-----+-----+-----+-----+-----+
|      family    |     118 |           2           5.7           27 |
|      subject   |     226 |           2           3.0           3  |
+-----+-----+-----+-----+-----+

Integration method: mvaghermite                                Integration pts. =           7
Wald chi2(3) =          83.28
Log likelihood = -304.88845                                     Prob > chi2      =          0.0000

+-----+-----+-----+-----+-----+-----+
|      dtlm      | Coef.  | Std. Err. | z    | P>|z| | [95% Conf. Interval] |
+-----+-----+-----+-----+-----+-----+
| difficulty     | - .9329891 | .1037376 | -8.99 | 0.000 | -1.136311  -.7296672 |
| group         |           |           |       |       |           |           |
|   2           | - .1632243 | .204265  | -0.80 | 0.424 | -.5635763   .2371276 |
|   3           | - .6220196 | .228063  | -2.73 | 0.006 | -1.069015  -.1750244 |
| _cons         | - .8405154 | .1597223 | -5.26 | 0.000 | -1.153565  -.5274654 |
+-----+-----+-----+-----+-----+-----+
| family        |           |           |       |       |           |           |
|   var(_cons)  | .2120948   | .1736281 |       |       | .0426292    1.055244 |
+-----+-----+-----+-----+-----+-----+
| family>       |           |           |       |       |           |           |
|   subject     |           |           |       |       |           |           |
|   var(_cons)  | .3559141   | .219331  |       |       | .106364     1.190956 |
+-----+-----+-----+-----+-----+-----+

LR test vs. probit model: chi2(2) = 19.23                      Prob > chi2 = 0.0001
Note: LR test is conservative and provided only for reference.
```

We see that we have 226 subjects from 118 families. After adjusting for the random-effects structure, the probability of successful completion of the Tower of London decreases dramatically as the level of difficulty increases. Also, people with schizophrenia (group==3) tended not to perform as well as the control subjects.



The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by ||.

Stored results

meprobit stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>meprobit</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	<code>probit</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>probit</code>
<code>e(family)</code>	<code>bernoulli</code> or <code>binomial</code>
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	<code>offset</code>
<code>e(binomial)</code>	binomial number of trials
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vctype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	<code>max</code> or <code>min</code> ; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of <code>ml</code> method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>

<code>e(marginswexp)</code>	weight expression for margins
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>
Matrices	
<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance
Functions	
<code>e(sample)</code>	marks estimation sample

Methods and formulas

`meprobit` is a convenience command for `meglm` with a `probit` link and a `bernoulli` or `binomial` family; see [ME] [meglm](#).

Model (1) assumes Bernoulli data, a special case of the binomial. Because binomial data are also supported by `meprobit` (option `binomial()`), the methods presented below are in terms of the more general binomial mixed-effects model.

For a two-level binomial model, consider the response y_{ij} as the number of successes from a series of r_{ij} Bernoulli trials (replications). For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$\begin{aligned} f(\mathbf{y}_j | \mathbf{u}_j) &= \prod_{i=1}^{n_j} \left[\binom{r_{ij}}{y_{ij}} \{ \Phi(\boldsymbol{\eta}_{ij}) \}^{y_{ij}} \{ 1 - \Phi(\boldsymbol{\eta}_{ij}) \}^{r_{ij} - y_{ij}} \right] \\ &= \exp \left(\sum_{i=1}^{n_j} \left[y_{ij} \log \{ \Phi(\boldsymbol{\eta}_{ij}) \} - (r_{ij} - y_{ij}) \log \{ \Phi(-\boldsymbol{\eta}_{ij}) \} + \log \binom{r_{ij}}{y_{ij}} \right] \right) \end{aligned}$$

for $\boldsymbol{\eta}_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \text{offset}_{ij}$.

Defining $\mathbf{r}_j = (r_{j1}, \dots, r_{jn_j})'$ and

$$c(\mathbf{y}_j, \mathbf{r}_j) = \sum_{i=1}^{n_j} \log \binom{r_{ij}}{y_{ij}}$$

where $c(\mathbf{y}_j, \mathbf{r}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j | \mathbf{u}_j) = \exp \left[\mathbf{y}_j' \log \{ \Phi(\boldsymbol{\eta}_j) \} - (\mathbf{r}_j - \mathbf{y}_j)' \log \{ \Phi(-\boldsymbol{\eta}_j) \} + c(\mathbf{y}_j, \mathbf{r}_j) \right]$$

where $\boldsymbol{\eta}_j$ is formed by stacking the row vectors $\boldsymbol{\eta}_{ij}$. We extend the definitions of $\Phi(\cdot)$, $\log(\cdot)$, and $\exp(\cdot)$ to be vector functions where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ , the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned}\mathcal{L}_j(\beta, \Sigma) &= (2\pi)^{-q/2} |\Sigma|^{-1/2} \int f(\mathbf{y}_j | \mathbf{u}_j) \exp(-\mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp\{c(\mathbf{y}_j, \mathbf{r}_j)\} (2\pi)^{-q/2} |\Sigma|^{-1/2} \int \exp\{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j\end{aligned}\quad (2)$$

where

$$h(\beta, \Sigma, \mathbf{u}_j) = \mathbf{y}_j' \log\{\Phi(\boldsymbol{\eta}_j)\} - (\mathbf{r}_j - \mathbf{y}_j)' \log\{\Phi(-\boldsymbol{\eta}_j)\} - \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

meprobit supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

References

- Huq, N. M., and J. Cleland. 1990. *Bangladesh Fertility Survey 1989 (Main Report)*. National Institute of Population Research and Training.
- Ng, E. S.-W., J. R. Carpenter, H. Goldstein, and J. Rasbash. 2006. Estimation in generalised linear mixed models with binary outcomes by simulated maximum likelihood. *Statistical Modelling* 6: 23–42.
- Rabe-Hesketh, S., T. Touloupoulou, and R. M. Murray. 2001. Multilevel modeling of cognitive function in schizophrenic patients and their first degree relatives. *Multivariate Behavioral Research* 36: 279–298.

Also see

- [ME] **meprobit postestimation** — Postestimation tools for meprobit
- [ME] **mecloglog** — Multilevel mixed-effects complementary log-log regression
- [ME] **melogit** — Multilevel mixed-effects logistic regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: meprobit** — Bayesian multilevel probit regression
- [SEM] **intro 5** — Tour of models ([Multilevel mixed-effects models](#))
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtprobit** — Random-effects and population-averaged probit models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation commands are of special interest after `meprobit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat icc</code>	estimate intraclass correlations
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; density and distribution functions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
density	predicted density function
distribution	predicted distribution function
pearson	Pearson residuals
deviance	deviance residuals
anscombe	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description

Main	
<code>conditional(<i>ctype</i>)</code>	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<code>marginal</code>	compute <i>statistic</i> marginally with respect to the random effects
<code>nooffset</code>	make calculation ignoring offset or exposure

Integration	
<code>int_options</code>	integration options
pearson, deviance, anscombe may not be combined with <code>marginal</code> .	

<i>ctype</i>	Description
<code>ebmeans</code>	empirical Bayes means of random effects; the default
<code>ebmodes</code>	empirical Bayes modes of random effects
<code>fixedonly</code>	prediction for the fixed portion of the model only

<i>re_options</i>	Description

Main	
<code>ebmeans</code>	use empirical Bayes means of random effects; the default
<code>ebmodes</code>	use empirical Bayes modes of random effects
<code>reses(<i>stub*</i> <i>newvarlist</i>)</code>	calculate standard errors of empirical Bayes estimates

Integration	
<code>int_options</code>	integration options

<i>int_options</i>	Description
<code>intpoints(#)</code>	use # quadrature points to compute marginal predictions and empirical Bayes means
<code>iterate(#)</code>	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<code>tolerance(#)</code>	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main	
mu, the default, calculates the predicted mean, that is, the probability of a positive outcome.	
eta, xb, stdp, density, distribution, pearson, deviance, anscombe, scores, conditional(), marginal, and nooffset; see [ME] meglm postestimation .	
reffects, ebmeans, ebmodes, and reses(); see [ME] meglm postestimation .	

Integration	
intpoints(), iterate(), and tolerance(); see [ME] meglm postestimation .	

margins

Description for margins

margins estimates margins of response for mean responses and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mu	mean response; the default
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	not allowed with margins
<u>density</u>	not allowed with margins
<u>distribution</u>	not allowed with margins
<u>pearson</u>	not allowed with margins
<u>deviance</u>	not allowed with margins
<u>anscombe</u>	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options conditional(ebmeans) and conditional(ebmodes) are not allowed with margins.
Option marginal is assumed where applicable if conditional(fixedonly) is not specified.

Statistics not allowed with margins are functions of stochastic quantities other than e(b).
For the full syntax, see [R] margins.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects probit model using `meprobit`. Here we show a short example of predicted probabilities and predicted random effects; refer to [\[ME\] meglm postestimation](#) for additional examples.

➤ Example 1: Predicting random effects and estimating intraclass correlations

In [example 2](#) of [\[ME\] meprobit](#), we analyzed the cognitive ability (`dtlm`) of patients with schizophrenia compared with their relatives and control subjects, by using a three-level probit model with random effects at the family and subject levels. Cognitive ability was measured as the successful completion of the “Tower of London”, a computerized task, measured at three levels of difficulty.

```
. use http://www.stata-press.com/data/r15/towerlondon
(Tower of London data)

. meprobit dtlm difficulty i.group || family: || subject:
(output omitted)
```

We obtain predicted probabilities based on the contribution of both fixed effects and random effects by typing

```
. predict pr
(predictions based on fixed effects and posterior means of random effects)
(option mu assumed)
(using 7 quadrature points)
```

As the note says, the predicted values are based on the posterior means of random effects. You can use the `modes` option to obtain predictions based on the posterior modes of random effects.

We obtain predictions of the posterior means themselves by typing

```
. predict re*, reffects
(calculating posterior means of random effects)
(using 7 quadrature points)
```

Because we have one random effect at the family level and another random effect at the subject level, Stata saved the predicted posterior means in the variables `re1` and `re2`, respectively. If you are not sure which prediction corresponds to which level, you can use the `describe` command to show the variable labels.

Here we list the data for family 16:

```
. list family subject dtlm pr re1 re2 if family==16, sepby(subject)
```

	family	subject	dtlm	pr	re1	re2
208.	16	5	1	.5301687	.5051272	.1001124
209.	16	5	0	.1956408	.5051272	.1001124
210.	16	5	0	.0367041	.5051272	.1001124
211.	16	34	1	.8876646	.5051272	.7798247
212.	16	34	1	.6107262	.5051272	.7798247
213.	16	34	1	.2572725	.5051272	.7798247
214.	16	35	0	.6561904	.5051272	-.0322885
215.	16	35	1	.2977437	.5051272	-.0322885
216.	16	35	0	.071612	.5051272	-.0322885

The predicted random effects at the family level (`re1`) are the same for all members of the family. Similarly, the predicted random effects at the individual level (`re2`) are constant within each individual. The predicted probabilities (`pr`) for this family seem to be in fair agreement with the response (`dt1m`) based on a cutoff of 0.5.

We can use `estat icc` to estimate the residual intraclass correlation (conditional on the difficulty level and the individual's category) between the latent responses of subjects within the same family or between the latent responses of the same subject and family:

```
. estat icc
```

```
Residual intraclass correlation
```

Level	ICC	Std. Err.	[95% Conf. Interval]	
family	.1352637	.1050492	.0261998	.4762821
subject family	.3622485	.0877459	.2124808	.5445812

`estat icc` reports two intraclass correlations for this three-level nested model. The first is the level-3 intraclass correlation at the family level, the correlation between latent measurements of the cognitive ability in the same family. The second is the level-2 intraclass correlation at the subject-within-family level, the correlation between the latent measurements of cognitive ability in the same subject and family.

There is not a strong correlation between individual realizations of the latent response, even within the same subject.

◀

Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [ME] [meglm postestimation](#).

Also see

[ME] [meprobit](#) — Multilevel mixed-effects probit regression

[ME] [meglm postestimation](#) — Postestimation tools for meglm

[U] [20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

meqrlogit, like melogit, fits mixed-effects models for binary or binomial responses. The conditional distribution of the response given the random effects is assumed to be Bernoulli, with success probability determined by the logistic cumulative distribution function.

meqrlogit provides an alternative estimation method, which uses the QR decomposition of the variance-components matrix. This method may aid convergence when variance components are near the boundary of the parameter space.

Quick start

Two-level logistic regression of y on x with random intercepts by lev2 using QR decomposition

```
meqrlogit y x || lev2:
```

Add random coefficients for x

```
meqrlogit y x || lev2: x
```

As above, but allow the random effects to be correlated

```
meqrlogit y x || lev2: x, covariance(unstructured)
```

Three-level random-intercept model of y on x with lev2 nested within lev3

```
meqrlogit y x || lev3: || lev2:
```

Crossed-effects model of y on x with two-way crossed random effects by factors a and b

```
meqrlogit y x || _all:R.a || b:
```

Menu

Statistics > Multilevel mixed-effects models > Estimation by QR decomposition > Logistic regression

Syntax

```
meqrlogit depvar fe_equation || re_equation [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[indepvars] [if] [in] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname [ , re_options ]
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description

Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1

<i>re_options</i>	Description

Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>collinear</code>	keep collinear variables

<i>options</i>	Description
Model	
<u>binomial</u> (<i>varname</i> #)	set binomial trials if data are in binomial form
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
or	report fixed-effects coefficients as odds ratios
<u>variance</u>	show random-effects parameter estimates as variances and covariances; the default
<u>stddeviations</u>	show random-effects parameter estimates as standard deviations and correlations
<u>noretable</u>	suppress random-effects table
<u>nofetable</u>	suppress fixed-effects table
<u>estmetric</u>	show parameter estimates in the estimation metric
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intpoints</u> (# [<i># ...</i>])	set the number of integration (quadrature) points; default is <code>intpoints(7)</code>
<u>laplace</u>	use Laplacian approximation; equivalent to <code>intpoints(1)</code>
Maximization	
<i>maximize_options</i>	control the maximization process; seldom used
<u>retolerance</u> (#)	tolerance for random-effects estimates; default is <code>retolerance(1e-8)</code> ; seldom used
<u>reiterate</u> (#)	maximum number of iterations for random-effects estimation; default is <code>reiterate(50)</code> ; seldom used
<u>matsqrt</u>	parameterize variance components using matrix square roots; the default
<u>matlog</u>	parameterize variance components using matrix logarithms
<u>refineopts</u> (<i>maximize_options</i>)	control the maximization process during refinement of starting values
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the R. notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the R. notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated

indepvars may contain factor variables; see [U] 11.4.3 [Factor variables](#).

indepvars and *varlist* may contain time-series operators; see [U] 11.4.4 [Time-series varlists](#).

bootstrap, *by*, *jackknife*, *mi estimate*, *rolling*, and *statsby* are allowed; see [U] 11.1.10 [Prefix commands](#).

coeflegend does not appear in the dialog box.

See [U] 20 [Estimation and postestimation commands](#) for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: **independent**, **exchangeable**, **identity**, and **unstructured**.

covariance(independent) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is **covariance(independent)**, except when the *R*. notation is used, in which case the default is **covariance(identity)** and only **covariance(identity)** and **covariance(exchangeable)** are allowed.

covariance(exchangeable) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(identity) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(unstructured) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p+1)/2$ unique parameters.

collinear specifies that meqrlogit not omit collinear variables from the random-effects equation. Usually, there is no reason to leave collinear variables in place; in fact, doing so usually causes the estimation to fail because of the matrix singularity caused by the collinearity. However, with certain models (for example, a random-effects model with a full set of contrasts), the variables may be collinear, yet the model is fully identified because of restrictions on the random-effects covariance structure. In such cases, using the **collinear** option allows the estimation to take place with the random-effects equation intact.

binomial(*varname* | #) specifies that the data are in binomial form; that is, *depvar* records the number of successes from a series of binomial trials. This number of trials is given either as *varname*, which allows this number to vary over the observations, or as the constant #. If **binomial()** is not specified (the default), *depvar* is treated as Bernoulli, with any nonzero, nonmissing values indicating positive responses.

Reporting

level(#); see [R] [estimation options](#).

or reports estimated fixed-effects coefficients transformed to odds ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated. **or** may be specified either at estimation or upon replay.

`variance`, the default, displays the random-effects parameter estimates as variances and covariances. `stddeviations` displays the random-effects parameter estimates as standard deviations and correlations.

`noretable` suppresses the random-effects table.

`nofetable` suppresses the fixed-effects table.

`estmetric` displays all parameter estimates in the estimation metric. Fixed-effects estimates are unchanged from those normally displayed, but random-effects parameter estimates are displayed as log-standard deviations and hyperbolic arctangents of correlations, with equation names that organize them by model level.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intpoints(#[ # ...])` sets the number of integration points for adaptive Gaussian quadrature. The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases with the number of quadrature points, and in models with many levels or many random coefficients, this increase can be substantial.

You may specify one number of integration points applying to all levels of random effects in the model, or you may specify distinct numbers of points for each level. `intpoints(7)` is the default; that is, by default seven quadrature points are used for each level.

`laplace` specifies that log likelihoods be calculated using the Laplacian approximation, equivalent to adaptive Gaussian quadrature with one integration point for each level in the model; `laplace` is equivalent to `intpoints(1)`. Computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. The computational time saved by using `laplace` can thus be substantial, especially when you have many levels or random coefficients.

The Laplacian approximation has been known to produce biased parameter estimates, but the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects. If your interest lies primarily with the fixed-effects estimates, the Laplace approximation may be a viable faster alternative to adaptive quadrature with multiple integration points.

When the `R.varname` notation is used, the dimension of the random effects increases by the number of distinct values of `varname`. Even when this number is small to moderate, it increases the total random-effects dimension to the point where estimation with more than one quadrature point is prohibitively intensive.

For this reason, when you use the `R.` notation in your random-effects equations, the `laplace` option is assumed. You can override this behavior by using the `intpoints()` option, but doing so is not recommended.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonnrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for meqrlogit are listed below.

For the `technique()` option, the default is `technique(nr)`. The bhhh algorithm may not be specified.

`from(init_specs)` is particularly useful when combined with `refineopts(iterate(0))` (see the description [below](#)), which bypasses the initial optimization stage.

`retolerance(#)` specifies the convergence tolerance for the estimated random effects used by adaptive Gaussian quadrature. Although not estimated as model parameters, random-effects estimators are used to adapt the quadrature points. Estimating these random effects is an iterative procedure, with convergence declared when the maximum relative change in the random effects is less than `retolerance()`. The default is `retolerance(1e-8)`. You should seldom have to use this option.

`reiterate(#)` specifies the maximum number of iterations used when estimating the random effects to be used in adapting the Gaussian quadrature points; see the `retolerance()` option. The default is `reiterate(50)`. You should seldom have to use this option.

`matsqrt` (the default), during optimization, parameterizes variance components by using the matrix square roots of the variance–covariance matrices formed by these components at each model level.

`matlog`, during optimization, parameterizes variance components by using the matrix logarithms of the variance–covariance matrices formed by these components at each model level.

The `matsqrt` parameterization ensures that variance–covariance matrices are positive semidefinite, while `matlog` ensures matrices that are positive definite. For most problems, the matrix square root is more stable near the boundary of the parameter space. However, if convergence is problematic, one option may be to try the alternate `matlog` parameterization. When convergence is not an issue, both parameterizations yield equivalent results.

`refineopts(maximize_options)` controls the maximization process during the refinement of starting values. Estimation in meqrlogit takes place in two stages. In the first stage, starting values are refined by holding the quadrature points fixed between iterations. During the second stage, quadrature points are adapted with each evaluation of the log likelihood. Maximization options specified within `refineopts()` control the first stage of optimization; that is, they control the refining of starting values.

maximize_options specified outside `refineopts()` control the second stage.

The one exception to the above rule is the `nolog` option, which when specified outside `refineopts()` applies globally.

`from(init_specs)` is not allowed within `refineopts()` and instead must be specified globally.

Refining starting values helps make the iterations of the second stage (those that lead toward the solution) more numerically stable. In this regard, of particular interest is `refineopts(iterate(#))`, with two iterations being the default. Should the maximization fail because of instability in the Hessian calculations, one possible solution may be to increase the number of iterations here.

The following option is available with meqrlogit but is not shown in the dialog box:

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects logistic regression is logistic regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

`meqrlogit` allows for many levels of random effects. For example, in a three-level model you can specify random effects for schools and then random effects for classes nested within schools. In this model, the observations (presumably, the students) comprise the first level, the classes comprise the second level, and the schools comprise the third.

However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = 1 | \mathbf{u}_j) = H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) \quad (1)$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are the binary-valued y_{ij} , and we follow the standard Stata convention of treating $y_{ij} = 1$ if `depvarij` $\neq 0$ and treating $y_{ij} = 0$ otherwise. The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard logistic regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Finally, because this is logistic regression, $H(\cdot)$ is the logistic cumulative distribution function, which maps the linear predictor to the probability of a success ($y_{ij} = 1$), with $H(v) = \exp(v) / \{1 + \exp(v)\}$.

Model (1) may also be stated in terms of a latent linear response, where only $y_{ij} = I(y_{ij}^* > 0)$ is observed for the latent

$$y_{ij}^* = \mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j + \epsilon_{ij}$$

The errors ϵ_{ij} are distributed as logistic with mean 0 and variance $\pi^2/3$ and are independent of $\mathbf{u}_j$.

Model (1) is an example of a generalized linear mixed model (GLMM), which generalizes the linear mixed-effects (LME) model to non-Gaussian responses. You can fit LMEs in Stata by using `mixed` and fit GLMMs by using `meqglm`. Because of the relationship between LMEs and GLMMs, there is insight to be gained through examination of the linear mixed model. This is especially true for Stata users because the terminology, syntax, options, and output for fitting these types of models are nearly identical. See [ME] `mixed` and the references therein, particularly in [Introduction](#), for more information.

Multilevel models with binary responses have been used extensively in the health and social sciences. As just one example, [Leyland and Goldstein \(2001, sec. 3.6\)](#) describe a study of equity of healthcare in Great Britain. Multilevel models with binary and other limited dependent responses also have a long history in econometrics; [Rabe-Hesketh, Skrondal, and Pickles \(2005\)](#) provide an excellent survey.

Log-likelihood calculations for fitting any generalized mixed-effects model require integrating out the random effects. One widely used modern method is to directly estimate the integral required to calculate the log likelihood by Gauss–Hermite quadrature or some variation thereof. The estimation method

used by `meqrlogit` is a multicoefficient and multilevel extension of one of these quadrature types, namely, adaptive Gaussian quadrature (AGQ) based on conditional modes, with the multicoefficient extension from [Pinheiro and Bates \(1995\)](#) and the multilevel extension from [Pinheiro and Chao \(2006\)](#); see [Methods and formulas](#).

► Example 1: Two-level random-intercept model

[Ng et al. \(2006\)](#) analyze a subsample of data from the 1989 Bangladesh fertility survey ([Huq and Cleland 1990](#)), which polled 1,934 Bangladeshi women on their use of contraception.

```
. use http://www.stata-press.com/data/r15/bangladesh
(Bangladesh Fertility Survey, 1989)

. describe

Contains data from http://www.stata-press.com/data/r15/bangladesh.dta
  obs:      1,934      Bangladesh Fertility Survey, 1989
  vars:      7         28 May 2016 20:27
  size:     19,340      (_dta has notes)
```

variable name	storage type	display format	value label	variable label
district	byte	%9.0g		District
c_use	byte	%9.0g	yesno	Use contraception
urban	byte	%9.0g	urban	Urban or rural
age	float	%6.2f		Age, mean centered
child1	byte	%9.0g		1 child
child2	byte	%9.0g		2 children
child3	byte	%9.0g		3 or more children

Sorted by: district

The women sampled were from 60 districts, identified by the variable `district`. Each district contained either urban or rural areas (variable `urban`) or both. The variable `c_use` is the binary response, with a value of 1 indicating contraceptive use. Other covariates include mean-centered `age` and three indicator variables recording number of children.

Consider a standard logistic regression model, amended to have random effects for each district. Defining $\pi_{ij} = \Pr(c\_use_{ij} = 1)$, we have

$$\text{logit}(\pi_{ij}) = \beta_0 + \beta_1 \text{urban}_{ij} + \beta_2 \text{age}_{ij} + \beta_3 \text{child1}_{ij} + \beta_4 \text{child2}_{ij} + \beta_5 \text{child3}_{ij} + u_j \quad (2)$$

for $j = 1, \dots, 60$ districts, with $i = 1, \dots, n_j$ women in district j .

```
. meqrlogit c_use urban age child* || district:
Refining starting values:
Iteration 0:   log likelihood = -1219.2682
Iteration 1:   log likelihood = -1209.3544
Iteration 2:   log likelihood = -1207.1895
Performing gradient-based optimization:
Iteration 0:   log likelihood = -1207.1895
Iteration 1:   log likelihood = -1206.8323
Iteration 2:   log likelihood = -1206.8322
Iteration 3:   log likelihood = -1206.8322
Mixed-effects logistic regression
Group variable: district
Number of obs      =      1,934
Number of groups   =        60
Obs per group:
    min =          2
    avg =         32.2
    max =         118
Integration points =       7
Log likelihood = -1206.8322
Wald chi2(5)      =      109.60
Prob > chi2       =      0.0000
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
urban	.7322764	.1194857	6.13	0.000	.4980887	.9664641
age	-.0264982	.0078916	-3.36	0.001	-.0419654	-.0110309
child1	1.116002	.1580921	7.06	0.000	.8061466	1.425856
child2	1.365895	.174669	7.82	0.000	1.02355	1.70824
child3	1.344031	.1796549	7.48	0.000	.9919141	1.696148
_cons	-1.68929	.1477592	-11.43	0.000	-1.978892	-1.399687

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
district: Identity				
var(_cons)	.2156188	.0733234	.1107202	.4199007

LR test vs. logistic model: chibar2(01) = 43.39 Prob >= chibar2 = 0.0000

Notes:

1. The estimation log consists of two parts:
 - (a) A set of iterations aimed at refining starting values. These are designed to be relatively quick iterations aimed at getting the parameter estimates within a neighborhood of the eventual solution, making the iterations in (b) more numerically stable.
 - (b) A set of gradient-based iterations. By default, these are Newton–Raphson iterations, but other methods are available by specifying the appropriate *maximize_options*; see [\[R\] maximize](#).
2. The first estimation table reports the fixed effects, and these can be interpreted just as you would the output from `logit`. You can also specify the `or` option at estimation or on replay to display the fixed effects as odds ratios instead.

If you did display results as odds ratios, you would find urban women to have roughly double the odds of using contraception as that of their rural counterparts. Having any number of children will increase the odds from three- to fourfold when compared with the base category of no children. Contraceptive use also decreases with age.

3. The second estimation table shows the estimated variance components. The first section of the table is labeled `district: Identity`, meaning that these are random effects at the `district` level and that their variance–covariance matrix is a multiple of the identity matrix; that is, $\Sigma = \sigma_u^2 \mathbf{I}$.

Because we have only one random effect at this level, meqrlogit knew that Identity is the only possible covariance structure. In any case, σ_u^2 was estimated as 0.22 with standard error 0.07.

4. A likelihood-ratio test comparing the model to ordinary logistic regression, (2) without u_j , is provided and is highly significant for these data.

◀

The above extends to models with more than two levels of nesting by adding more random-effects equations, each separated by | |.

Stored results

meqrlogit stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(reparm_rc)</code>	return code, final reparameterization
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	meqrlogit
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	logistic
<code>e(title)</code>	title in estimation output
<code>e(offset)</code>	offset
<code>e(binomial)</code>	binomial number of trials
<code>e(redim)</code>	random-effects dimensions
<code>e(vartypes)</code>	variance-structure types
<code>e(revars)</code>	random-effects covariates
<code>e(n_quad)</code>	number of integration points
<code>e(laplace)</code>	laplace, if Laplace approximation
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	bootstrap or jackknife if defined
<code>e(vctype)</code>	title used to label Std. Err.
<code>e(method)</code>	ML
<code>e(opt)</code>	type of optimization
<code>e(ml_method)</code>	type of ml method
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	b V
<code>e(estat_cmd)</code>	program used to implement estat
<code>e(predict)</code>	program used to implement predict
<code>e(marginsnotok)</code>	predictions disallowed by margins

<code>e(marginsdefault)</code>	default <code>predict()</code> specification for margins
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

Model (1) assumes Bernoulli data, a special case of the binomial. Because binomial data are also supported by `meqrlogit` (option `binomial()`), the methods presented below are in terms of the more general binomial mixed-effects model.

For a two-level binomial model, consider the response y_{ij} as the number of successes from a series of r_{ij} Bernoulli trials (replications). For cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$\begin{aligned} f(\mathbf{y}_j | \mathbf{u}_j) &= \prod_{i=1}^{n_j} \left[\binom{r_{ij}}{y_{ij}} \{H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\}^{y_{ij}} \{1 - H(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\}^{r_{ij}-y_{ij}} \right] \\ &= \exp \left(\sum_{i=1}^{n_j} \left[y_{ij} (\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) - r_{ij} \log \{1 + \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\} + \log \binom{r_{ij}}{y_{ij}} \right] \right) \end{aligned}$$

for $H(v) = \exp(v) / \{1 + \exp(v)\}$.

Defining $\mathbf{r}_j = (r_{j1}, \dots, r_{jn_j})'$ and

$$c(\mathbf{y}_j, \mathbf{r}_j) = \sum_{i=1}^{n_j} \log \binom{r_{ij}}{y_{ij}}$$

where $c(\mathbf{y}_j, \mathbf{r}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j | \mathbf{u}_j) = \exp [\mathbf{y}_j' (\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) - \mathbf{r}_j' \log \{1 + \exp(\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j)\} + c(\mathbf{y}_j, \mathbf{r}_j)]$$

where $\mathbf{X}_j$ is formed by stacking the row vectors $\mathbf{x}_{ij}$ and $\mathbf{Z}_j$ is formed by stacking the row vectors $\mathbf{z}_{ij}$. We extend the definitions of the functions $\log(\cdot)$ and $\exp(\cdot)$ to be vector functions where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{y}_j | \mathbf{u}_j) \exp(-\mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp \{c(\mathbf{y}_j, \mathbf{r}_j)\} (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int \exp \{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \tag{5}$$

where

$$h(\beta, \Sigma, \mathbf{u}_j) = \mathbf{y}_j' (\mathbf{X}_j \beta + \mathbf{Z}_j \mathbf{u}_j) - \mathbf{r}_j' \log \{1 + \exp(\mathbf{X}_j \beta + \mathbf{Z}_j \mathbf{u}_j)\} - \mathbf{u}_j' \Sigma^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{r}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (5) has no closed form and thus must be approximated. The Laplacian approximation (Tierney and Kadane 1986; Pinheiro and Bates 1995) is based on a second-order Taylor expansion of $h(\beta, \Sigma, \mathbf{u}_j)$ about the value of $\mathbf{u}_j$ that maximizes it. Taking first and second derivatives, we obtain

$$\begin{aligned} h'(\beta, \Sigma, \mathbf{u}_j) &= \frac{\partial h(\beta, \Sigma, \mathbf{u}_j)}{\partial \mathbf{u}_j} = \mathbf{Z}_j' \{\mathbf{y}_j - \mathbf{m}(\beta, \mathbf{u}_j)\} - \Sigma^{-1} \mathbf{u}_j \\ h''(\beta, \Sigma, \mathbf{u}_j) &= \frac{\partial^2 h(\beta, \Sigma, \mathbf{u}_j)}{\partial \mathbf{u}_j \partial \mathbf{u}_j'} = -\{\mathbf{Z}_j' \mathbf{V}(\beta, \mathbf{u}_j) \mathbf{Z}_j + \Sigma^{-1}\} \end{aligned}$$

where $\mathbf{m}(\beta, \mathbf{u}_j)$ is the vector function with the i th element equal to the conditional mean of y_{ij} given $\mathbf{u}_j$, that is, $r_{ij}H(\mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j)$. $\mathbf{V}(\beta, \mathbf{u}_j)$ is the diagonal matrix whose diagonal entries v_{ij} are the conditional variances of y_{ij} given $\mathbf{u}_j$, namely,

$$v_{ij} = r_{ij}H(\mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j) \{1 - H(\mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j)\}$$

The maximizer of $h(\beta, \Sigma, \mathbf{u}_j)$ is $\hat{\mathbf{u}}_j$ such that $h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{0}$. The integrand in (5) is proportional to the posterior density $f(\mathbf{u}_j | \mathbf{y}_j)$, so $\hat{\mathbf{u}}_j$ also represents the posterior mode, a plausible estimator of $\mathbf{u}_j$ in its own right.

Given the above derivatives, the second-order Taylor approximation then takes the form

$$h(\beta, \Sigma, \mathbf{u}_j) \approx h(\beta, \Sigma, \hat{\mathbf{u}}_j) + \frac{1}{2} (\mathbf{u}_j - \hat{\mathbf{u}}_j)' h''(\beta, \Sigma, \hat{\mathbf{u}}_j) (\mathbf{u}_j - \hat{\mathbf{u}}_j) \quad (6)$$

The first-derivative term vanishes because $h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{0}$. Therefore,

$$\begin{aligned} \int \exp \{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j &\approx \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j)\} \\ &\quad \times \int \exp \left[-\frac{1}{2} (\mathbf{u}_j - \hat{\mathbf{u}}_j)' \{-h''(\beta, \Sigma, \hat{\mathbf{u}}_j)\} (\mathbf{u}_j - \hat{\mathbf{u}}_j) \right] d\mathbf{u}_j \quad (7) \\ &= \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j)\} (2\pi)^{q/2} | -h''(\beta, \Sigma, \hat{\mathbf{u}}_j) |^{-1/2} \end{aligned}$$

because the latter integrand can be recognized as the “kernel” of a multivariate normal density.

Combining the above with (5) (and taking logs) gives the Laplacian log-likelihood contribution of the j th cluster,

$$\mathcal{L}_j^{\text{Lap}}(\beta, \Sigma) = -\frac{1}{2} \log |\Sigma| - \log |\mathbf{R}_j| + h(\beta, \Sigma, \hat{\mathbf{u}}_j) + c(\mathbf{y}_j, \mathbf{r}_j)$$

where $\mathbf{R}_j$ is an upper-triangular matrix such that $-h''(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{R}_j \mathbf{R}_j'$. Pinheiro and Chao (2006) show that $\hat{\mathbf{u}}_j$ and $\mathbf{R}_j$ can be efficiently computed as the iterative solution to a least-squares problem by using matrix decomposition methods similar to those used in fitting LME models (Bates and Pinheiro 1998; Pinheiro and Bates 2000; [ME] mixed).

The fidelity of the Laplacian approximation is determined wholly by the accuracy of the approximation in (6). An alternative that does not depend so heavily on this approximation is integration via AGQ (Naylor and Smith 1982; Liu and Pierce 1994).

The application of AGQ to this particular problem is from Pinheiro and Bates (1995). When we reexamine the integral in question, a transformation of integration variables yields

$$\begin{aligned} \int \exp \{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j &= |\mathbf{R}_j|^{-1} \int \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1}\mathbf{t})\} d\mathbf{t} \\ &= (2\pi)^{q/2} |\mathbf{R}_j|^{-1} \int \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1}\mathbf{t}) + \mathbf{t}'\mathbf{t}/2\} \phi(\mathbf{t}) d\mathbf{t} \end{aligned} \quad (8)$$

where $\phi(\cdot)$ is the standard multivariate normal density. Because the integrand is now expressed as some function multiplied by a normal density, it can be estimated by applying the rules of standard Gauss–Hermite quadrature. For a predetermined number of quadrature points N_Q , define $a_k = \sqrt{2}a_k^*$ and $w_k = w_k^*/\sqrt{\pi}$, for $k = 1, \dots, N_Q$, where (a_k^*, w_k^*) are a set of abscissas and weights for Gauss–Hermite quadrature approximations of $\int \exp(-x^2)f(x)dx$, as obtained from Abramowitz and Stegun (1964, 924).

Define $\mathbf{a}_k = (a_{k_1}, a_{k_2}, \dots, a_{k_q})'$; that is, $\mathbf{a}_k$ is a vector that spans the N_Q abscissas over the dimension q of the random effects. Applying quadrature rules to (8) yields the AGQ approximation,

$$\begin{aligned} \int \exp \{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j &\approx (2\pi)^{q/2} |\mathbf{R}_j|^{-1} \sum_{k_1=1}^{N_Q} \cdots \sum_{k_q=1}^{N_Q} \left[\exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1}\mathbf{a}_k) + \mathbf{a}_k'\mathbf{a}_k/2\} \prod_{p=1}^q w_{k_p} \right] \\ &\equiv (2\pi)^{q/2} \hat{G}_j(\beta, \Sigma) \end{aligned}$$

resulting in the AGQ log-likelihood contribution of the i th cluster,

$$\mathcal{L}_j^{\text{AGQ}}(\beta, \Sigma) = -\frac{1}{2} \log |\Sigma| + \log \left\{ \hat{G}_j(\beta, \Sigma) \right\} + c(\mathbf{y}_j, \mathbf{r}_j)$$

The “adaptive” part of adaptive Gaussian quadrature lies in the translation and rescaling of the integration variables in (8) by using $\hat{\mathbf{u}}_j$ and $\mathbf{R}_j^{-1}$, respectively. This transformation of quadrature abscissas (centered at 0 in standard form) is chosen to better capture the features of the integrand, which through (7) can be seen to resemble a multivariate normal distribution with mean $\hat{\mathbf{u}}_j$ and variance $\mathbf{R}_j^{-1}\mathbf{R}_j^{-T}$. AGQ is therefore not as dependent as the Laplace method upon the approximation in (6). In AGQ, (6) serves merely to redirect the quadrature abscissas, with the AGQ approximation improving as the number of quadrature points N_Q increases. In fact, Pinheiro and Bates (1995) point out that AGQ with only one quadrature point ($a = 0$ and $w = 1$) reduces to the Laplacian approximation.

The log likelihood for the entire dataset is then simply the sum of the contributions of the M individual clusters, namely, $\mathcal{L}(\beta, \Sigma) = \sum_{j=1}^M \mathcal{L}_j^{\text{Lap}}(\beta, \Sigma)$ for Laplace and $\mathcal{L}(\beta, \Sigma) = \sum_{j=1}^M \mathcal{L}_j^{\text{AGQ}}(\beta, \Sigma)$ for AGQ.

Maximization of $\mathcal{L}(\beta, \Sigma)$ is performed with respect to (β, θ) , where θ is a vector comprising the unique elements of the matrix square root of Σ . This is done to ensure that Σ is always positive semidefinite. If the `matlog` option is specified, then θ instead consists of the unique elements of the matrix logarithm of Σ . For well-conditioned problems, both methods produce equivalent results, yet our experience deems the former as more numerically stable near the boundary of the parameter space.

Once maximization is achieved, parameter estimates are mapped from $(\hat{\beta}, \hat{\theta})$ to $(\hat{\beta}, \hat{\gamma})$, where $\hat{\gamma}$ is a vector containing the unique (estimated) elements of Σ , expressed as logarithms of standard deviations for the diagonal elements and hyperbolic arctangents of the correlations for off-diagonal elements. This last step is necessary to (a) obtain a parameterization under which parameter estimates can be displayed and interpreted individually, rather than as elements of a matrix square root (or logarithm), and (b) parameterize these elements such that their ranges each encompass the entire real line.

Parameter estimates are stored in `e(b)` as $(\hat{\beta}, \hat{\gamma})$, with the corresponding variance–covariance matrix stored in `e(V)`. Parameter estimates can be displayed in this metric by specifying the `estmetric` option. However, in `meqrlogit` output, variance components are displayed as variances and covariances.

The approach outlined above can be extended from two-level models to higher-level models; see [Pinheiro and Chao \(2006\)](#) for details.

References

- Abramowitz, M., and I. A. Stegun, ed. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Washington, DC: National Bureau of Standards.
- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies. <http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Harbord, R. M., and P. Whiting. 2009. [metandi: Meta-analysis of diagnostic accuracy using hierarchical logistic regression](#). *Stata Journal* 9: 211–229.
- Huq, N. M., and J. Cleland. 1990. *Bangladesh Fertility Survey 1989 (Main Report)*. National Institute of Population Research and Training.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- Leyland, A. H., and H. Goldstein, ed. 2001. *Multilevel Modelling of Health Statistics*. New York: Wiley.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Liu, Q., and D. A. Pierce. 1994. A note on Gauss–Hermite quadrature. *Biometrika* 81: 624–629.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Naylor, J. C., and A. F. M. Smith. 1982. Applications of a method for the efficient computation of posterior distributions. *Journal of the Royal Statistical Society, Series C* 31: 214–225.
- Ng, E. S.-W., J. R. Carpenter, H. Goldstein, and J. Rasbash. 2006. Estimation in generalised linear mixed models with binary outcomes by simulated maximum likelihood. *Statistical Modelling* 6: 23–42.
- Palmer, T. M., C. M. Macdonald-Wallis, D. A. Lawlor, and K. Tilling. 2014. [Estimating adjusted associations between random effects from multilevel models: The reffadjust package](#). *Stata Journal* 14: 119–140.
- Pinheiro, J. C., and D. M. Bates. 1995. Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics* 4: 12–35.

- . 2000. *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- Pinheiro, J. C., and E. C. Chao. 2006. Efficient Laplacian and adaptive Gaussian quadrature algorithms for multilevel generalized linear mixed models. *Journal of Computational and Graphical Statistics* 15: 58–81.
- Rabe-Hesketh, S., A. Skrondal, and A. Pickles. 2005. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics* 128: 301–323.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.
- Tierney, L., and J. B. Kadane. 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* 81: 82–86.

Also see

- [ME] **meqrlogit postestimation** — Postestimation tools for meqrlogit
- [ME] **mecloglog** — Multilevel mixed-effects complementary log-log regression
- [ME] **melogit** — Multilevel mixed-effects logistic regression
- [ME] **meprobit** — Multilevel mixed-effects probit regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [MI] **estimation** — Estimation commands for use with mi estimate
- [SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)
- [XT] **xtlogit** — Fixed-effects, random-effects, and population-averaged logit models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples
Also see

[predict](#)
Methods and formulas

[margins](#)
References

Postestimation commands

The following postestimation commands are of special interest after `meqrlogit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat icc</code>	estimate intraclass correlations
<code>estat recovariance</code>	display the estimated random-effects covariance matrices
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estimates</code>	cataloging estimation results
<code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
<code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [reses(newvarsspec)
      relevel(levelvar) ]
```

Syntax for obtaining other predictions

```
predict [type] newvar [if] [in] [, statistic nooffset fixedonly]
```

newvarsspec is *stub** or *newvarlist*.

<i>statistic</i>	Description
Main	
<code>mu</code>	mean response; the default
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code><u>pearson</u></code>	Pearson residuals
<code><u>deviance</u></code>	deviance residuals
<code><u>anscombe</u></code>	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

Options for predict

Main

`reffects` calculates posterior modal estimates of the random effects. By default, estimates for all random effects in the model are calculated. However, if the `relevel(levelvar)` option is specified, then estimates for only level *levelvar* in the model are calculated. For example, if `classes` are nested within `schools`, then typing

```
. predict b*, reffects relevel(school)
```

would yield random-effects estimates at the school level. You must specify *q* new variables, where *q* is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you.

reses(*newvarsspec*) calculates standard errors for the random-effects estimates. By default, standard errors for all random effects in the model are calculated. However, if the **relevel**(*levelvar*) option is specified, then standard errors for only level *levelvar* in the model are calculated; see the **reffects** option.

You must specify q new variables, where q is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you. The new variables will have the same storage type as the corresponding random-effects variables.

The **reffects** and **reses**() options often generate multiple new variables at once. When this occurs, the random effects (or standard errors) contained in the generated variables correspond to the order in which the variance components are listed in the output of **meqrlogit**. Still, examining the variable labels of the generated variables (with the **describe** command, for instance) can be useful in deciphering which variables correspond to which terms in the model.

relevel(*levelvar*) specifies the level in the model at which predictions for random effects and their standard errors are to be obtained. *levelvar* is the name of the model level and is either the name of the variable describing the grouping at that level or is **_all**, a special designation for a group comprising all the estimation data.

mu, the default, calculates the predicted mean. By default, this is based on a linear predictor that includes both the fixed effects and the random effects, and the predicted mean is conditional on the values of the random effects. Use the **fixedonly** option (see [below](#)) if you want predictions that include only the fixed portion of the model, that is, if you want random effects set to 0.

xb calculates the linear prediction $\mathbf{x}\beta$ based on the estimated fixed effects (coefficients) in the model. This is equivalent to fixing all random effects in the model to their theoretical (prior) mean value of 0.

stdp calculates the standard error of the fixed-effects linear predictor $\mathbf{x}\beta$.

pearson calculates Pearson residuals. Pearson residuals large in absolute value may indicate a lack of fit. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

deviance calculates deviance residuals. Deviance residuals are recommended by [McCullagh and Nelder \(1989\)](#) as having the best properties for examining the goodness of fit of a GLM. They are approximately normally distributed if the model is correctly specified. They may be plotted against the fitted values or against a covariate to inspect the model's fit. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

anscombe calculates Anscombe residuals, which are designed to closely follow a normal distribution. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

nooffset is relevant only if you specified **offset**(*varname*) with **meqrlogit**. It modifies the calculations made by **predict** so that they ignore the offset variable; the linear prediction is treated as $\mathbf{X}\beta + \mathbf{Zu}$ rather than $\mathbf{X}\beta + \mathbf{Zu} + \text{offset}$.

fixedonly modifies predictions to include only the fixed portion of the model, equivalent to setting all random effects equal to 0; see the **mu** option.

margins

Description for margins

`margins` estimates margins of response for linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [options]
```

<i>statistic</i>	Description
<code>xb</code>	linear predictor for the fixed portion of the model only; the default
<code><u>r</u>effects</code>	not allowed with <code>margins</code>
<code>mu</code>	not allowed with <code>margins</code>
<code>stdp</code>	not allowed with <code>margins</code>
<code><u>p</u>earson</code>	not allowed with <code>margins</code>
<code><u>d</u>eviance</code>	not allowed with <code>margins</code>
<code><u>a</u>nscombe</code>	not allowed with <code>margins</code>

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.

For the full syntax, see [\[R\] margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a logistic mixed-effects model with `meqrlogit`. For the most part, calculation centers around obtaining estimates of the subject/group-specific random effects. Random effects are not provided as estimates when the model is fit but instead need to be predicted after estimation. Calculation of intraclass correlations, estimating the dependence between latent linear responses for different levels of nesting, may also be of interest.

► Example 1: Obtaining predicted probabilities and random effects

In [example 3](#) of [\[ME\] melogit](#), we represented the probability of contraceptive use among Bangladeshi women by using the model (stated with slightly different notation here)

$$\text{logit}(\pi_{ij}) = \beta_0\text{rural}_{ij} + \beta_1\text{urban}_{ij} + \beta_2\text{age}_{ij} + \beta_3\text{child1}_{ij} + \beta_4\text{child2}_{ij} + \beta_5\text{child3}_{ij} + a_j\text{rural}_{ij} + b_j\text{urban}_{ij}$$

where π_{ij} is the probability of contraceptive use, $j = 1, \dots, 60$ districts, $i = 1, \dots, n_j$ women within each district, and a_j and b_j are normally distributed with mean 0 and variance–covariance matrix

$$\Sigma = \text{Var} \begin{bmatrix} a_j \\ b_j \end{bmatrix} = \begin{bmatrix} \sigma_a^2 & \sigma_{ab} \\ \sigma_{ab} & \sigma_b^2 \end{bmatrix}$$

The purpose of using this particular model was to allow for district random effects that were specific to the rural and urban areas of that district and that could be interpreted as such. Below we fit the model using meqrlogit.

```
. use http://www.stata-press.com/data/r15/bangladesh
(Bangladesh Fertility Survey, 1989)

. generate byte rural = 1 - urban

. meqrlogit c_use rural urban age child*, noconstant || district: rural urban,
> noconstant covariance(unstructured)

Refining starting values:
Iteration 0:   log likelihood = -1208.3924
Iteration 1:   log likelihood = -1204.1317
Iteration 2:   log likelihood = -1200.6012

Performing gradient-based optimization:
Iteration 0:   log likelihood = -1200.6012
Iteration 1:   log likelihood = -1199.3332
Iteration 2:   log likelihood = -1199.315
Iteration 3:   log likelihood = -1199.315

Mixed-effects logistic regression
Group variable: district

Number of obs      =      1,934
Number of groups   =         60
Obs per group:
    min            =         2
    avg            =      32.2
    max            =      118

Integration points =    7
Log likelihood     = -1199.315

Wald chi2(6)       =      120.24
Prob > chi2        =      0.0000
```

c_use	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
rural	-1.71165	.1605618	-10.66	0.000	-2.026345	-1.396954
urban	-.8958623	.1704961	-5.25	0.000	-1.230028	-.5616961
age	-.026415	.008023	-3.29	0.001	-.0421398	-.0106902
child1	1.13252	.1603285	7.06	0.000	.818282	1.446758
child2	1.357739	.1770522	7.67	0.000	1.010724	1.704755
child3	1.353827	.1828801	7.40	0.000	.9953882	1.712265

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
district: Unstructured				
var(rural)	.3897439	.1292459	.2034726	.7465393
var(urban)	.2442965	.1450674	.0762886	.782303
cov(rural,urban)	-.0161411	.1057469	-.2234012	.1911189

LR test vs. logistic model: chi2(3) = 58.42 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

We can now obtain predicted probabilities, the default prediction:

```
. predict p
(option mu assumed; predicted means)
```

These predictions are based on a linear predictor that includes both the fixed effects and the random effects due to district. Specifying the `fixedonly` option gives predictions that set the random effects to their prior mean of 0. Below we compare both over the first 20 observations:

```
. predict p_fixed, fixedonly
(option mu assumed; predicted means)
. list c_use p p_fixed age child1 child2 child3
```

	c_use	p	p_fixed	age	child1	child2	child3
1.	no	.3579543	.4927183	18.44	0	0	1
2.	no	.2134724	.3210403	-5.56	0	0	0
3.	no	.4672256	.6044016	1.44	0	1	0
4.	no	.4206505	.5584864	8.44	0	0	1
5.	no	.2510909	.3687281	-13.56	0	0	0
6.	no	.2412878	.3565185	-11.56	0	0	0
7.	no	.3579543	.4927183	18.44	0	0	1
8.	no	.4992191	.6345999	-3.56	0	0	1
9.	no	.4572049	.594723	-5.56	1	0	0
10.	no	.4662518	.6034657	1.44	0	0	1
11.	yes	.2412878	.3565185	-11.56	0	0	0
12.	no	.2004691	.3040173	-2.56	0	0	0
13.	no	.4506573	.5883407	-4.56	1	0	0
14.	no	.4400747	.5779263	5.44	0	0	1
15.	no	.4794194	.6160359	-0.56	0	0	1
16.	yes	.4465936	.5843561	4.44	0	0	1
17.	no	.2134724	.3210403	-5.56	0	0	0
18.	yes	.4794194	.6160359	-0.56	0	0	1
19.	yes	.4637673	.6010735	-6.56	1	0	0
20.	no	.5001973	.6355067	-3.56	0	1	0



Technical note

Out-of-sample predictions are permitted after `meqrlglogit`, but if these predictions involve estimated random effects, the integrity of the estimation data must be preserved. If the estimation data have changed since the model was fit, `predict` will be unable to obtain predicted random effects that are appropriate for the fitted model and will give an error. Thus to obtain out-of-sample predictions that contain random-effects terms, be sure that the data for these predictions are in observations that augment the estimation data.



Methods and formulas

Continuing the discussion in *Methods and formulas* of [ME] `meqrlglogit`, and using the definitions and formulas defined there, we begin by considering the prediction of the random effects u_j for the j th cluster in a two-level model.

Given a set of estimated `meqrlogit` parameters $(\widehat{\beta}, \widehat{\Sigma})$, a profile likelihood in $\mathbf{u}_j$ is derived from the joint distribution $f(\mathbf{y}_j, \mathbf{u}_j)$ as

$$\mathcal{L}_j(\mathbf{u}_j) = \exp\{c(\mathbf{y}_j, \mathbf{r}_j)\} (2\pi)^{-q/2} |\widehat{\Sigma}|^{-1/2} \exp\left\{g\left(\widehat{\beta}, \widehat{\Sigma}, \mathbf{u}_j\right)\right\} \quad (1)$$

The conditional maximum likelihood estimator of $\mathbf{u}_j$ —conditional on fixed $(\widehat{\beta}, \widehat{\Sigma})$ —is the maximizer of $\mathcal{L}_j(\mathbf{u}_j)$ or, equivalently, the value of $\widehat{\mathbf{u}}_j$ that solves

$$\mathbf{0} = g'\left(\widehat{\beta}, \widehat{\Sigma}, \widehat{\mathbf{u}}_j\right) = \mathbf{Z}'_j \left\{ \mathbf{y}_j - \mathbf{m}(\widehat{\beta}, \widehat{\mathbf{u}}_j) \right\} - \widehat{\Sigma}^{-1} \widehat{\mathbf{u}}_j$$

Because (1) is proportional to the conditional density $f(\mathbf{u}_j | \mathbf{y}_j)$, you can also refer to $\widehat{\mathbf{u}}_j$ as the conditional mode (or posterior mode if you lean toward Bayesian terminology). Regardless, you are referring to the same estimator.

Conditional standard errors for the estimated random effects are derived from standard theory of maximum likelihood, which dictates that the asymptotic variance matrix of $\widehat{\mathbf{u}}_j$ is the negative inverse of the Hessian, which is estimated as

$$g''\left(\widehat{\beta}, \widehat{\Sigma}, \widehat{\mathbf{u}}_j\right) = -\left\{ \mathbf{Z}'_j \mathbf{V}(\widehat{\beta}, \widehat{\mathbf{u}}_j) \mathbf{Z}_j + \widehat{\Sigma}^{-1} \right\}$$

Similar calculations extend to models with more than one level of random effects; see [Pineiro and Chao \(2006\)](#).

For any observation i in the j th cluster in a two-level model, define the linear predictor as

$$\widehat{\eta}_{ij} = \mathbf{x}_{ij} \widehat{\beta} + \mathbf{z}_{ij} \widehat{\mathbf{u}}_j$$

In a three-level model, for the i th observation within the j th level-two cluster within the k th level-three cluster,

$$\widehat{\eta}_{ijk} = \mathbf{x}_{ijk} \widehat{\beta} + \mathbf{z}_{ijk}^{(3)} \widehat{\mathbf{u}}_k^{(3)} + \mathbf{z}_{ijk}^{(2)} \widehat{\mathbf{u}}_{jk}^{(2)}$$

where $\mathbf{z}^{(p)}$ and $\mathbf{u}^{(p)}$ refer to the level p design variables and random effects, respectively. For models with more than three levels, the definition of $\widehat{\eta}$ extends in the natural way, with only the notation becoming more complicated.

If the `fixedonly` option is specified, $\widehat{\eta}$ contains the linear predictor for only the fixed portion of the model, for example, in a two-level model $\widehat{\eta}_{ij} = \mathbf{x}_{ij} \widehat{\beta}$. In what follows, we assume a two-level model, with the only necessary modification for multilevel models being the indexing.

The predicted mean conditional on the random effects $\widehat{\mathbf{u}}_j$ is

$$\widehat{\mu}_{ij} = r_{ij} H(\widehat{\eta}_{ij})$$

Pearson residuals are calculated as

$$\nu_{ij}^P = \frac{y_{ij} - \widehat{\mu}_{ij}}{\{V(\widehat{\mu}_{ij})\}^{1/2}}$$

for $V(\widehat{\mu}_{ij}) = \widehat{\mu}_{ij}(1 - \widehat{\mu}_{ij}/r_{ij})$.

Deviance residuals are calculated as

$$\nu_{ij}^D = \text{sign}(y_{ij} - \hat{\mu}_{ij}) \sqrt{\hat{d}_{ij}^2}$$

where

$$\hat{d}_{ij}^2 = \begin{cases} 2r_{ij} \log \left(\frac{r_{ij}}{r_{ij} - \hat{\mu}_{ij}} \right) & \text{if } y_{ij} = 0 \\ 2y_{ij} \log \left(\frac{y_{ij}}{\hat{\mu}_{ij}} \right) + 2(r_{ij} - y_{ij}) \log \left(\frac{r_{ij} - y_{ij}}{r_{ij} - \hat{\mu}_{ij}} \right) & \text{if } 0 < y_{ij} < r_{ij} \\ 2r_{ij} \log \left(\frac{r_{ij}}{\hat{\mu}_{ij}} \right) & \text{if } y_{ij} = r_{ij} \end{cases}$$

Anscombe residuals are calculated as

$$\nu_{ij}^A = \frac{3 \left\{ y_{ij}^{2/3} \mathcal{H}(y_{ij}/r_{ij}) - \hat{\mu}_{ij}^{2/3} \mathcal{H}(\hat{\mu}_{ij}/r_{ij}) \right\}}{2 \left(\hat{\mu}_{ij} - \hat{\mu}_{ij}^2/r_{ij} \right)^{1/6}}$$

where $\mathcal{H}(t)$ is a specific univariate case of the Hypergeometric2F1 function (Wolfram 2003, 780). For Anscombe residuals for binomial regression, the specific form of the Hypergeometric2F1 function that we require is $\mathcal{H}(t) = {}_2F_1(2/3, 1/3, 5/3, t)$.

For a discussion of the general properties of the above residuals, see Hardin and Hilbe (2012, chap. 4).

References

- Hardin, J. W., and J. M. Hilbe. 2012. *Generalized Linear Models and Extensions*. 3rd ed. College Station, TX: Stata Press.
- McCullagh, P., and J. A. Nelder. 1989. *Generalized Linear Models*. 2nd ed. London: Chapman & Hall/CRC.
- Pinheiro, J. C., and E. C. Chao. 2006. Efficient Laplacian and adaptive Gaussian quadrature algorithms for multilevel generalized linear mixed models. *Journal of Computational and Graphical Statistics* 15: 58–81.
- Wolfram, S. 2003. *The Mathematica Book*. 5th ed. Champaign, IL: Wolfram Media.

Also see

- [ME] **meqrlgit** — Multilevel mixed-effects logistic regression (QR decomposition)
- [U] **20 Estimation and postestimation commands**

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

meqrpoisson, like [mepoisson](#), fits mixed-effects models for count responses, for which the conditional distribution of the response given the random effects is assumed to be Poisson.

meqrpoisson provides an alternative estimation method that uses the QR decomposition of the variance-components matrix. This method may aid convergence when variance components are near the boundary of the parameter space.

Quick start

Two-level Poisson regression of y on x with random intercepts by lev2 using QR decomposition

```
meqrpoisson y x || lev2:
```

As above, but report incidence-rate ratios

```
meqrpoisson y x || lev2:, irr
```

Add [indicator variables](#) for each level of categorical variable a and random coefficients on x

```
meqrpoisson y x i.a || lev2: x, irr
```

Three-level random-intercept model of y on x with lev2 nested within lev3

```
meqrpoisson y x || lev3: || lev2:
```

Menu

Statistics > Multilevel mixed-effects models > Estimation by QR decomposition > Poisson regression

Syntax

```
meqrpoisson depvar fe_equation || re_equation [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[indepvars] [if] [in] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname [ , re_options ]
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>exposure(<i>varname</i>_e)</code>	include $\ln(\text{varname}_e)$ in model with coefficient constrained to 1
<code>offset(<i>varname</i>_o)</code>	include <i>varname</i> _o in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>collinear</code>	keep collinear variables

<i>options</i>	Description
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>irr</u>	report fixed-effects coefficients as incidence-rate ratios
<u>variance</u>	show random-effects parameter estimates as variances and covariances; the default
<u>stddeviations</u>	show random-effects parameter estimates as standard deviations and correlations
<u>noretable</u>	suppress random-effects table
<u>nofetable</u>	suppress fixed-effects table
<u>estmetric</u>	show parameter estimates in the estimation metric
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>intpoints</u> (# [<i># ...</i>])	set the number of integration (quadrature) points; default is <code>intpoints(7)</code>
<u>laplace</u>	use Laplacian approximation; equivalent to <code>intpoints(1)</code>
Maximization	
<i>maximize_options</i>	control the maximization process; seldom used
<u>retolerance</u> (#)	tolerance for random-effects estimates; default is <code>retolerance(1e-8)</code> ; seldom used
<u>reiterate</u> (#)	maximum number of iterations for random-effects estimation; default is <code>reiterate(50)</code> ; seldom used
<u>matsqrt</u>	parameterize variance components using matrix square roots; the default
<u>matlog</u>	parameterize variance components using matrix logarithms
<u>refineopts</u> (<i>maximize_options</i>)	control the maximization process during refinement of starting values
<u>coeflegend</u>	display legend instead of statistics
<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the R. notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the R. notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated

indepvars may contain factor variables; see [U] 11.4.3 **Factor variables**.

indepvars and *varlist* may contain time-series operators; see [U] 11.4.4 **Time-series varlists**.

bootstrap, *by*, *jackknife*, *mi estimate*, *rolling*, and *statsby* are allowed; see [U] 11.1.10 **Prefix commands**.

coeflegend does not appear in the dialog box.

See [U] 20 **Estimation and postestimation commands** for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

exposure(*varname_e*) specifies a variable that reflects the amount of exposure over which the *depvar* events were observed for each observation; $\ln(\text{varname}_e)$ is included in the fixed-effects portion of the model with the coefficient constrained to be 1.

offset(*varname_o*) specifies that *varname_o* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: **independent**, **exchangeable**, **identity**, and **unstructured**.

covariance(**independent**) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is **covariance**(**independent**), except when the R. notation is used, in which case the default is **covariance**(**identity**) and only **covariance**(**identity**) and **covariance**(**exchangeable**) are allowed.

covariance(**exchangeable**) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(**identity**) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(**unstructured**) allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p+1)/2$ unique parameters.

collinear specifies that **meqrpoisson** not omit collinear variables from the random-effects equation.

Usually, there is no reason to leave collinear variables in place; in fact, doing so usually causes the estimation to fail because of the matrix singularity caused by the collinearity. However, with certain models (for example, a random-effects model with a full set of contrasts), the variables may be collinear, yet the model is fully identified because of restrictions on the random-effects covariance structure. In such cases, using the **collinear** option allows the estimation to take place with the random-effects equation intact.

Reporting

level(#); see [R] **estimation options**.

irr reports estimated fixed-effects coefficients transformed to incidence-rate ratios, that is, $\exp(\beta)$ rather than β . Standard errors and confidence intervals are similarly transformed. This option affects how results are displayed, not how they are estimated. **irr** may be specified at estimation or upon replay.

variance, the default, displays the random-effects parameter estimates as variances and covariances.

`stddeviations` displays the random-effects parameter estimates as standard deviations and correlations.

`noretable` suppresses the random-effects table.

`nofetable` suppresses the fixed-effects table.

`estmetric` displays all parameter estimates in the estimation metric. Fixed-effects estimates are unchanged from those normally displayed, but random-effects parameter estimates are displayed as log-standard deviations and hyperbolic arctangents of correlations, with equation names that organize them by model level.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `noci`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intpoints(#[ # ... ])` sets the number of integration points for adaptive Gaussian quadrature. The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases with the number of quadrature points, and in models with many levels or many random coefficients, this increase can be substantial.

You may specify one number of integration points applying to all levels of random effects in the model, or you may specify distinct numbers of points for each level. `intpoints(7)` is the default; that is, by default seven quadrature points are used for each level.

`laplace` specifies that log likelihoods be calculated using the Laplacian approximation, equivalent to adaptive Gaussian quadrature with one integration point for each level in the model; `laplace` is equivalent to `intpoints(1)`. Computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. The computational time saved by using `laplace` can thus be substantial, especially when you have many levels or random coefficients.

The Laplacian approximation has been known to produce biased parameter estimates, but the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects. If your interest lies primarily with the fixed-effects estimates, the Laplace approximation may be a viable faster alternative to adaptive quadrature with multiple integration points.

When the `R.varname` notation is used, the dimension of the random effects increases by the number of distinct values of `varname`. Even when this number is small to moderate, it increases the total random-effects dimension to the point where estimation with more than one quadrature point is prohibitively intensive.

For this reason, when you use the `R.` notation in your random-effects equations, the `laplace` option is assumed. You can override this behavior by using the `intpoints()` option, but doing so is not recommended.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrntolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `meqrpoisson` are listed below.

For the `technique()` option, the default is `technique(nr)`. The `bhhh` algorithm may not be specified.

`from(init_specs)` is particularly useful when combined with `refineopts(iterate(0))` (see the description [below](#)), which bypasses the initial optimization stage.

`retolerance(#)` specifies the convergence tolerance for the estimated random effects used by adaptive Gaussian quadrature. Although not estimated as model parameters, random-effects estimators are used to adapt the quadrature points. Estimating these random effects is an iterative procedure, with convergence declared when the maximum relative change in the random effects is less than `retolerance()`. The default is `retolerance(1e-8)`. You should seldom have to use this option.

`reiterate(#)` specifies the maximum number of iterations used when estimating the random effects to be used in adapting the Gaussian quadrature points; see the `retolerance()` option. The default is `reiterate(50)`. You should seldom have to use this option.

`matsqrt` (the default), during optimization, parameterizes variance components by using the matrix square roots of the variance–covariance matrices formed by these components at each model level.

`matlog`, during optimization, parameterizes variance components by using the matrix logarithms of the variance–covariance matrices formed by these components at each model level.

The `matsqrt` parameterization ensures that variance–covariance matrices are positive semidefinite, while `matlog` ensures matrices that are positive definite. For most problems, the matrix square root is more stable near the boundary of the parameter space. However, if convergence is problematic, one option may be to try the alternate `matlog` parameterization. When convergence is not an issue, both parameterizations yield equivalent results.

`refineopts(maximize_options)` controls the maximization process during the refinement of starting values. Estimation in `meqrpoisson` takes place in two stages. In the first stage, starting values are refined by holding the quadrature points fixed between iterations. During the second stage, quadrature points are adapted with each evaluation of the log likelihood. Maximization options specified within `refineopts()` control the first stage of optimization; that is, they control the refining of starting values.

maximize_options specified outside `refineopts()` control the second stage.

The one exception to the above rule is the `nolog` option, which when specified outside `refineopts()` applies globally.

`from(init_specs)` is not allowed within `refineopts()` and instead must be specified globally.

Refining starting values helps make the iterations of the second stage (those that lead toward the solution) more numerically stable. In this regard, of particular interest is `refineopts(iterate(#))`, with two iterations being the default. Should the maximization fail because of instability in the Hessian calculations, one possible solution may be to increase the number of iterations here.

The following option is available with `meqrpoisson` but is not shown in the dialog box:

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects Poisson regression is Poisson regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

`meqrpoisson` allows for many levels of random effects. For example, in a three-level model you can specify random effects for schools and then random effects for classes nested within schools. The observations (students, presumably) would comprise level one of the model, the classes would comprise level two, and the schools would comprise level three.

However, for simplicity, for now we consider the two-level model, where for a series of M independent clusters, and conditional on a set of random effects $\mathbf{u}_j$,

$$\Pr(y_{ij} = y | \mathbf{u}_j) = \exp(-\mu_{ij}) \mu_{ij}^y / y! \quad (1)$$

for $\mu_{ij} = \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)$, $j = 1, \dots, M$ clusters, and with cluster j consisting of $i = 1, \dots, n_j$ observations. The responses are counts y_{ij} . The $1 \times p$ row vector $\mathbf{x}_{ij}$ are the covariates for the fixed effects, analogous to the covariates you would find in a standard Poisson regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$.

The $1 \times q$ vector $\mathbf{z}_{ij}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ij}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{z}_{ij} = \mathbf{x}_{ij}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Model (1) is an example of a generalized linear mixed model (GLMM), which generalizes the linear mixed-effects (LME) model to non-Gaussian responses. You can fit LMEs in Stata by using `mixed` and fit GLMMs by using `meglm`. Because of the relationship between LMEs and GLMMs, there is insight to be gained through examination of the linear mixed model. This is especially true for Stata users because the terminology, syntax, options, and output for fitting these types of models are nearly identical. See [ME] [mixed](#) and the references therein, particularly in the [Introduction](#), for more information.

Log-likelihood calculations for fitting any generalized mixed-effects model require integrating out the random effects. One widely used modern method is to directly estimate the integral required to calculate the log likelihood by Gauss–Hermite quadrature or some variation thereof. The estimation method used by `meqrpoisson` is a multicoefficient and multilevel extension of one of these quadrature types, namely, adaptive Gaussian quadrature (AGQ) based on conditional modes, with the multicoefficient extension from [Pinheiro and Bates \(1995\)](#) and the multilevel extension from [Pinheiro and Chao \(2006\)](#); see [Methods and formulas](#).

► Example 1: Two-level random-intercept model

[Breslow and Clayton \(1993\)](#) fit a mixed-effects Poisson model to data from a randomized trial of the drug progabide for the treatment of epilepsy.


```
. use http://www.stata-press.com/data/r15/epilepsy
(Epilepsy data; progabide drug treatment)

. describe

Contains data from http://www.stata-press.com/data/r15/epilepsy.dta
   obs:                236                Epilepsy data; progabide drug
                                         treatment
   vars:                 8                31 May 2016 14:09
   size:               4,956              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
subject	byte	%9.0g		Subject ID: 1-59
seizures	int	%9.0g		No. of seizures
treat	byte	%9.0g		1: progabide; 0: placebo
visit	float	%9.0g		Dr. visit; coded as (-.3, -.1, .1, .3)
lage	float	%9.0g		log(age), mean-centered
lbas	float	%9.0g		log(0.25*baseline seizures), mean-centered
lbas_trt	float	%9.0g		lbas/treat interaction
v4	byte	%8.0g		Fourth visit indicator

Sorted by: subject

Originally from [Thall and Vail \(1990\)](#), data were collected on 59 subjects (31 on progabide, 28 on placebo). The number of epileptic seizures (`seizures`) was recorded during the two weeks prior to each of four doctor visits (`visit`). The treatment group is identified by the indicator variable `treat`. Data were also collected on the logarithm of age (`lage`) and the logarithm of one-quarter the number of seizures during the eight weeks prior to the study (`lbas`). The variable `lbas_trt` represents the interaction between `lbas` and treatment. `lage`, `lbas`, and `lbas_trt` are mean centered. Because the study originally noted a substantial decrease in seizures prior to the fourth doctor visit, an indicator, `v4`, for the fourth visit was also recorded.

[Breslow and Clayton \(1993\)](#) fit a random-effects Poisson model for the number of observed seizures

$$\log(\mu_{ij}) = \beta_0 + \beta_1 \text{treat}_{ij} + \beta_2 \text{lbas}_{ij} + \beta_3 \text{lbas\_trt}_{ij} + \beta_4 \text{lage}_{ij} + \beta_5 \text{v4}_{ij} + u_j$$

for $j = 1, \dots, 59$ subjects and $i = 1, \dots, 4$ visits. The random effects u_j are assumed to be normally distributed with mean 0 and variance σ_u^2 .

```
. meqrpoisson seizures treat lbas lbas_trt lage v4 || subject:
Refining starting values:
Iteration 0:   log likelihood = -680.40577   (not concave)
Iteration 1:   log likelihood = -668.60112
Iteration 2:   log likelihood = -666.37635
Performing gradient-based optimization:
Iteration 0:   log likelihood = -666.37635
Iteration 1:   log likelihood = -665.4589
Iteration 2:   log likelihood = -665.29074
Iteration 3:   log likelihood = -665.29068
Mixed-effects Poisson regression
Group variable: subject
Number of obs      =      236
Number of groups   =      59
Obs per group:
    min =      4
    avg  =     4.0
    max  =      4
Integration points =    7
Log likelihood = -665.29068
Wald chi2(5)      =     121.67
Prob > chi2       =     0.0000
```

seizures	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
treat	-.9330388	.4008346	-2.33	0.020	-1.71866	-.1474175
lbas	.884433	.1312313	6.74	0.000	.6272244	1.141642
lbas_trt	.3382609	.2033384	1.66	0.096	-.0602752	.7367969
lage	.4842389	.3472775	1.39	0.163	-.1964126	1.16489
v4	-.1610871	.0545758	-2.95	0.003	-.2680537	-.0541206
_cons	2.154574	.2200426	9.79	0.000	1.723299	2.58585

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
subject: Identity				
var(_cons)	.2528265	.058956	.1600785	.3993118

LR test vs. Poisson model: chibar2(01) = 304.74 Prob >= chibar2 = 0.0000

The number of seizures before the fourth visit does exhibit a significant drop, and the patients on progabide demonstrate a decrease in frequency of seizures compared with the placebo group. The subject-specific random effects also appear significant: $\hat{\sigma}_u^2 = 0.25$ with standard error 0.06.

Because this is a simple random-intercept model, you can obtain equivalent results by using `xtpoisson` with the `re` and `normal` options.



➤ Example 2: Three-level random-intercept model

Rabe-Hesketh and Skrondal (2012, exercise 13.7) describe data from the *Atlas of Cancer Mortality in the European Economic Community* (EEC) (Smans, Mair, and Boyle 1993). The data were analyzed in Langford, Bentham, and McDonald (1998) and record the number of deaths among males due to malignant melanoma during 1971–1980.

```
. use http://www.stata-press.com/data/r15/melanoma
(Skin cancer (melanoma) data)

. describe

Contains data from http://www.stata-press.com/data/r15/melanoma.dta
  obs:           354                Skin cancer (melanoma) data
 vars:           6                  30 May 2016 17:10
 size:          4,956              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
nation	byte	%11.0g	n	Nation ID
region	byte	%9.0g		Region ID: EEC level-I areas
county	int	%9.0g		County ID: EEC level-II/level-III areas
deaths	int	%9.0g		No. deaths during 1971-1980
expected	float	%9.0g		No. expected deaths
uv	float	%9.0g		UV dose, mean-centered

Sorted by:

Nine European nations (variable `nation`) are represented, and data were collected over geographical regions defined by EEC statistical services as level I areas (variable `region`), with deaths being recorded for each of 354 counties, which are level II or level III EEC-defined areas (variable `county`, which identifies the observations). Counties are nested within regions, and regions are nested within nations.

The variable `deaths` records the number of deaths for each county, and `expected` records the expected number of deaths (the exposure) on the basis of crude rates for the combined countries. Finally, the variable `uv` is a measure of exposure to ultraviolet (UV) radiation.

In modeling the number of deaths, one possibility is to include dummy variables for the nine nations as fixed effects. Another is to treat these as random effects and fit the three-level random-intercept Poisson model,

$$\log(\mu_{ijk}) = \log(\text{expected}_{ijk}) + \beta_0 + \beta_1 \text{uv}_{ijk} + \beta_2 \text{uv}_{ijk}^2 + u_k + v_{jk}$$

for nation k , region j , and county i . The model includes an exposure term for expected deaths.

```
. meqrpoisson deaths c.uv##c.uv, exposure(expected) || nation: || region:
Refining starting values:
Iteration 0:   log likelihood = -1169.4088   (not concave)
Iteration 1:   log likelihood = -1156.8957   (not concave)
Iteration 2:   log likelihood = -1101.8457
Performing gradient-based optimization:
Iteration 0:   log likelihood = -1101.8457
Iteration 1:   log likelihood = -1090.6214
Iteration 2:   log likelihood = -1089.4198
Iteration 3:   log likelihood = -1089.411
Iteration 4:   log likelihood = -1089.411
Mixed-effects Poisson regression                                Number of obs      =      354
```

Group Variable	No. of Groups	Observations per Group			Integration Points
		Minimum	Average	Maximum	
nation	9	3	39.3	95	7
region	78	1	4.5	13	7

```
Log likelihood = -1089.411                                Wald chi2(2)      =      25.69
                                                         Prob > chi2       =      0.0000
```

deaths	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
uv	.0056975	.0137931	0.41	0.680	-.0213364	.0327314
c.uv#c.uv	-.0058374	.001388	-4.21	0.000	-.0085579	-.0031169
_cons ln(expected)	.1289976 1	.1581122 (exposure)	0.82	0.415	-.1808967	.4388918

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
nation: Identity	var(_cons)	.1840722	.094531	.0672745	.5036466
region: Identity	var(_cons)	.0382743	.0087869	.0244057	.0600237

```
LR test vs. Poisson model: chi2(2) = 1267.13                Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

By including an exposure variable that is an expected rate, we are in effect specifying a linear model for the log of the standardized mortality ratio, the ratio of observed deaths to expected deaths that is based on a reference population. Here the reference population is all nine nations.

Looking at the estimated variance components, we can see that there is more unexplained variability between nations than between regions within each nation. This may be due to, for example, country-specific informational campaigns on the risks of sun exposure.



Stored results

meqrpoisson stores the following in **e()**:

Scalars

e(N)	number of observations
e(k)	number of parameters
e(k_f)	number of fixed-effects parameters
e(k_r)	number of random-effects parameters
e(k_rs)	number of variances
e(k_rc)	number of covariances
e(df_m)	model degrees of freedom
e(ll)	log likelihood
e(chi2)	χ^2
e(p)	significance
e(ll_c)	log likelihood, comparison model
e(chi2_c)	χ^2 , comparison test
e(df_c)	degrees of freedom, comparison test
e(p_c)	significance, comparison test
e(rank)	rank of e(V)
e(reparam_rc)	return code, final reparameterization
e(rc)	return code
e(converged)	1 if converged, 0 otherwise

Macros

e(cmd)	meqrpoisson
e(cmdline)	command as typed
e(depvar)	name of dependent variable
e(ivars)	grouping variables
e(exposurevar)	exposure variable
e(model)	Poisson
e(title)	title in estimation output
e(offset)	offset
e(redim)	random-effects dimensions
e(vartypes)	variance-structure types
e(revars)	random-effects covariates
e(n_quad)	number of integration points
e(laplace)	laplace, if Laplace approximation
e(chi2type)	Wald; type of model χ^2
e(vce)	bootstrap or jackknife if defined
e(vcetype)	title used to label Std. Err.
e(method)	ML
e(opt)	type of optimization
e(ml_method)	type of ml method
e(technique)	maximization technique
e(datasignature)	the checksum
e(datasignaturevars)	variables used in calculation of checksum
e(properties)	b V
e(estat_cmd)	program used to implement estat
e(predict)	program used to implement predict
e(marginsnotok)	predictions disallowed by margins
e(marginsdefault)	default predict() specification for margins
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved

Matrices

e(b)	coefficient vector
e(Cns)	constraints matrix
e(N_g)	group counts
e(g_min)	group-size minimums
e(g_avg)	group-size averages
e(g_max)	group-size maximums
e(V)	variance-covariance matrix of the estimators

Functions

e(sample)	marks estimation sample

Methods and formulas

In a two-level Poisson model, for cluster j , $j = 1, \dots, M$, the conditional distribution of $\mathbf{y}_j = (y_{j1}, \dots, y_{jn_j})'$, given a set of cluster-level random effects $\mathbf{u}_j$, is

$$\begin{aligned} f(\mathbf{y}_j | \mathbf{u}_j) &= \prod_{i=1}^{n_j} [\{\exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\}^{y_{ij}} \exp\{-\exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)\} / y_{ij}!] \\ &= \exp \left[\sum_{i=1}^{n_j} \{y_{ij}(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) - \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j) - \log(y_{ij}!)\} \right] \end{aligned}$$

Defining $c(\mathbf{y}_j) = \sum_{i=1}^{n_j} \log(y_{ij}!)$, where $c(\mathbf{y}_j)$ does not depend on the model parameters, we can express the above compactly in matrix notation,

$$f(\mathbf{y}_j | \mathbf{u}_j) = \exp \{ \mathbf{y}_j' (\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j) - \mathbf{1}' \exp(\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j) - c(\mathbf{y}_j) \}$$

where $\mathbf{X}_j$ is formed by stacking the row vectors $\mathbf{x}_{ij}$ and $\mathbf{Z}_j$ is formed by stacking the row vectors $\mathbf{z}_{ij}$. We extend the definition of $\exp(\cdot)$ to be a vector function where necessary.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{y}_j, \mathbf{u}_j)$,

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{y}_j | \mathbf{u}_j) \exp(-\mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2) d\mathbf{u}_j \\ &= \exp\{-c(\mathbf{y}_j)\} (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int \exp\{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \end{aligned} \quad (2)$$

where

$$h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) = \mathbf{y}_j' (\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j) - \mathbf{1}' \exp(\mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j) - \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j / 2$$

and for convenience, in the arguments of $h(\cdot)$ we suppress the dependence on the observable data $(\mathbf{y}_j, \mathbf{X}_j, \mathbf{Z}_j)$.

The integration in (2) has no closed form and thus must be approximated. The Laplacian approximation (Tierney and Kadane 1986; Pinheiro and Bates 1995) is based on a second-order Taylor expansion of $h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)$ about the value of $\mathbf{u}_j$ that maximizes it. Taking first and second derivatives, we obtain

$$\begin{aligned} h'(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) &= \frac{\partial h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)}{\partial \mathbf{u}_j} = \mathbf{Z}_j' \{\mathbf{y}_j - \mathbf{m}(\boldsymbol{\beta}, \mathbf{u}_j)\} - \boldsymbol{\Sigma}^{-1} \mathbf{u}_j \\ h''(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j) &= \frac{\partial^2 h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)}{\partial \mathbf{u}_j \partial \mathbf{u}_j'} = -\{\mathbf{Z}_j' \mathbf{V}(\boldsymbol{\beta}, \mathbf{u}_j) \mathbf{Z}_j + \boldsymbol{\Sigma}^{-1}\} \end{aligned}$$

where $\mathbf{m}(\boldsymbol{\beta}, \mathbf{u}_j)$ is the vector function with the i th element equal to the conditional mean of y_{ij} given $\mathbf{u}_j$, that is, $\exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)$. $\mathbf{V}(\boldsymbol{\beta}, \mathbf{u}_j)$ is the diagonal matrix whose diagonal entries v_{ij} are the conditional variances of y_{ij} given $\mathbf{u}_j$, namely,

$$v_{ij} = \exp(\mathbf{x}_{ij}\boldsymbol{\beta} + \mathbf{z}_{ij}\mathbf{u}_j)$$

because equality of mean and variance is a characteristic of the Poisson distribution.

The maximizer of $h(\beta, \Sigma, \mathbf{u}_j)$ is $\hat{\mathbf{u}}_j$ such that $h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{0}$. The integrand in (2) is proportional to the posterior density $f(\mathbf{u}_j | \mathbf{y}_j)$, so $\hat{\mathbf{u}}_j$ also represents the posterior mode, a plausible estimator of $\mathbf{u}_j$ in its own right.

Given the above derivatives, the second-order Taylor approximation then takes the form

$$h(\beta, \Sigma, \mathbf{u}_j) \approx h(\beta, \Sigma, \hat{\mathbf{u}}_j) + \frac{1}{2} (\mathbf{u}_j - \hat{\mathbf{u}}_j)' h''(\beta, \Sigma, \hat{\mathbf{u}}_j) (\mathbf{u}_j - \hat{\mathbf{u}}_j) \quad (3)$$

The first-derivative term vanishes because $h'(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{0}$. Therefore,

$$\begin{aligned} \int \exp \{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j &\approx \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j)\} \\ &\times \int \exp \left[-\frac{1}{2} (\mathbf{u}_j - \hat{\mathbf{u}}_j)' \{-h''(\beta, \Sigma, \hat{\mathbf{u}}_j)\} (\mathbf{u}_j - \hat{\mathbf{u}}_j) \right] d\mathbf{u}_j \quad (4) \\ &= \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j)\} (2\pi)^{q/2} | -h''(\beta, \Sigma, \hat{\mathbf{u}}_j) |^{-1/2} \end{aligned}$$

because the latter integrand can be recognized as the “kernel” of a multivariate normal density.

Combining the above with (2) (and taking logs) gives the Laplacian log-likelihood contribution of the j th cluster,

$$\mathcal{L}_j^{\text{Lap}}(\beta, \Sigma) = -\frac{1}{2} \log |\Sigma| - \log |\mathbf{R}_j| + h(\beta, \Sigma, \hat{\mathbf{u}}_j) - c(\mathbf{y}_j)$$

where $\mathbf{R}_j$ is an upper-triangular matrix such that $-h''(\beta, \Sigma, \hat{\mathbf{u}}_j) = \mathbf{R}_j \mathbf{R}_j'$. Pinheiro and Chao (2006) show that $\hat{\mathbf{u}}_j$ and $\mathbf{R}_j$ can be efficiently computed as the iterative solution to a least-squares problem by using matrix decomposition methods similar to those used in fitting LME models (Bates and Pinheiro 1998; Pinheiro and Bates 2000; [ME] mixed).

The fidelity of the Laplacian approximation is determined wholly by the accuracy of the approximation in (3). An alternative that does not depend so heavily on this approximation is integration via AGQ (Naylor and Smith 1982; Liu and Pierce 1994).

The application of AGQ to this particular problem is from Pinheiro and Bates (1995). When we reexamine the integral in question, a transformation of integration variables yields

$$\begin{aligned} \int \exp \{h(\beta, \Sigma, \mathbf{u}_j)\} d\mathbf{u}_j &= |\mathbf{R}_j|^{-1} \int \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1} \mathbf{t})\} d\mathbf{t} \\ &= (2\pi)^{q/2} |\mathbf{R}_j|^{-1} \int \exp \{h(\beta, \Sigma, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1} \mathbf{t}) + \mathbf{t}' \mathbf{t} / 2\} \phi(\mathbf{t}) d\mathbf{t} \quad (5) \end{aligned}$$

where $\phi(\cdot)$ is the standard multivariate normal density. Because the integrand is now expressed as some function multiplied by a normal density, it can be estimated by applying the rules of standard Gauss–Hermite quadrature. For a predetermined number of quadrature points N_Q , define $a_k = \sqrt{2} a_k^*$ and $w_k = w_k^* / \sqrt{\pi}$, for $k = 1, \dots, N_Q$, where (a_k^*, w_k^*) are a set of abscissas and weights for Gauss–Hermite quadrature approximations of $\int \exp(-x^2) f(x) dx$, as obtained from Abramowitz and Stegun (1964, 924).

Define $\mathbf{a}_k = (a_{k_1}, a_{k_2}, \dots, a_{k_q})'$; that is, $\mathbf{a}_k$ is a vector that spans the N_Q abscissas over the dimension q of the random effects. Applying quadrature rules to (5) yields the AGQ approximation,

$$\begin{aligned} & \int \exp \{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \mathbf{u}_j)\} d\mathbf{u}_j \\ & \approx (2\pi)^{q/2} |\mathbf{R}_j|^{-1} \sum_{k_1=1}^{N_Q} \cdots \sum_{k_q=1}^{N_Q} \left[\exp \{h(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \hat{\mathbf{u}}_j + \mathbf{R}_j^{-1} \mathbf{a}_k) + \mathbf{a}_k' \mathbf{a}_k / 2\} \prod_{p=1}^q w_{k_p} \right] \\ & \equiv (2\pi)^{q/2} \hat{G}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) \end{aligned}$$

resulting in the AGQ log-likelihood contribution of the j th cluster,

$$\mathcal{L}_j^{\text{AGQ}}(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = -\frac{1}{2} \log |\boldsymbol{\Sigma}| + \log \left\{ \hat{G}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) \right\} - c(\mathbf{y}_j)$$

The “adaptive” part of adaptive Gaussian quadrature lies in the translation and rescaling of the integration variables in (5) by using $\hat{\mathbf{u}}_j$ and $\mathbf{R}_j^{-1}$, respectively. This transformation of quadrature abscissas (centered at 0 in standard form) is chosen to better capture the features of the integrand, through which (4) can be seen to resemble a multivariate normal distribution with mean $\hat{\mathbf{u}}_j$ and variance $\mathbf{R}_j^{-1} \mathbf{R}_j^{-T}$. AGQ is therefore not as dependent as the Laplace method upon the approximation in (3). In AGQ, (3) serves merely to redirect the quadrature abscissas, with the AGQ approximation improving as the number of quadrature points, N_Q , increases. In fact, [Pinheiro and Bates \(1995\)](#) point out that AGQ with only one quadrature point ($a = 0$ and $w = 1$) reduces to the Laplacian approximation.

The log likelihood for the entire dataset is then simply the sum of the contributions of the M individual clusters, namely, $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \sum_{j=1}^M \mathcal{L}_j^{\text{Lap}}(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ for Laplace and $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \sum_{j=1}^M \mathcal{L}_j^{\text{AGQ}}(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ for AGQ.

Maximization of $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ is performed with respect to $(\boldsymbol{\beta}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a vector comprising the unique elements of the matrix square root of $\boldsymbol{\Sigma}$. This is done to ensure that $\boldsymbol{\Sigma}$ is always positive semidefinite. If the `matlog` option is specified, then $\boldsymbol{\theta}$ instead consists of the unique elements of the matrix logarithm of $\boldsymbol{\Sigma}$. For well-conditioned problems, both methods produce equivalent results, yet our experience deems the former as more numerically stable near the boundary of the parameter space.

Once maximization is achieved, parameter estimates are mapped from $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\theta}})$ to $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$, where $\hat{\boldsymbol{\gamma}}$ is a vector containing the unique (estimated) elements of $\boldsymbol{\Sigma}$, expressed as logarithms of standard deviations for the diagonal elements and hyperbolic arctangents of the correlations for off-diagonal elements. This last step is necessary to (a) obtain a parameterization under which parameter estimates can be displayed and interpreted individually, rather than as elements of a matrix square root (or logarithm), and (b) parameterize these elements such that their ranges each encompass the entire real line.

Parameter estimates are stored in `e(b)` as $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$, with the corresponding variance–covariance matrix stored in `e(V)`. Parameter estimates can be displayed in this metric by specifying the `estmetric` option. However, in `meqrpoisson` output, variance components are displayed as variances and covariances.

The approach outlined above can be extended from two-level models to models with three or more levels; see [Pinheiro and Chao \(2006\)](#) for details.

References

- Abramowitz, M., and I. A. Stegun, ed. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Washington, DC: National Bureau of Standards.
- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies. <http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Breslow, N. E., and D. G. Clayton. 1993. Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* 88: 9–25.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- Langford, I. H., G. Benthams, and A. McDonald. 1998. Multi-level modelling of geographically aggregated health data: A case study on malignant melanoma mortality and UV exposure in the European community. *Statistics in Medicine* 17: 41–57.
- Leyland, A. H., and H. Goldstein, ed. 2001. *Multilevel Modelling of Health Statistics*. New York: Wiley.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Liu, Q., and D. A. Pierce. 1994. A note on Gauss–Hermite quadrature. *Biometrika* 81: 624–629.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. *Mixture Models: Inference and Applications to Clustering*. New York: Dekker.
- Naylor, J. C., and A. F. M. Smith. 1982. Applications of a method for the efficient computation of posterior distributions. *Journal of the Royal Statistical Society, Series C* 31: 214–225.
- Palmer, T. M., C. M. Macdonald-Wallis, D. A. Lawlor, and K. Tilling. 2014. [Estimating adjusted associations between random effects from multilevel models: The reffadjust package](#). *Stata Journal* 14: 119–140.
- Pinheiro, J. C., and D. M. Bates. 1995. Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics* 4: 12–35.
- . 2000. *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- Pinheiro, J. C., and E. C. Chao. 2006. Efficient Laplacian and adaptive Gaussian quadrature algorithms for multilevel generalized linear mixed models. *Journal of Computational and Graphical Statistics* 15: 58–81.
- Rabe-Hesketh, S., and A. Skrondal. 2012. [Multilevel and Longitudinal Modeling Using Stata](#). 3rd ed. College Station, TX: Stata Press.
- Rabe-Hesketh, S., A. Skrondal, and A. Pickles. 2005. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics* 128: 301–323.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.
- Skrondal, A., and S. Rabe-Hesketh. 2004. [Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models](#). Boca Raton, FL: Chapman & Hall/CRC.
- Smans, M., C. S. Mair, and P. Boyle. 1993. *Atlas of Cancer Mortality in the European Economic Community*. Lyon, France: IARC Scientific Publications.

- Thall, P. F., and S. C. Vail. 1990. Some covariance models for longitudinal count data with overdispersion. *Biometrics* 46: 657–671.
- Tierney, L., and J. B. Kadane. 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* 81: 82–86.

Also see

- [ME] [meqrpoisson postestimation](#) — Postestimation tools for meqrpoisson
- [ME] [menbreg](#) — Multilevel mixed-effects negative binomial regression
- [ME] [mepoisson](#) — Multilevel mixed-effects Poisson regression
- [ME] [me](#) — Introduction to multilevel mixed-effects models
- [MI] [estimation](#) — Estimation commands for use with mi estimate
- [SEM] [intro 5](#) — Tour of models (*Multilevel mixed-effects models*)
- [XT] [xtpoisson](#) — Fixed-effects, random-effects, and population-averaged Poisson models
- [U] [20 Estimation and postestimation commands](#)

Postestimation commands	predict	margins
Remarks and examples	Methods and formulas	References
Also see		

Postestimation commands

The following postestimation commands are of special interest after `meqrpoisson`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat recovariance</code>	display the estimated random-effects covariance matrices
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike’s and Schwarz’s Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estimates</code>	cataloging estimation results
<code>hausman</code>	Hausman’s specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
<code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

predict

Description for predict

`predict` creates a new variable containing predictions such as mean responses; linear predictions; standard errors; and Pearson, deviance, and Anscombe residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [reses(newvarsspec)
      relevel(levelvar) ]
```

Syntax for obtaining other predictions

```
predict [type] newvar [if] [in] [, statistic nooffset fixedonly]
```

newvarsspec is *stub** or *newvarlist*.

<i>statistic</i>	Description

Main	
<code>mu</code>	predicted mean; the default
<code>xb</code>	linear predictor for the fixed portion of the model only
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code>pearson</code>	Pearson residuals
<code>deviance</code>	deviance residuals
<code>anscombe</code>	Anscombe residuals

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

Options for predict

Main

`reffects` calculates posterior modal estimates of the random effects. By default, estimates for all random effects in the model are calculated. However, if the `relevel`(*levelvar*) option is specified, then estimates for only level *levelvar* in the model are calculated. For example, if `classes` are nested within `schools`, then typing

```
. predict b*, reffects relevel(school)
```

would yield random-effects estimates at the school level. You must specify *q* new variables, where *q* is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you.

reses(*newvarsspec*) calculates standard errors for the random-effects estimates. By default, standard errors for all random effects in the model are calculated. However, if the **relevel**(*levelvar*) option is specified, then standard errors for only level *levelvar* in the model are calculated; see the **reffects** option.

You must specify q new variables, where q is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you. The new variables will have the same storage type as the corresponding random-effects variables.

The **reffects** and **reses**() options often generate multiple new variables at once. When this occurs, the random effects (or standard errors) contained in the generated variables correspond to the order in which the variance components are listed in the output of **meqrpoisson**. Still, examining the variable labels of the generated variables (with the **describe** command, for instance) can be useful in deciphering which variables correspond to which terms in the model.

relevel(*levelvar*) specifies the level in the model at which predictions for random effects and their standard errors are to be obtained. *levelvar* is the name of the model level and is either the name of the variable describing the grouping at that level or is **_all**, a special designation for a group comprising all the estimation data.

mu, the default, calculates the predicted mean, that is, the predicted count. By default, this is based on a linear predictor that includes both the fixed effects and the random effects, and the predicted mean is conditional on the values of the random effects. Use the **fixedonly** option (see [below](#)) if you want predictions that include only the fixed portion of the model, that is, if you want random effects set to 0.

xb calculates the linear prediction $\mathbf{x}\beta$ based on the estimated fixed effects (coefficients) in the model. This is equivalent to fixing all random effects in the model to their theoretical (prior) mean value of 0.

stdp calculates the standard error of the fixed-effects linear predictor $\mathbf{x}\beta$.

pearson calculates Pearson residuals. Pearson residuals large in absolute value may indicate a lack of fit. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

deviance calculates deviance residuals. Deviance residuals are recommended by [McCullagh and Nelder \(1989\)](#) as having the best properties for examining the goodness of fit of a GLM. They are approximately normally distributed if the model is correctly specified. They may be plotted against the fitted values or against a covariate to inspect the model's fit. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

anscombe calculates Anscombe residuals, which are designed to closely follow a normal distribution. By default, residuals include both the fixed portion and the random portion of the model. The **fixedonly** option modifies the calculation to include the fixed portion only.

nooffset is relevant only if you specified **offset**(*varname_o*) or **exposure**(*varname_e*) for **meqrpoisson**. It modifies the calculations made by **predict** so that they ignore the offset/exposure variable; the linear prediction is treated as $\mathbf{X}\beta + \mathbf{Z}\mathbf{u}$ rather than $\mathbf{X}\beta + \mathbf{Z}\mathbf{u} + \text{offset}$, or $\mathbf{X}\beta + \mathbf{Z}\mathbf{u} + \ln(\text{exposure})$, whichever is relevant.

fixedonly modifies predictions to include only the fixed portion of the model, equivalent to setting all random effects equal to 0; see the [mu](#) option.

margins

Description for margins

`margins` estimates margins of response for linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [options]
```

<i>statistic</i>	Description
<code>xb</code>	linear predictor for the fixed portion of the model only; the default
<code><u>r</u>effects</code>	not allowed with <code>margins</code>
<code>mu</code>	not allowed with <code>margins</code>
<code>stdp</code>	not allowed with <code>margins</code>
<code><u>p</u>earson</code>	not allowed with <code>margins</code>
<code><u>d</u>eviance</code>	not allowed with <code>margins</code>
<code><u>a</u>nscombe</code>	not allowed with <code>margins</code>

Statistics not allowed with `margins` are functions of stochastic quantities other than $e(b)$.

For the full syntax, see [R] [margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a Poisson mixed-effects model with `meqrpoisson`. For the most part, calculation centers around obtaining estimates of the subject/group-specific random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation.

► Example 1: Obtaining predicted counts and estimates of random effects

In [example 2](#) of [ME] [meqrpoisson](#), we modeled the number of deaths among males in nine European nations as a function of exposure to ultraviolet radiation (`uv`). We used a three-level Poisson model with random effects at the nation and region levels.

```
. use http://www.stata-press.com/data/r15/melanoma
(Skin cancer (melanoma) data)

. meqrpoisson deaths c.uv##c.uv, exposure(expected) || nation: || region:
(output omitted)
```

We can use `predict` to obtain the predicted counts as well as the estimates of the random effects at the nation and region levels.

```
. predict mu
(option mu assumed; predicted means)

. predict re_nat re_reg, reffects
```

Stata displays a note that the predicted values of `mu` are based on the posterior means of random effects. You can use the `modes` option to obtain predictions based on the posterior modes of random effects.

Here we list the data for the first nation in the dataset, which happens to be Belgium:

```
. list nation region deaths mu re_nat re_reg if nation==1, sepyby(region)
```

	nation	region	deaths	mu	re_nat	re_reg
1.	Belgium	1	79	69.53809	-.1195986	.3621477
2.	Belgium	2	80	78.28663	-.1195986	.0478318
3.	Belgium	2	51	46.30186	-.1195986	.0478318
4.	Belgium	2	43	54.35935	-.1195986	.0478318
5.	Belgium	2	89	66.90427	-.1195986	.0478318
6.	Belgium	2	19	34.89818	-.1195986	.0478318
7.	Belgium	3	19	8.199643	-.1195986	-.4348477
8.	Belgium	3	15	41.09571	-.1195986	-.4348477
9.	Belgium	3	33	30.90989	-.1195986	-.4348477
10.	Belgium	3	9	6.942485	-.1195986	-.4348477
11.	Belgium	3	12	12.21362	-.1195986	-.4348477

We can see that the predicted random effects at the nation level, `re_nat`, are the same for all the observations listed. Similarly, the predicted random effects at the region level, `re_reg`, are the same within each region.

□ Technical note

Out-of-sample predictions are permitted after `meqrpoisson`, but if these predictions involve estimated random effects, the integrity of the estimation data must be preserved. If the estimation data have changed since the model was fit, `predict` will be unable to obtain predicted random effects that are appropriate for the fitted model and will give an error. Thus to obtain out-of-sample predictions that contain random-effects terms, be sure that the data for these predictions are in observations that augment the estimation data.

Methods and formulas

Continuing the discussion in [Methods and formulas](#) of [ME] `meqrpoisson` and using the definitions and formulas defined there, we begin by considering the prediction of the random effects u_j for the j th cluster in a two-level model.

Given a set of estimated `meqrpoisson` parameters, $(\hat{\beta}, \hat{\Sigma})$, a profile likelihood in $\mathbf{u}_j$ is derived from the joint distribution $f(\mathbf{y}_j, \mathbf{u}_j)$ as

$$\mathcal{L}_j(\mathbf{u}_j) = \exp\{-c(\mathbf{y}_j)\} (2\pi)^{-q/2} |\hat{\Sigma}|^{-1/2} \exp\left\{g\left(\hat{\beta}, \hat{\Sigma}, \mathbf{u}_j\right)\right\} \quad (1)$$

The conditional maximum likelihood estimator of $\mathbf{u}_j$ —conditional on fixed $(\hat{\beta}, \hat{\Sigma})$ —is the maximizer of $\mathcal{L}_j(\mathbf{u}_j)$ or, equivalently, the value of $\hat{\mathbf{u}}_j$ that solves

$$\mathbf{0} = g'\left(\hat{\beta}, \hat{\Sigma}, \hat{\mathbf{u}}_j\right) = \mathbf{Z}'_j \left\{\mathbf{y}_j - \mathbf{m}(\hat{\beta}, \hat{\mathbf{u}}_j)\right\} - \hat{\Sigma}^{-1} \hat{\mathbf{u}}_j$$

Because (1) is proportional to the conditional density $f(\mathbf{u}_j|\mathbf{y}_j)$, you can also refer to $\hat{\mathbf{u}}_j$ as the conditional mode (or posterior mode if you lean toward Bayesian terminology). Regardless, you are referring to the same estimator.

Conditional standard errors for the estimated random effects are derived from standard theory of maximum likelihood, which dictates that the asymptotic variance matrix of $\hat{\mathbf{u}}_j$ is the negative inverse of the Hessian, which is estimated as

$$g''\left(\hat{\beta}, \hat{\Sigma}, \hat{\mathbf{u}}_j\right) = -\left\{\mathbf{Z}'_j \mathbf{V}(\hat{\beta}, \hat{\mathbf{u}}_j) \mathbf{Z}_j + \hat{\Sigma}^{-1}\right\}$$

Similar calculations extend to models with more than one level of random effects; see [Pinehiro and Chao \(2006\)](#).

For any observation i in the j th cluster in a two-level model, define the linear predictor as

$$\hat{\eta}_{ij} = \mathbf{x}_{ij} \hat{\beta} + \mathbf{z}_{ij} \hat{\mathbf{u}}_j$$

In a three-level model, for the i th observation within the j th level-two cluster within the k th level-three cluster,

$$\hat{\eta}_{ijk} = \mathbf{x}_{ijk} \hat{\beta} + \mathbf{z}_{ijk}^{(3)} \hat{\mathbf{u}}_k^{(3)} + \mathbf{z}_{ijk}^{(2)} \hat{\mathbf{u}}_{jk}^{(2)}$$

where $\mathbf{z}^{(p)}$ and $\mathbf{u}^{(p)}$ refer to the level p design variables and random effects, respectively. For models with more than three levels, the definition of $\hat{\eta}$ extends in the natural way, with only the notation becoming more complicated.

If the `fixedonly` option is specified, $\hat{\eta}$ contains the linear predictor for only the fixed portion of the model, for example, in a two-level model $\hat{\eta}_{ij} = \mathbf{x}_{ij} \hat{\beta}$. In what follows, we assume a two-level model, with the only necessary modification for multilevel models being the indexing.

The predicted mean conditional on the random effects $\hat{\mathbf{u}}_j$ is

$$\hat{\mu}_{ij} = \exp(\hat{\eta}_{ij})$$

Pearson residuals are calculated as

$$\nu_{ij}^P = \frac{y_{ij} - \hat{\mu}_{ij}}{\{V(\hat{\mu}_{ij})\}^{1/2}}$$

for $V(\hat{\mu}_{ij}) = \hat{\mu}_{ij}$.

Deviance residuals are calculated as

$$\nu_{ij}^D = \text{sign}(y_{ij} - \hat{\mu}_{ij}) \sqrt{\hat{d}_{ij}^2}$$

where

$$\hat{d}_{ij}^2 = \begin{cases} 2\hat{\mu}_{ij} & \text{if } y_{ij} = 0 \\ 2 \left\{ y_{ij} \log \left(\frac{y_{ij}}{\hat{\mu}_{ij}} \right) - (y_{ij} - \hat{\mu}_{ij}) \right\} & \text{otherwise} \end{cases}$$

Anscombe residuals are calculated as

$$\nu_{ij}^A = \frac{3 \left(y_{ij}^{2/3} - \hat{\mu}_{ij}^{2/3} \right)}{2\hat{\mu}_{ij}^{1/6}}$$

For a discussion of the general properties of the above residuals, see [Hardin and Hilbe \(2012, chap. 4\)](#).

References

- Hardin, J. W., and J. M. Hilbe. 2012. *Generalized Linear Models and Extensions*. 3rd ed. College Station, TX: Stata Press.
- McCullagh, P., and J. A. Nelder. 1989. *Generalized Linear Models*. 2nd ed. London: Chapman & Hall/CRC.
- Pinheiro, J. C., and E. C. Chao. 2006. Efficient Laplacian and adaptive Gaussian quadrature algorithms for multilevel generalized linear mixed models. *Journal of Computational and Graphical Statistics* 15: 58–81.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.

Also see

[ME] **meqrpoisson** — Multilevel mixed-effects Poisson regression (QR decomposition)

[U] **20 Estimation and postestimation commands**

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`mestreg` fits a mixed-effects parametric survival-time model. The conditional distribution of the response given the random effects is assumed to be an exponential, Weibull, lognormal, loglogistic, or gamma distribution. `mestreg` can be used with single- or multiple-record `st` data.

Quick start

Without weights

Two-level Weibull survival model with covariates `x1` and `x2` and random intercepts by `lev2` using `stset` data

```
mestreg x1 x2 || lev2:, distribution(weibull)
```

Mixed-effects model adding random coefficients for `x1`

```
mestreg x1 x2 || lev2:x1, distribution(weibull)
```

Three-level random-intercept model with `lev2` nested within `lev3`

```
mestreg x1 x2 || lev3: || lev2:, distribution(weibull)
```

With weights

Two-level Weibull survival model with covariates `x1` and `x2`, random intercepts by `lev2`, and observation-level frequency weights `wvar1` using `stset` data

```
mestreg x1 x2 [fweight=wvar1] || lev2:, distribution(weibull)
```

Two-level random-intercept model from a two-stage sampling design with PSUs identified by `psu` using PSU-level and observation-level sampling weights `wvar2` and `wvar1`

```
mestreg x1 x2 [pweight=wvar1] || psu:, pweight(wvar2)
```

Same as above, but `svyset` the data first

```
svyset psu, weight(wvar2) || _n, weight(wvar1)  
svy: mestreg x1 x2 || psu:, distribution(weibull)
```

Note: Any supported parametric survival [distribution](#) may be specified in place of `weibull` above.

Menu

Statistics > Multilevel mixed-effects models > Parametric survival regression

Syntax

```
mestreg fe_equation [|| re_equation] [|| re_equation ...],
      distribution(distname) [options]
```

where the syntax of *fe_equation* is

```
[indepvars] [if] [in] [weight] [, fe_options]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [, re_options]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code><u>noconstant</u></code>	suppress constant term from the fixed-effects equation
<code><u>offset</u>(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code><u>covariance</u>(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code><u>noconstant</u></code>	suppress constant term from the random-effects equation
<code><u>fweight</u>(<i>varname</i>)</code>	frequency weights at higher levels
<code><u>iweight</u>(<i>varname</i>)</code>	importance weights at higher levels
<code><u>pweight</u>(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
* <u>d</u> istribution(<i>distname</i>)	specify survival distribution
<u>t</u> ime	use accelerated failure-time metric
<u>c</u> onstraints(<i>constraints</i>)	apply specified linear constraints
<u>c</u> ollinear	keep collinear variables
SE/Robust	
<u>v</u> ce(<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code>
Reporting	
<u>l</u> evel(#)	set confidence level; default is <code>level(95)</code>
<u>n</u> ohr	do not report hazard ratios
<u>t</u> ratio	report time ratios
<u>n</u> oshow	do not show st setting information
<u>n</u> ocnsreport	do not display constraints
<u>n</u> otable	suppress coefficient table
<u>n</u> oheader	suppress output header
<u>n</u> ogroup	suppress table summarizing groups
<i>display_options</i>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<u>i</u> ntmethod(<i>intmethod</i>)	integration method
<u>i</u> ntpoints(#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<i>maximize_options</i>	control the maximization process; seldom used
<u>s</u> tartvalues(<i>svmethod</i>)	method for obtaining starting values
<u>s</u> tartgrid[(<i>gridspec</i>)]	perform a grid search to improve starting values
<u>n</u> oestimate	do not fit the model; show starting values instead
<u>d</u> numerical	use numerical derivative techniques
<u>c</u> oeflegend	display legend instead of statistics

* `distribution(distname)` is required.

<i>vartype</i>	Description
<u>i</u> ndependent	one unique variance parameter per random effect, all covariances 0; the default unless the R. notation is used
<u>e</u> xchangeable	equal variances for random effects, and one common pairwise covariance
<u>i</u> dentify	equal variances for random effects, all covariances 0; the default if the R. notation is used
<u>u</u> nstructured	all variances and covariances to be distinctly estimated
<u>f</u> ixed(<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<u>p</u> attern(<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>distname</i>	Description
<u>e</u> xponential	exponential survival distribution
<u>l</u> oglogistic	loglogistic survival distribution
<u>l</u> logistic	synonym for loglogistic
<u>w</u> eibull	Weibull survival distribution
<u>l</u> ognormal	lognormal survival distribution
<u>l</u> normal	synonym for lognormal
<u>g</u> amma	gamma survival distribution

<i>intmethod</i>	Description
<u>m</u> vaghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>m</u> caghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

You must `stset` your data before using `mestreg`; see [ST] [stset](#).
indepvars may contain factor variables; see [U] [11.4.3 Factor variables](#).
`bayes`, `by`, and `svy` are allowed; see [U] [11.1.10 Prefix commands](#). For more details, see [BAYES] [bayes: mestreg](#).
`vce()` and weights are not allowed with the `svy` prefix; see [SVY] [svy](#).
`fweights`, `iweights`, and `pweights` are allowed; see [U] [11.1.6 weight](#). Only one type of weight may be specified.
Weights are not supported under the Laplacian approximation or for crossed models.
`startvalues()`, `startgrid`, `noestimate`, `dnumerical`, and `coeflegend` do not appear in the dialog box.
See [U] [20 Estimation and postestimation commands](#) for more capabilities of estimation commands.

Options

Model

`noconstant` suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

`offset(varname)` specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

`covariance(vartype)` specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: `independent`, `exchangeable`, `identity`, `unstructured`, `fixed(matname)`, or `pattern(matname)`.

`covariance(independent)` covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is `covariance(independent)` unless a crossed random-effects model is fit, in which case the default is `covariance(identity)`.

`covariance(exchangeable)` structure specifies one common variance for all random effects and one common pairwise covariance.

`covariance(identity)` is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

`covariance(unstructured)` allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p+1)/2$ unique parameters.

`covariance(fixed(matname))` and `covariance(pattern(matname))` covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if $matname[i, j] = matname[k, l]$.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`distribution(distname)` specifies the survival model to be fit. *distname* is one of the following: `exponential`, `loglogistic`, `llogistic`, `weibull`, `lognormal`, `lnormal`, or `gamma`. This option is required.

`time` specifies that the model be fit in the accelerated failure-time metric rather than in the log relative-hazard metric. This option is valid only for the exponential and Weibull models because these are the only models that have both a proportional-hazards and an accelerated failure-time parameterization. Regardless of metric, the likelihood function is the same, and models are equally appropriate in either metric; it is just a matter of changing interpretation.

`time` must be specified at estimation.

`constraints(constraints)`, `collinear`; see [R] [estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`; see [R] [estimation options](#).

`nohr`, which may be specified at estimation or upon redisplaying results, specifies that coefficients rather than exponentiated coefficients be displayed, that is, that coefficients rather than hazard ratios be displayed. This option affects only how coefficients are displayed, not how they are estimated.

This option is valid only for the exponential and Weibull models because they have a natural proportional-hazards parameterization. These two models, by default, report hazards ratios (exponentiated coefficients).

`tratio` specifies that exponentiated coefficients, which are interpreted as time ratios, be displayed. `tratio` is appropriate only for the loglogistic, lognormal, and gamma models or for the exponential and Weibull models when fit in the accelerated failure-time metric.

`tratio` may be specified at estimation or upon replay.

`noshow` prevents `mestreg` from showing the key `st` variables. This option is rarely used because most users type `stset`, `show` or `stset, noshow` to set once and for all whether they want to see these variables mentioned at the top of the output of every `st` command; see [ST] [stset](#).

`nocnsreport`; see [R] [estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `mestreg` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `mestreg` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[gridspec]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

For a general introduction to `me` commands, see [ME] [me](#).

Remarks are presented under the following headings:

[Introduction](#)

[Two-level models](#)

[Three-level models](#)

Introduction

Mixed-effects survival models contain both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

mestreg allows for many levels of random effects. However, for simplicity, we now consider two-level models, where we have a series of M independent clusters and a set of random effects $\mathbf{u}_j$ corresponding to those clusters. Two often-used models for adjusting survivor functions for the effects of covariates are the accelerated failure-time (AFT) model and the multiplicative or proportional hazards (PH) model.

In the AFT model, the natural logarithm of the survival time, $\log t$, is expressed as a linear function of the covariates; when we incorporate random-effects, this yields the model

$$\log t_{ji} = \mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j + v_{ji}$$

for $j = 1, \dots, M$ clusters, with cluster j consisting of $i = 1, \dots, n_j$ observations. The $1 \times p$ row vector $\mathbf{x}_{ji}$ contains the covariates for the fixed effects, with regression coefficients (fixed effects) $\boldsymbol{\beta}$.

The $1 \times q$ vector $\mathbf{z}_{ji}$ contains the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercept model, $\mathbf{z}_{ji}$ is simply the scalar 1. The random effects $\mathbf{u}_j$ are M realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components.

Finally, v_{ji} are the observation-level errors with density $\varphi(\cdot)$. The distributional form of the error term determines the regression model. Five regression models are implemented in **mestreg** using the AFT parameterization: exponential, gamma, loglogistic, lognormal, and Weibull. The lognormal regression model is obtained by letting $\varphi(\cdot)$ be the normal density. Similarly, by letting $\varphi(\cdot)$ be the logistic density, one obtains the loglogistic regression. Setting $\varphi(\cdot)$ equal to the extreme-value density yields the exponential and the Weibull regression models.

In the PH models fit by **mestreg**, the covariates have a multiplicative effect on the hazard function

$$h(t_{ji}) = h_0(t_{ji}) \exp(\mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j)$$

for some baseline hazard function $h_0(t)$. For the **mestreg** command, $h_0(t)$ is assumed to be parametric. The exponential and Weibull models are implemented in **mestreg** for the PH parameterization. These two models are implemented using both the AFT and PH parameterizations.

mestreg is suitable only for data that have been **stset**. By using **stset** on your data, you define the variables `_t0`, `_t`, and `_d`, which serve as the trivariate response variable (t_0, t, d) . Each response corresponds to a period under observation, $(t_0, t]$, resulting in either failure ($d = 1$) or right-censoring ($d = 0$) at time t .

mestreg does not allow delayed entry or gaps. However, **mestreg** is appropriate for data exhibiting multiple records per subject and time-varying covariates. **mestreg** requires subjects to be nested within clusters.

stset weights are not used; instead, weights must be specified at estimation. Weights are not allowed with crossed models or the Laplacian approximation. See [Survey estimation](#) in *Methods and formulas* for details.

Two-level models

► Example 1: Two-level random-intercept PH model

In [example 11](#) of [\[ST\] streg](#), we fit a Weibull model with an inverse-Gaussian shared frailty to the recurrence times for catheter-insertion point infection for 38 kidney dialysis patients. In this example, the subjects are the catheter insertions, not the patients themselves. This is a function of how the data were recorded—the onset of risk occurs at the time the catheter is inserted and not, say, at the time of admission of the patient into the study. Thus we have two subjects (insertions) within each group (patient). Each catheter insertion results in either infection (`infect==1`) or right-censoring (`infect==0`). The `stset` results are shown below.

```
. use http://www.stata-press.com/data/r15/catheter
(Kidney data, McGilchrist and Aisbett, Biometrics, 1991)
. stset
-> stset time, failure(infect)

      failure event:  infect != 0 & infect < .
obs. time interval:  (0, time]
exit on or before:   failure
```

76	total observations	
0	exclusions	

76	observations remaining, representing	
58	failures in single-record/single-failure data	
7,424	total analysis time at risk and under observation	
	at risk from t =	0
	earliest observed entry t =	0
	last observed exit t =	562

While it is reasonable to assume independence of patients, we would not want to assume that recurrence times within each patient are independent. The model used in [\[ST\] streg](#) allowed us to model the correlation by assuming that it was the result of a latent patient-level effect, or frailty.

The random-effects approach used by `mestreg` is more flexible because it allows you to experiment with several levels of random effects, including random coefficients, or both. You can then choose the model that best suits your data. Here we use `mestreg` to fit a random-effects Weibull model with normally distributed random effects. This model can be viewed as a shared frailty model with lognormal frailty.

```
. mestreg age female || patient:, distribution(weibull)
      failure _d: infect
      analysis time _t: time
Fitting fixed-effects model:
Iteration 0:   log likelihood = -1700989.9
Iteration 1:   log likelihood = -440.1998
Iteration 2:   log likelihood = -336.62162
Iteration 3:   log likelihood = -334.64937
Iteration 4:   log likelihood = -334.57959
Iteration 5:   log likelihood = -334.57944
Iteration 6:   log likelihood = -334.57944
Refining starting values:
Grid node 0:   log likelihood = -336.03604
Fitting full model:
Iteration 0:   log likelihood = -336.03604 (not concave)
Iteration 1:   log likelihood = -333.14043
Iteration 2:   log likelihood = -330.40952
Iteration 3:   log likelihood = -329.89242
Iteration 4:   log likelihood = -329.87847
Iteration 5:   log likelihood = -329.87832
Iteration 6:   log likelihood = -329.87832
Mixed-effects Weibull PH regression
Group variable:      patient
Number of obs       =      76
Number of groups    =      38
Obs per group:
      min =      2
      avg =     2.0
      max =      2
Integration method: mvaghermite
Integration pts.    =      7
Wald chi2(2)       =     10.12
Prob > chi2        =     0.0063
Log likelihood = -329.87832
```

_t	Haz. Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
age	1.007348	.013788	0.53	0.593	.9806828	1.034737
female	.1904727	.099992	-3.16	0.002	.0680737	.5329493
_cons	.0072901	.0072274	-4.96	0.000	.0010444	.0508881
/ln_p	.2243233	.1402795			-.0506195	.4992661
patient var(_cons)	.8308495	.4978425			.256735	2.688808

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline hazard (conditional on zero random effects).
LR test vs. Weibull model: chibar2(01) = 9.40 Prob >= chibar2 = 0.0011

The results are similar to those in [\[ST\] streg](#). The likelihood-ratio test compares the random-effects model with a survival model with fixed-effects only. The results support the random-effects model.

By default, when fitting a model with the PH parameterization, **mestreg** displays exponentiated coefficients, labeled as hazard ratios. These hazard ratios should be interpreted as “conditional hazard ratios”, that is, conditional on the random effects.

For example, the hazard ratio for **age** is 1.01. This means that according to the model, for a given patient, the hazard would increase 1% with each year of age. However, at the population level, marginal hazards corresponding to different levels of the covariates are not necessarily proportional. [Example 5](#) in [\[ME\] mestreg postestimation](#) illustrates this point with simulated data.

The exponentiated coefficients of covariates that usually remain constant within a group do not have a natural interpretation as conditional hazard ratios. However, the magnitude of the exponentiated coefficients always gives an idea of the effect of the covariates. In this example, `female` is constant within the group. The estimated hazard ratio for `female` is 0.19, which indicates that hazard functions for females tend to be smaller than hazard functions for males. Both conditional and unconditional predictions can be obtained with `predict`. Unconditional predictions can be visualized by using `stcurve`. Unconditional effects can be tested and visualized by using `margins` and `marginsplot`. See [example 1](#) in [\[ME\] mestreg postestimation](#) for an example using `predict`, `margins`, and `marginsplot`.



► Example 2: Two-level random-intercept AFT model

Although the PH parameterization is more popular in the literature because the output is easier to interpret, the AFT parameterization is useful when we need to make comparisons with other models that have only an AFT parameterization. For example, we might want to compare the Weibull results from [example 1](#) with the results from a gamma model.

Let's redisplay the results of a Weibull PH model from [example 1](#) as coefficients:

```
. mestreg, nohr
Mixed-effects Weibull PH regression          Number of obs      =          76
Group variable:      patient                Number of groups   =          38
                                           Obs per group:
                                           min =              2
                                           avg =             2.0
                                           max =              2
Integration method: mvaghermite              Integration pts.   =           7
Log likelihood = -329.87832                  Wald chi2(2)      =         10.12
                                           Prob > chi2       =         0.0063
```

_t	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
age	.0073207	.0136874	0.53	0.593	-.0195062	.0341476
female	-1.658247	.5249676	-3.16	0.002	-2.687164	-.629329
_cons	-4.921236	.9914009	-4.96	0.000	-6.864346	-2.978126
/ln_p	.2243233	.1402795			-.0506195	.4992661
patient var(_cons)	.8308495	.4978425			.256735	2.688808

```
LR test vs. Weibull model: chibar2(01) = 9.40          Prob >= chibar2 = 0.0011
```

We can refit the Weibull model using the AFT parameterization by specifying option `time`.

```
. mestreg age female || patient:, distribution(weibull) time
      failure _d:  infect
      analysis time _t:  time
Fitting fixed-effects model:
Iteration 0:  log likelihood = -346.46486
Iteration 1:  log likelihood = -343.29515
Iteration 2:  log likelihood = -335.0513
Iteration 3:  log likelihood = -334.58308
Iteration 4:  log likelihood = -334.57944
Iteration 5:  log likelihood = -334.57944
Refining starting values:
Grid node 0:  log likelihood = -335.10428
Fitting full model:
Iteration 0:  log likelihood = -335.10428
Iteration 1:  log likelihood = -332.13546
Iteration 2:  log likelihood = -330.01623
Iteration 3:  log likelihood = -329.88013
Iteration 4:  log likelihood = -329.87832
Iteration 5:  log likelihood = -329.87832
Mixed-effects Weibull AFT regression
Group variable:      patient
Number of obs      =      76
Number of groups   =      38
Obs per group:
      min =      2
      avg =     2.0
      max =      2
Integration method: mvaghermite
Integration pts.   =      7
Wald chi2(2)      =     13.00
Prob > chi2       =     0.0015
Log likelihood = -329.87832
```

_t	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
age	-.0058496	.010872	-0.54	0.591	-.0271585	.0154592
female	1.325034	.3719102	3.56	0.000	.596103	2.053964
_cons	3.932346	.5663757	6.94	0.000	2.82227	5.042422
/ln_p	.2243237	.1402794			-.0506189	.4992663
patient var(_cons)	.5304902	.2343675			.2231626	1.261053

```
LR test vs. Weibull model: chibar2(01) = 9.40      Prob >= chibar2 = 0.0011
```

The estimates of coefficients and variance components are different between the two models. In fact, the coefficients have the opposite signs. This is expected because the two models have different parameterizations. The relationship between the coefficients and variances of the two parameterizations for the Weibull model is

$$\beta_{PH} = -p \times \beta_{AFT}$$
$$\text{var}_{PH} = p^2 \times \text{var}_{AFT}$$

where p denotes the ancillary parameter. It is estimated in the logarithmic metric and is displayed in the output as `/ln_p`.

For example, we could calculate β_{PH} for `female` as approximately $-\exp(0.22) \times 1.33 = -1.66$. If we exponentiate this to obtain the hazard ratio that was reported in [example 1](#), we obtain the same reported result, 0.19.

For a discussion of the differences between the PH and AFT parameterizations, see, for example, Cleves, Gould, and Marchenko (2016).

Now, we can compare the results from our Weibull specification with the results from a gamma specification.

```
. mestreg age female || patient:, distribution(gamma)
      failure _d:  infect
      analysis time _t:  time
Fitting fixed-effects model:
Iteration 0:   log likelihood = -351.17349
Iteration 1:   log likelihood = -337.04571
Iteration 2:   log likelihood = -335.10167
Iteration 3:   log likelihood = -335.09115
Iteration 4:   log likelihood = -335.09115
Refining starting values:
Grid node 0:   log likelihood = -334.49759
Fitting full model:
Iteration 0:   log likelihood = -334.49759
Iteration 1:   log likelihood = -331.87827
Iteration 2:   log likelihood = -329.64795
Iteration 3:   log likelihood = -329.52682
Iteration 4:   log likelihood = -329.52635
Iteration 5:   log likelihood = -329.52634
Mixed-effects gamma AFT regression
Group variable:      patient
Number of obs       =      76
Number of groups    =      38
Obs per group:
      min =      2
      avg =     2.0
      max =      2
Integration method: mvaghermite
Integration pts.    =      7
Wald chi2(2)       =     13.23
Prob > chi2        =     0.0013
Log likelihood = -329.52634
```

_t	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
age	-.0060276	.0108267	-0.56	0.578	-.0272475	.0151924
female	1.324745	.3685132	3.59	0.000	.6024726	2.047018
_cons	3.873854	.5628993	6.88	0.000	2.770592	4.977117
/logs	-.1835075	.1008892			-.3812467	.0142317
patient var(_cons)	.5071823	.2241959			.213254	1.206232

```
LR test vs. gamma model: chibar2(01) = 11.13      Prob >= chibar2 = 0.0004
```

The coefficients and the random-effects variance are very similar for the two AFT models.

We can compare the marginal distributions or hazard functions for the two models by using `stcurve`; see example 2 in [ME] [mestreg postestimation](#).

➤ **Example 3: Two-level random-slope model**

In this example, we use a modified form of the dataset from [Rabe-Hesketh and Skrondal \(2012, chap. 15.7\)](#), previously published in [Danahy et al. \(1977\)](#) and analyzed by Pickles and Crouchley (1994,1995) and [Rabe-Hesketh, Skrondal, and Pickles \(2004\)](#).

`angina.dta` includes data on 21 patients with coronary heart disease who participated in a randomized crossover trial comparing a drug to prevent angina (chest pain) with a placebo. The participants are identified by `pid`.

Before receiving the drug (or placebo), participants were asked to exercise on exercise bikes to the onset of angina or, if angina did not occur, to exhaustion. The exercise time, `seconds`, and the result of the exercise, `angina`—angina (`angina=1`) or exhaustion (`angina=0`)—were recorded. The drug (`treat=1`) or placebo (`treat=0`) was then taken orally, and the exercise test was repeated one, three, and five hours (variable `occasion`) after drug or placebo administration. Because each exercise test can have a failure (the occurrence of angina), the test is the subject. Each test is identified by `tid`. Failure is indicated by the variable `angina`. In this case, we have eight repeated measures per study participant.

Before fitting the model, we `stset` our data:

```
. use http://www.stata-press.com/data/r15/angina
(Angina drug data, Rabe-Hesketh and Skrondal, 2012, ch 15.7)
. stset seconds, failure(angina) id(tid)

      id:  tid
failure event:  angina != 0 & angina < .
obs. time interval:  (seconds[_n-1], seconds]
exit on or before:  failure
```

168	total observations	
0	exclusions	
168	observations remaining, representing	
168	subjects	
155	failures in single-failure-per-subject data	
47,267	total analysis time at risk and under observation	
	at risk from t =	0
	earliest observed entry t =	0
	last observed exit t =	743

To reiterate, we specify `seconds` as the time variable, `angina` as the failure variable, and `tid` as the variable identifying multiple observations per test.

[Rabe-Hesketh and Skrondal \(2012\)](#) apply several models to this dataset, including a lognormal model and a Cox model with random effects. We fit a Weibull model with covariates `occasion` and `treat` and interaction between `occasion` and `treat`. We include a random effect at the subject level.

```

. mestreg occasion##treat || pid:, distribution(weibull)
      failure _d: angina
      analysis time _t: seconds
      id: tid
note: 1.occasion#1.treat identifies no observations in the sample
note: 4.occasion#1.treat omitted because of collinearity
(output omitted)
Mixed-effects Weibull PH regression      Number of obs      =      168
Group variable:      pid                  Number of groups    =      21
                                           Obs per group:
                                           min =      8
                                           avg =     8.0
                                           max =      8
Integration method: mvaghermite          Integration pts.    =      7
                                           Wald chi2(6)       =     78.14
Log likelihood = -885.67135               Prob > chi2        =     0.0000

```

_t	Haz. Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
occasion						
2	.719456	.2031744	-1.17	0.244	.4136423	1.251364
3	.902988	.2542476	-0.36	0.717	.5200146	1.568009
4	1.264262	.3516347	0.84	0.399	.7329746	2.180648
treat						
Drug	.3825531	.128784	-2.85	0.004	.1977608	.7400195
occasion#						
treat						
1#Drug	1 (empty)					
2#Drug	.1576401	.0804767	-3.62	0.000	.0579589	.4287586
3#Drug	.4512793	.2127706	-1.69	0.091	.1791093	1.137032
4#Drug	1 (omitted)					
_cons	4.90e-13	9.98e-13	-13.91	0.000	9.03e-15	2.66e-11
/ln_p	1.640297	.0689544			1.505149	1.775445
pid						
var(_cons)	4.529641	1.544175			2.322124	8.835725

Note: Estimates are transformed only in the first equation.

Note: _cons estimates baseline hazard (conditional on zero random effects).

LR test vs. Weibull model: chibar2(01) = 177.40 Prob >= chibar2 = 0.0000

Because individuals were exercising without the administration of a placebo or treatment at the first occasion (occasion==1), the category for interaction between occasion==1 and treat==1 is empty.

The estimated variance at the individual level (that is, the variance between individuals) is equal to 4.53. The likelihood-ratio test shows evidence in favor of the random-effects model versus the fixed-effects model.

The parameter $p = \exp(1.640297) = 5.16$, which is larger than 1. This means that the estimated hazard (conditional on the covariates and on the random effects) is a monotonically increasing function if we assume a Weibull distribution.

The model contains interaction terms for `occasion` and `treat`. Interpretation of interaction terms is usually less straightforward. Briefly, to interpret the exponentiated coefficients as conditional hazard ratios, we need to examine all the covariates in the interaction. The hazard function for `pid = j`, when we set `occasion = k` and `treat = l`, will be

$$h(t) = h_0(t) \times \exp(\beta_{\text{occ}_k} + \beta_{\text{treat}_l} + \beta_{\text{occ}_k \times \text{treat}_l} + \text{\_cons} + u_j)$$

where β_{occ_k} , β_{treat_l} , and $\beta_{\text{occ}_k \times \text{treat}_l}$ are, respectively, the coefficients for the dummies for `occasion = k` and `treat = l` and the interaction (`occasion = k × treatment = l`).

For example, when `treat = 0`, the hazard function is

$$h(t|\text{treat} = 0, \text{occasion} = k, \text{pid} = j) = h_0(t) \times \exp(\beta_{\text{occ}_k} + \text{\_cons} + u_j)$$

where β_{occ_1} is equal to 0 because `occasion = 1` is the base category. This means that for a given `pid`,

$$\frac{h(t|\text{treat} = 0, \text{occ} = k, \text{pid} = j)}{h(t|\text{treat} = 0, \text{occ} = 1, \text{pid} = j)} = \exp(\beta_{\text{occ}_k})$$

Notice that this is only true within `pid`, because different participants have different u_j s.

The coefficients have already been exponentiated, so we can see clearly that according to this model, when there is no treatment, the hazard for occasion 2 is smaller than the hazard for occasion 1. The increasing ratios indicate that the hazard increases with the occasion. Similar calculations could be performed for other interaction terms.

The easiest way to interpret models with interactions is by using `margins` and `marginsplot`, which allow us to compute and then visualize unconditional predictions and marginal effects. See [\[R\] margins](#) for more information.

Above we assumed a constant treatment effect for all individuals for each occasion. However, we may instead believe that the treatment effect varies also with individuals. This can be modeled by adding a random coefficient for the treatment, `treat`, at the individual level; we also include the `covariance(unstructured)` option to estimate a covariance term between the random intercept and the random slope for `treat`.

```
. mestreg occasion##treat || pid: treat, distribution(weibull)
> covariance(unstructured)
      failure _d:  angina
      analysis time _t:  seconds
      id:  tid
note: 1.occasion#1.treat identifies no observations in the sample
note: 4.occasion#1.treat omitted because of collinearity
```

(output omitted)

```
Mixed-effects Weibull PH regression      Number of obs      =      168
Group variable:      pid                  Number of groups   =      21
                                           Obs per group:
                                           min =      8
                                           avg  =     8.0
                                           max  =      8
Integration method: mvaghermite          Integration pts.   =      7
Log likelihood = -859.50038                Wald chi2(6)      =     50.18
                                           Prob > chi2       =     0.0000
```

_t	Haz. Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
occasion						
2	.5993591	.1861745	-1.65	0.099	.3260503	1.101766
3	.8643306	.2560242	-0.49	0.623	.483665	1.544597
4	1.333201	.3843218	1.00	0.318	.7577392	2.345694
treat						
Drug	.2147751	.1280091	-2.58	0.010	.0667814	.6907365
occasion# treat						
1#Drug	1 (empty)					
2#Drug	.1594337	.0885644	-3.31	0.001	.0536714	.4736058
3#Drug	.4632936	.2273925	-1.57	0.117	.1770402	1.212385
4#Drug	1 (omitted)					
_cons	6.21e-17	1.75e-16	-13.20	0.000	2.44e-19	1.58e-14
/ln_p	1.91931	.0736166			1.775024	2.063596
pid						
var(treat)	4.682507	1.956897			2.064178	10.62208
var(_cons)	6.939041	2.372975			3.549852	13.56403
pid						
cov(treat, _cons)	1.73782	1.313054	1.32	0.186	-.8357182	4.311357

Note: Estimates are transformed only in the first equation.

Note: _cons estimates baseline hazard (conditional on zero random effects).

LR test vs. Weibull model: chi2(3) = 229.74 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

We obtain somewhat different estimates of hazard ratios, but our inferential conclusions remain the same. We now observe two variances in the output, the variance for the intercept at the individual level and the variance for the coefficient for treatment at the individual level. The variance for the intercept is smaller because some of the variability is now explained by varying coefficients for treatment. The covariance is positive, meaning that the random slope tends to be larger for individuals who have a larger random intercept. See [example 4](#) in [\[ME\] mestreg postestimation](#) for an application of `predict` that presents a graphical analysis of this relationship.



Three-level models

➤ Example 4: Three-level random-slope model

[Blossfeld, Golsch, and Rohwer \(2007\)](#) analyze a dataset based on the German Life History Study of [Mayer and Brückner \(1989\)](#), collected in the years 1981–1983. The `jobhistory` dataset contains a modified version of [Blossfeld, Golsch, and Rohwer](#)’s anonymization of a random sample of 201 respondents from the original data. Each of the 600 observations in the dataset corresponds to a job episode. Variable `id` contains identification of the individual, `tstart` contains the starting point of the job (in months from the beginning of the century), `tend` is the end of the job episode, and `failure` indicates whether the date in `tend` corresponds to the actual end of the employment in a certain job or whether it is a censored observation.

We first `stset` the data. As explained in [Cleves \(1999\)](#) and [Therneau and Grambsch \(2000\)](#), when analyzing multiple-failure data, we can consider two main approaches. One approach is to define the study time from the first time that an individual starts being at risk. The second approach is to define the study time from the last failure. We will take the second approach, which means that we treat each job episode as the subject.

Therefore, the origin is defined as the start of each job episode, and the study time will be the time from the start of each episode until the jobs end or the episode is censored.

```
. use http://www.stata-press.com/data/r15/jobhistory
(Job history data, Blossfeld et al 2007)

. stset tend, origin(tstart) failure(failure)

      failure event:  failure != 0 & failure < .
obs. time interval:  (origin, tend]
exit on or before:   failure
   t for analysis:   (time-origin)
               origin: time tstart
```

600	total observations	
0	exclusions	
600	observations remaining, representing	
458	failures in single-record/single-failure data	
40,782	total analysis time at risk and under observation	
	at risk from t =	0
	earliest observed entry t =	0
	last observed exit t =	428

We want to fit a Weibull model using the education level, the number of previous jobs, the prestige of the current job, and gender as explanatory variables. `education` records the highest education level before entering the labor market, `njobs` contains the number of previous jobs for each individual, and `prestige` is an index for the prestige of the current job. The `birthyear` variable indicates the

year of birth. female is 1 for women, 0 for men. To account for individual heterogeneity, we include a random effect at the individual level.

```
. mestreg education njobs prestige i.female || id:, distribution(weibull)
      failure _d: failure
      analysis time _t: (tend-origin)
                origin: time tstart
```

Fitting fixed-effects model:

```
Iteration 0: log likelihood = -5736904.5
Iteration 1: log likelihood = -2664.7487
Iteration 2: log likelihood = -2484.7829
Iteration 3: log likelihood = -2477.4358
Iteration 4: log likelihood = -2477.3338
Iteration 5: log likelihood = -2477.3337
```

Refining starting values:

```
Grid node 0: log likelihood = -2491.2191
```

Fitting full model:

```
Iteration 0: log likelihood = -2491.2191 (not concave)
Iteration 1: log likelihood = -2468.3995
Iteration 2: log likelihood = -2450.0938
Iteration 3: log likelihood = -2443.0739
Iteration 4: log likelihood = -2442.875
Iteration 5: log likelihood = -2442.8747
Iteration 6: log likelihood = -2442.8746
```

Mixed-effects Weibull PH regression

Group variable:	id	Number of obs	=	600
		Number of groups	=	201
		Obs per group:		
		min	=	1
		avg	=	3.0
		max	=	9

Integration method: mvaghermite

	Integration pts.	=	7
	Wald chi2(4)	=	87.38
	Prob > chi2	=	0.0000

Log likelihood = -2442.8746

_t	Haz. Ratio	Std. Err.	z	P> z	[95% Conf. Interval]
education	1.11897	.0463468	2.71	0.007	1.031722 1.213597
njobs	.7071195	.0357624	-6.85	0.000	.6403884 .7808043
prestige	.9677567	.0069576	-4.56	0.000	.9542157 .98149
1.female	1.75651	.3185526	3.11	0.002	1.231063 2.506228
_cons	.0053352	.0029015	-9.62	0.000	.0018376 .0154904
/ln_p	.1695545	.0453649			.0806409 .2584681
id					
var(_cons)	1.016459	.2149037			.671623 1.538347

Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline hazard (conditional on zero random effects).
LR test vs. Weibull model: chibar2(01) = 68.92 Prob >= chibar2 = 0.0000

The estimated variance of the random intercept is equal to 1.02

According to this model, an increase in the number of previous jobs is negatively associated with job mobility; the same is true for an increase in the prestige of the current job. By contrast, an increase in the years of education is positively associated with job mobility. Also, women seem to be more mobile than men.

We now store our estimates for later use:

```
. estimates store randint

The dataset has only two natural levels. However, for illustration purposes, let's consider the following situation. Assume that we want to account for unobserved variables associated with the date of birth, such as life experience, level of familiarity with new technologies, and family situation. We therefore add a random effect for the year of birth. Now, individuals will be nested within birth years.
```

```
. mestreg education njobs prestige i.female || birthyear: || id:,
> distribution(weibull)

      failure _d:  failure
      analysis time _t:  (tend-origin)
                  origin:  time tstart

(output omitted)

Mixed-effects Weibull PH regression                                Number of obs      =           600



| Group Variable | No. of Groups | Observations per Group |         |         |
|----------------|---------------|------------------------|---------|---------|
|                |               | Minimum                | Average | Maximum |
| birthyear      | 12            | 3                      | 50.0    | 99      |
| id             | 201           | 1                      | 3.0     | 9       |



Integration method: mvaghermite                                Integration pts. =           7
                                                                Wald chi2(4)      =          83.20
                                                                Prob > chi2       =          0.0000

Log likelihood = -2439.9066



| _t                         | Haz. Ratio | Std. Err. | z     | P> z  | [95% Conf. Interval] |          |
|----------------------------|------------|-----------|-------|-------|----------------------|----------|
| education                  | 1.120373   | .045203   | 2.82  | 0.005 | 1.035189             | 1.212566 |
| njobs                      | .7181197   | .0372039  | -6.39 | 0.000 | .6487813             | .7948686 |
| prestige                   | .966567    | .0069189  | -4.75 | 0.000 | .9531009             | .9802234 |
| 1.female                   | 1.734236   | .3022479  | 3.16  | 0.002 | 1.232419             | 2.440384 |
| _cons                      | .0059091   | .0031758  | -9.55 | 0.000 | .0020609             | .0169429 |
| /ln_p                      | .1685641   | .0454824  |       |       | .0794203             | .257708  |
| birthyear<br>var(_cons)    | .0950371   | .0741445  |       |       | .0205976             | .4385006 |
| birthyear>id<br>var(_cons) | .8728384   | .2020938  |       |       | .5544339             | 1.374099 |



Note: Estimates are transformed only in the first equation.
Note: _cons estimates baseline hazard (conditional on zero random effects).
LR test vs. Weibull model: chi2(2) = 74.85                      Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
```

The results for the fixed part of the model are similar to the ones in the previous model. Now, we have two estimated variances—one estimate for the random intercept at the individual level and one estimate for the random intercept at the birth-year level. The variance component for the individual level is smaller for this model, and it looks as if the first model might have been trying to explain a variance component at the birth-year level by incorporating it into the individual-level variance. We can perform a likelihood-ratio test to compare the stored model `randint` with the current model:

```
. lrtest randint .
```

Likelihood-ratio test	LR chi2(1) =	5.94
(Assumption: randint nested in .)	Prob > chi2 =	0.0148

Note: The reported degrees of freedom assumes the null hypothesis is not on the boundary of the parameter space. If this is not true, then the reported test is conservative.

The test is conservative because we are testing on the boundary of the parameter space; see *Distribution theory for likelihood-ratio test* in [ME] [me](#) for details. Provided that we are testing only one variance component, we can adjust the p -value accordingly by dividing the reported value by two, which results in an adjusted p -value equal to 0.0074.

The test is significant at the 0.05 level. It supports the three-level model with the additional variance component at the birth-year level.

◀

Stored results

mestreg stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(N_clust)</code>	number of clusters
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>gsem</code>
<code>e(cmd2)</code>	<code>mestreg</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	model name
<code>e(title)</code>	title in estimation output
<code>e(distribution)</code>	distribution

<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	offset
<code>e(exposure)</code>	exposure variable
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2
<code>e(vce)</code>	<i>vcetype</i> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(frm2)</code>	hazard or time
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of <code>ml</code> method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	<code>b</code> <code>V</code>
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>
<code>e(marginswexp)</code>	weight expression for <code>margins</code>
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>
Matrices	
<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance
Functions	
<code>e(sample)</code>	marks estimation sample

Methods and formulas

Methods and formulas are presented under the following headings:

Survival models

Survey data

Survival models

Survival models have a trivariate response (t_0, t, d) :

t_0 is the starting time under observation $t_0 \geq 0$;

t is the ending time under observation $t \geq t_0$; and

d is an indicator for failure $d \in \{0, 1\}$.

The survival function for a given family is the complement of the cumulative distribution function, $S(t) = 1 - F(t)$. The unconditional density for a failure at time t is given by

$$g(t) = \frac{\partial F(t)}{\partial t} = -\frac{\partial S(t)}{\partial t}$$

Some distributions contain ancillary parameters that are not denoted here.

The conditional density for a failure at time t is

$$g(t|t \geq t_0, d = 1) = g(t)/S(t_0)$$

and the conditional probability of survival without failure up to time t is

$$P(T \geq t|t \geq t_0, d = 0) = S(t)/S(t_0)$$

The conditional likelihood is given by

$$L(t, t_0, d) = \left\{ \frac{g(t)}{S(t_0)} \right\}^d \left\{ \frac{S(t)}{S(t_0)} \right\}^{1-d}$$

See [Survival distributions](#) in [\[SEM\] methods and formulas for gsem](#) for the specific density function corresponding to each distribution.

Given a set of cluster-level random effects $\mathbf{u}_j$ for $j = 1, \dots, M$, the conditional distribution of $\mathbf{t}_j = (t_{j1}, \dots, t_{jn_j})'$ on $\boldsymbol{\eta}_j = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j = (\mathbf{x}_{j1}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j, \dots, \mathbf{x}_{jn_j}\boldsymbol{\beta} + \mathbf{z}_{jn_j}\mathbf{u}_j)$ for cluster j is

$$f(\mathbf{t}_j|\boldsymbol{\eta}_j) = \prod_{i=1}^{n_j} f(t_{ji}|\eta_{ji})$$

where $f(t_{ji}|\eta_{ji})$ is the contribution to the likelihood from observation ji ; that is,

$$f(t_{ji}|\eta_{ji}) = \left\{ \frac{g(t_{ji}|\mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j)}{S(t_{0ji}|\mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j)} \right\}^{d_{ji}} \left\{ \frac{S(t_{ji}|\mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j)}{S(t_{0ji}|\mathbf{x}_{ji}\boldsymbol{\beta} + \mathbf{z}_{ji}\mathbf{u}_j)} \right\}^{1-d_{ji}} \quad (1)$$

where $g(t|\eta)$ and $S(t|\eta)$ are, respectively, the density and the survivor function conditional on the linear prediction η .

As mentioned in [Introduction](#) under *Remarks and examples*, `mestreg` does not allow delayed entry or gaps. Therefore, the first observation for a given subject will have a value of $t_0 = 0$, and subsequent spells for the subject must start at the end of the previous spell. That is, if observations ji and $j, i + 1$ belong to the same subject, then $t_{0j, i+1} = t_{ji}$.

Because the prior distribution of $\mathbf{u}_j$ is multivariate normal with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$, the likelihood contribution for the j th cluster is obtained by integrating $\mathbf{u}_j$ out of the joint density $f(\mathbf{t}_j, \mathbf{u}_j)$,

$$\mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int f(\mathbf{t}_j|\mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j) \exp(-\mathbf{u}_j'\boldsymbol{\Sigma}^{-1}\mathbf{u}_j/2) d\mathbf{u}_j \quad (2)$$

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [\[ME\] meglm](#) for details.

Survey data

In the presence of sampling weights, following [Rabe-Hesketh and Skrondal \(2006\)](#), the weighted log pseudolikelihood for a two-level model is given as

$$\mathcal{L}(\beta, \Sigma) = \sum_{j=1}^M w_j \log \int_{-\infty}^{\infty} \exp \left\{ \sum_{i=1}^{n_j} w_{i|j} \log f(t_{ji} | \eta_{ji}) \right\} \phi(\mathbf{v}_{j1}) d\mathbf{v}_{j1}$$

where w_j is the inverse of the probability of selection for the j th cluster; $w_{i|j}$ is the inverse of the conditional probability of selection of individual i , given the selection of cluster j ; $f(t_{ji} | \eta_{ji})$ is as in (1); and η_{ji} , $\phi(\cdot)$, $\mathbf{v}_{j1}$ are defined as in [Methods and formulas](#) in [ME] **meglm**.

Weighted estimation is achieved through the direct application of w_j and $w_{i|j}$ into the likelihood calculations as detailed above to reflect replicated clusters for w_j and replicated observations within clusters for $w_{i|j}$. Because this estimation is based on replicated clusters and observations, frequency weights are handled similarly.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Blossfeld, H.-P., K. Golsch, and G. Rohwer. 2007. [Event History Analysis with Stata](#). Mahwah, NJ: Erlbaum.
- Cleves, M. A. 1999. FAQ: How do I analyze multiple failure-time data using Stata? <http://www.stata.com/support/faqs/statistics/multiple-failure-time-data>.
- Cleves, M. A., W. W. Gould, and Y. V. Marchenko. 2016. [An Introduction to Survival Analysis Using Stata](#). Rev. 3rd ed. College Station, TX: Stata Press.
- Danahy, D. T., D. T. Burwell, W. S. Aronow, and R. Prakash. 1977. Sustained hemodynamic and antianginal effect of high dose oral isosorbide dinitrate. *Circulation* 55: 381–387.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Joe, H. 2008. Accuracy of Laplace approximation for discrete response mixed models. *Computational Statistics & Data Analysis* 52: 5066–5074.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- Leyland, A. H., and H. Goldstein, ed. 2001. [Multilevel Modelling of Health Statistics](#). New York: Wiley.
- Lin, X., and N. E. Breslow. 1996. Bias correction in generalized linear mixed models with multiple components of dispersion. *Journal of the American Statistical Association* 91: 1007–1016.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- Mayer, K. U., and E. Brückner. 1989. Lebensverläufe und Wohlfahrtsentwicklung: Konzeption, Design und Methodik der Erhebung von Lebensverläufen der Geburtsjahrgänge 1929–1931, 1939–1941, 1949–1951. Materialien aus der Bildungsforschung 35, Max-Planck-Institut für Bildungsforschung, Berlin.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. [Generalized, Linear, and Mixed Models](#). 2nd ed. Hoboken, NJ: Wiley.
- McLachlan, G. J., and K. E. Basford. 1988. [Mixture Models: Inference and Applications to Clustering](#). New York: Dekker.
- Pickles, A., and R. Crouchley. 1994. Generalizations and applications of frailty models for survival and event data. *Statistical Methods in Medical Research* 3: 263–278.
- . 1995. A comparison of frailty models for multivariate survival data. *Statistics in Medicine* 14: 1447–1461.
- Rabe-Hesketh, S., and A. Skrondal. 2006. Multilevel modelling of complex survey data. *Journal of the Royal Statistical Society, Series A* 169: 805–827.

- . 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.
- Rabe-Hesketh, S., A. Skrondal, and A. Pickles. 2004. Generalized multilevel structural equation modeling. *Psychometrika* 69: 167–190.
- . 2005. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics* 128: 301–323.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Self, S. G., and K.-Y. Liang. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82: 605–610.
- Skrondal, A., and S. Rabe-Hesketh. 2004. *Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models*. Boca Raton, FL: Chapman & Hall/CRC.
- Therneau, T. M., and P. M. Grambsch. 2000. *Modeling Survival Data: Extending the Cox Model*. New York: Springer.

Also see

- [ME] **mestreg postestimation** — Postestimation tools for mestreg
- [ME] **me** — Introduction to multilevel mixed-effects models
- [BAYES] **bayes: mestreg** — Bayesian multilevel parametric survival model
- [ST] **streg** — Parametric survival models
- [ST] **Glossary**
- [SVY] **svy estimation** — Estimation commands for survey data
- [XT] **xtstreg** — Random-effects parametric survival models
- [U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation commands are of special interest after `mestreg`:

Command	Description
<code>stcurve</code>	plot the survivor, hazard, and cumulative hazard functions
<code>estat group</code>	summarize the composition of the nested groups
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance-covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as mean and median survival times, hazards, survivor functions, linear predictions, and standard errors.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
mean	mean survival time; the default
median	median survival time
hazard	hazard
eta	fitted linear predictor
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
surv	predicted survivor function
density	predicted density function
distribution	predicted distribution function

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure

Integration	
<i>int_options</i>	integration options

`median` may not be combined with `marginal`.

<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only

<i>re_options</i>	Description

Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates

Integration	
<i>int_options</i>	integration options

<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main	
<code>mean</code> ,	the default, calculates the mean survival time.
<code>median</code>	calculates the median survival time.
<code>hazard</code>	calculates the hazard. When <code>marginal</code> is specified, marginal hazard is calculated as a ratio of the marginal density to the marginal survivor function.
<code>surv</code>	calculates the predicted survivor function.
<code>eta</code> , <code>xb</code> , <code>stdp</code> , <code>density</code> , <code>distribution</code> , <code>scores</code> , <code>conditional()</code> , <code>marginal</code> , and <code>nooffset</code> ;	see [ME] meglm postestimation . <code>marginal</code> may not be specified with <code>median</code> .

reffects, ebmeans, ebmodes, and reses()); see [ME] [meglm postestimation](#).

Integration

intpoints(), iterate(), and tolerance()); see [ME] [meglm postestimation](#).

margins

Description for margins

margins estimates margins of response for mean and median survival times and linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [predict(statistic ...) ...] [options]
```

statistic	Description
mean	mean survival time; the default
median	median survival time
xb	linear predictor for the fixed portion of the model only
hazard	not allowed with margins
eta	not allowed with margins
stdp	not allowed with margins
surv	not allowed with margins
density	not allowed with margins
distribution	not allowed with margins
reffects	not allowed with margins
scores	not allowed with margins

Options conditional(ebmeans) and conditional(ebmodes) are not allowed with margins.
Option marginal is assumed where applicable if conditional(fixedonly) is not specified.

Statistics not allowed with margins are functions of stochastic quantities other than e(b).
For the full syntax, see [R] [margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects parametric survival model with `mestreg`. For the most part, predictions center on obtaining estimates of the survival times or hazard functions. Conditional predictions are based on the computation of the group-specific random effects, and marginal predictions are obtained by numerically integrating out the random effects.

► **Example 1 : Predicting conditional and marginal mean survival time**

In [example 1](#) of [\[ME\] mestreg](#), we analyzed the time to infection of the catheter insertion point for 38 kidney dialysis patients. We fit the following model:

```
. use http://www.stata-press.com/data/r15/catheter
(Kidney data, McGilchrist and Aisbett, Biometrics, 1991)
. stset time, failure(infect)
(output omitted)
. mestreg age female || patient:, distribution(weibull)
(output omitted)
```

The `predict` command allows us to compute marginal and conditional predictions. Unless stated differently, we use the word “conditional” to mean “conditional on the empirical Bayes predictions of the random effects”. Below we compute marginal and conditional means for the mean survival time.

```
. predict m_marg, mean marginal
. predict m_cond, mean conditional
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)
```

Now, we can display the predictions for some of the patients.

```
. sort female age patient
. list patient female age m_* in 15/20, sepby(patient)
```

	patient	female	age	m_marg	m_cond
15.	29	0	53	52.79355	22.36027
16.	29	0	53	52.79355	22.36027
17.	16	0	60	50.67546	28.01295
18.	16	0	60	50.67546	28.01295
19.	38	0	60	50.67546	49.47013
20.	38	0	60	50.67546	49.47013

We see in the output that the predicted expected conditional mean for patient 29 is equal to 22.36 (shown in `m_cond`). This is the expected time to infection for this patient. However, the predicted marginal mean for this patient is 52.79 (shown in `m_marg`). This is the expected time to infection for a patient from the population who is male and is 53 years old. This particular patient seems to be more prone to infection than would be expected based on his age and gender.


Conditional predictions are specific to each group, while marginal predictions are the same within each covariate pattern through the data. Patients 16 and 38 have the same covariate patterns; therefore, their marginal predicted means are the same. However, conditional predicted means differ.

margins and marginsplot show the changes in the marginal means for different ages.

```
. margins, predict(mean marginal) at(female=0 age=(20(5)70)) noatlegend
Adjusted predictions      Number of obs      =      76
Model VCE      : OIM
Expression      : Marginal predicted mean, predict(mean marginal)
```

	Delta-method		z	P> z	[95% Conf. Interval]	
	Margin	Std. Err.				
_at						
1	64.03481	28.99882	2.21	0.027	7.19816	120.8715
2	62.18903	26.33284	2.36	0.018	10.57761	113.8005
3	60.39646	24.11456	2.50	0.012	13.13279	107.6601
4	58.65556	22.37001	2.62	0.009	14.81116	102.5
5	56.96484	21.11488	2.70	0.007	15.58043	98.34925
6	55.32285	20.34538	2.72	0.007	15.44663	95.19908
7	53.7282	20.03192	2.68	0.007	14.46635	92.99004
8	52.17951	20.12	2.59	0.010	12.74503	91.61398
9	50.67546	20.53852	2.47	0.014	10.42071	90.93021
10	49.21476	21.21134	2.32	0.020	7.64129	90.78823
11	47.79617	22.06715	2.17	0.030	4.545348	91.04698

```
. marginsplot
Variables that uniquely identify margins: age
```



We see that the predicted marginal mean decreases with age; older patients are expected to have an event earlier. This is consistent with the findings from [example 1](#) of [\[ME\] mestreg](#) that the hazard is increasing with age.



➤ Example 2 : Predicting survivor functions

Continuing with [example 1](#), we now predict survivor functions.


```
. predict S_marg, surv marginal
(using 7 quadrature points)
. predict S_cond, surv conditional
(predictions based on fixed effects and posterior means of random effects)
(using 7 quadrature points)
. sort female age patient _t
. list patient female age _t S_* in 15/20, sepby(patient)
```

	patient	female	age	_t	S_marg	S_cond
15.	29	0	53	2	.9628581	.9564017
16.	29	0	53	25	.5165027	.3493623
17.	16	0	60	4	.9122225	.9230723
18.	16	0	60	17	.6273606	.6129264
19.	38	0	60	8	.8141544	.9107039
20.	38	0	60	63	.20487	.2900458

Survival predictions vary with the value of the study time variable because they are predictions of the survivor function at the study time `_t`. For example, patient 29 has a 0.96 probability that a new insertion remains at least 2 days without infection and a 0.35 probability that a new insertion remains at least 25 days without infection. For a random patient from the population, the probabilities to remain at least 2 or 25 days without infection are, respectively, 0.96 and 0.52.

We can use `stcurve` to plot these predictions simultaneously for males and females of the same age.

```
. stcurve, surv at1(female=0 age=53) at2(female=1 age=53)
(option unconditional assumed)
```



We see that the survivor function for females is above the survivor function for males, which means that females have a greater probability of not having an episode by study time `_t`.

► Example 3 : Comparing marginal hazards

In [example 2](#) of [\[ME\] mestreg](#), we estimated two different distributions with random effects on patient and covariates age and female. Here we compare the marginal hazards using `stcurve`. By default, `stcurve` plots predictions at the mean of the covariates, computed over the whole estimation sample. We plot the predictions for `female==1`.

```
. mestreg age female || patient:, dist(weibull) time
(output omitted)

. stcurve, hazard at(female=1)
(option unconditional assumed)


. graph save g1
(file g1.gph saved)

. mestreg age female || patient:, dist(gamma)
(output omitted)

. stcurve, hazard at(female=1)
(option unconditional assumed)

. graph save g2
(file g2.gph saved)

. graph combine g1.gph g2.gph
```



The two estimated marginal hazards are similar. The marginal hazard has a very different shape from the conditional hazards. The conditional hazard function for a Weibull or a gamma distribution are both monotonic (increasing, constant, or decreasing, depending on the parameters).


► Example 4 : Obtaining predictions of random effects

In [example 3](#) of [ME] **mestreg**, we fit a Weibull model with random intercepts and random coefficients at the subject level. We obtained a positive covariance between the random effects. We refit the model here and then use `predict` with the option `reffects` to obtain predictions of the random effects based on the empirical Bayes posterior means.

```
. use http://www.stata-press.com/data/r15/angina, clear
(Angina drug data, Rabe-Hesketh and Skrondal, 2012, ch 15.7)
. mestreg occasion##treat || pid: treat, distribution(weibull)
> covariance(unstructured)
(output omitted)
. predict re*, reffects
(calculating posterior means of random effects)
(using 7 quadrature points)
```

Plotting the predictions of the predicted random coefficient versus the random intercept shows the pattern we discussed in the main section: individuals with a larger random slope tend also to have a larger random intercept.

```
. twoway scatter re1 re2
```



Individuals with large random intercepts have individual hazards that are larger than those of other individuals with the same covariate patterns. Also, individuals with large random coefficients have individual conditional hazard ratios for treatment that are larger than those of other individuals with the same covariate pattern.

In other words, if the aim of the treatment is to decrease the hazard, then the positive correlation means that the treatment tends to be less effective for individuals who have a higher individual hazard (within the same occasion number).

◀

► Example 5 : Conditional and marginal hazards

In [example 1](#) of [ME] **mestreg**, we mentioned that hazard ratios should be interpreted as conditional on the random effects. Here we use `predict` to illustrate this concept. We use a simulated dataset for a Weibull model with random effects for `group` and a binary covariate `x`.

We show that for a given group, the conditional hazard function satisfies the proportional-hazards (PH) assumption. That is, for a given group j ,

$$h(t|x = 1, \text{group} = j) = \exp(\beta_x) \times h(t|x = 0, \text{group} = j)$$

is equivalent to

$$\log\{h(t|x = 1, \text{group} = j)\} = \beta_x + \log\{h(t|x = 0, \text{group} = j)\}$$

This property of the log hazard-function translates to one curve being a shifted version of the other, which is easier to see than the proportionality of the (untransformed) hazard function.

After fitting the model, we use `predict` to compute the conditional prediction of the hazard function for group 1; we create the variables `hcond0` and `hcond1`. `hcond0` will contain the conditional hazard for group 1 when `x==0`; `hcond1` will contain the conditional hazard for group 1 when `x==1`.

We also create `zcond = loghcond0 +  $\beta_x$` . If the PH assumption is satisfied, then the plotted values of `zcond` will be superimposed on those of `loghcond1`.

```
. use http://www.stata-press.com/data/r15/weibre, clear
. mestreg i.x || group:, distribution(weibull) nolog
      failure _d:  1 (meaning all fail)
      analysis time _t:  t
Mixed-effects Weibull PH regression          Number of obs    =    100,000
Group variable:                group          Number of groups  =      500
                                      Obs per group:
                                      min =          200
                                      avg  =         200.0
                                      max  =          200
Integration method: mvaghermite              Integration pts.  =          7
Log likelihood = 175196.47                    Wald chi2(1)     =   21447.86
                                      Prob > chi2      =    0.0000
```

_t	Haz. Ratio	Std. Err.	z	P> z	[95% Conf. Interval]	
1.x	2.713138	.0184908	146.45	0.000	2.677137	2.749622
_cons	2.564135	.0797385	30.28	0.000	2.412518	2.725281
/ln_p	-.6925791	.0024746			-.6974291	-.687729
group						
var(_cons)	.472804	.0303096			.4169789	.536103


Note: Estimates are transformed only in the first equation.

Note: `_cons` estimates baseline hazard (conditional on zero random effects).

LR test vs. Weibull model: `chibar2(01) = 35800.39` Prob >= `chibar2` = 0.0000

```
. predict hcond, hazard conditional(ebmeans)
(predictions based on fixed effects and posterior means of random effects)
. gen loghcond0 = log(hcond) if x==0
(49,991 missing values generated)
. gen loghcond1 = log(hcond) if x==1
(50,009 missing values generated)
. gen zcond = loghcond0 + _b[_t:1.x]
(49,991 missing values generated)
. sort _t group
```


```
. twoway line loghcond0 loghcond1 zcond _t if group==1
```



In the graph above, the line for `loghcond1` cannot be distinguished from the line for `zcond` for most of the distribution. This illustrates that the PH assumption is satisfied for the conditional hazard. Notice that you can still see a part of `loghcond1` near the origin. This is because the two variables correspond to different values of `_t` and only `loghcond1` happens to be defined at the early values.

Now, we make the same computation for the marginal hazard.

```
. predict hmarg, hazard marginal
. gen loghmarg0 = log(hmarg) if x==0
(49,991 missing values generated)
. gen loghmarg1 = log(hmarg) if x==1
(50,009 missing values generated)
. gen zmarg = loghmarg0 + _b[_t:1.x]
(49,991 missing values generated)
. sort _t group
. twoway line loghmarg0 loghmarg1 zmarg _t
```



The curve for `zmarg` is clearly different from the curve for `loghmarg1`, demonstrating that the marginal distribution does not meet the PH assumption. Notice that the line for `loghmarg1` is shorter than the others. This is because predictions are obtained at the values of `_t` in the dataset. These values of `_t` were simulated based on the model, which determines that observations with `x==1` fail earlier.

◀

Methods and formulas

Methods and formulas for predicting random effects and other statistics are given in [Methods and formulas](#) of [ME] [meglm postestimation](#). Statistics of special interest for survival analysis are described below.

`predict newvar` with the `conditional()` option computes the following predictions:

`median:`

$$newvar_{ji} = \{t : \hat{S}(t|\mathbf{x}_{ji}, \hat{u}_{ji}) = 1/2\}$$

where $\hat{S}(t|\mathbf{x}_{ji}, \hat{u}_{ji})$ is $S(t|\mathbf{x}_{ji}\hat{\beta} + \hat{u}_{ji})$, where $\hat{u}_{ji}$ are the empirical Bayes predictions for u_{ji} . If `conditional(fixedonly)` is specified, then 0 is substituted for $\hat{u}_{ji}$.

`mean:`

$$newvar_{ji} = \int_0^{\infty} \hat{S}(t|\mathbf{x}_{ji}, u_{ji}) dt$$

`surv:`

$$newvar_{ji} = \hat{S}(t_{ji}|\mathbf{x}_{ji}, \hat{u}_{ji})$$

`hazard:`

$$newvar_{ji} = \hat{g}(t_{ji}|\mathbf{x}_{ji}, \hat{u}_{ji}) / \hat{S}(t_{ji}|\mathbf{x}_{ji}, \hat{u}_{ji})$$

where $\hat{g}(t|\mathbf{x}_{ji}, u_{ji})$ is the density $g(t|\mathbf{x}_{ji}\hat{\beta} + \hat{u}_{ji})$.

When the `marginal` option is used with `mean` or `surv`, the prediction is computed marginally with respect to the random effects. That is, the prediction is integrated over the random-effects distributions. When the `marginal` option is used with `hazard`, the hazard for the marginal distribution is computed. That is, the predicted hazard is computed as the quotient of the marginal hazard and the marginal survivor function.

Also see

[ME] [mestreg](#) — Multilevel mixed-effects parametric survival models

[ME] [meglm postestimation](#) — Postestimation tools for `meglm`

[ME] [mixed postestimation](#) — Postestimation tools for `mixed`

[ST] [stcurve](#) — Plot survivor, hazard, cumulative hazard, or cumulative incidence function

[U] [20 Estimation and postestimation commands](#)

[Description](#)
[Options](#)
[Also see](#)

[Quick start](#)
[Remarks and examples](#)

[Menu](#)
[Stored results](#)

[Syntax](#)
[Methods and formulas](#)

Description

`metobit` fits mixed-effects models for continuous responses where the outcome variable is censored. Censoring limits may be fixed for all observations or vary across observations.

Quick start

Without weights

Two-level tobit regression of `y` on `x` with random intercepts by `lev2` where `y` is censored at a lower limit of 5

```
metobit y x || lev2:, ll(5)
```

As above, but specify that left-censoring occurs at 5 and right-censoring occurs at 25

```
metobit y x || lev2:, ll(5) ul(25)
```

As above, but where `lower` and `upper` are variables containing the censoring limits

```
metobit y x || lev2:, ll(lower) ul(upper)
```

Mixed-effects model adding random coefficients for `x`

```
metobit y x || lev2: x, ll(5)
```

Three-level random-intercept model of `y` on `x` with `lev2` nested within `lev3`

```
metobit y x || lev3: || lev2:, ll(5)
```

Crossed-effects model of `y` on `x` with two-way crossed random effects by factors `a` and `b`

```
metobit y x || _all:R.a || b:, ll(5)
```

With weights

Two-level tobit regression of `y` on `x` with random intercepts by `lev2` and observation-level frequency weights `wvar1`

```
metobit y x [fweight=wvar1] || lev2:, ll(5)
```

Two-level random-intercept model from a two-stage sampling design with PSUs identified by `psu` using PSU-level and observation-level sampling weights `wvar2` and `wvar1`, respectively

```
metobit y x [pweight=wvar1] || psu:, pweight(wvar2) ll(5)
```

Same as above, but `svyset` data first

```
svyset psu, weight(wvar2) || _n, weight(wvar1)
svy: metobit y x || psu:, ll(5)
```

Menu

Statistics > Multilevel mixed-effects models > Tobit regression

Syntax

```
metobit depvar fe_equation [ || re_equation ] [ || re_equation ... ] [ , options ]
```

where the syntax of *fe_equation* is

```
[ indepvars ] [ if ] [ in ] [ weight ] [ , fe_options ]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [ varlist ] [ , re_options ]
```

for random effects among the values of a factor variable

```
levelvar: R.varname
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<code>noconstant</code>	suppress constant term from the fixed-effects equation
<code>offset(<i>varname</i>)</code>	include <i>varname</i> in model with coefficient constrained to 1

<i>re_options</i>	Description
Model	
<code>covariance(<i>vartype</i>)</code>	variance–covariance structure of the random effects
<code>noconstant</code>	suppress constant term from the random-effects equation
<code>fweight(<i>varname</i>)</code>	frequency weights at higher levels
<code>iweight(<i>varname</i>)</code>	importance weights at higher levels
<code>pweight(<i>varname</i>)</code>	sampling weights at higher levels

<i>options</i>	Description
Model	
<code>ll</code> [(<i>varname</i> #)]	left-censoring variable or limit
<code>ul</code> [(<i>varname</i> #)]	right-censoring variable or limit
<code>constraints</code> (<i>constraints</i>)	apply specified linear constraints
<code>collinear</code>	keep collinear variables
SE/Robust	
<code>vce</code> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster</code> <i>clustvar</i>
Reporting	
<code>level</code> (#)	set confidence level; default is <code>level(95)</code>
<code>nocnsreport</code>	do not display constraints
<code>notable</code>	suppress coefficient table
<code>noheader</code>	suppress output header
<code>nogroup</code>	suppress table summarizing groups
<code>display_options</code>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Integration	
<code>intmethod</code> (<i>intmethod</i>)	integration method
<code>intpoints</code> (#)	set the number of integration (quadrature) points for all levels; default is <code>intpoints(7)</code>
Maximization	
<code>maximize_options</code>	control the maximization process; seldom used
<code>startvalues</code> (<i>svmethod</i>)	method for obtaining starting values
<code>startgrid</code> [(<i>gridspec</i>)]	perform a grid search to improve starting values
<code>noestimate</code>	do not fit the model; show starting values instead
<code>dnumerical</code>	use numerical derivative techniques
<code>coeflegend</code>	display legend instead of statistics
<i>vartype</i>	Description
<code>independent</code>	one unique variance parameter per random effect, all covariances 0; the default unless the <code>R.</code> notation is used
<code>exchangeable</code>	equal variances for random effects, and one common pairwise covariance
<code>identity</code>	equal variances for random effects, all covariances 0; the default if the <code>R.</code> notation is used
<code>unstructured</code>	all variances and covariances to be distinctly estimated
<code>fixed</code> (<i>matname</i>)	user-selected variances and covariances constrained to specified values; the remaining variances and covariances unrestricted
<code>pattern</code> (<i>matname</i>)	user-selected variances and covariances constrained to be equal; the remaining variances and covariances unrestricted

<i>intmethod</i>	Description
<u>mv</u> aghermite	mean–variance adaptive Gauss–Hermite quadrature; the default unless a crossed random-effects model is fit
<u>mc</u> aghermite	mode-curvature adaptive Gauss–Hermite quadrature
<u>g</u> hermite	nonadaptive Gauss–Hermite quadrature
<u>l</u> aplace	Laplacian approximation; the default for crossed random-effects models

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

depvar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *by*, and *svy* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] *bayes: metobit*.

vce() and *weights* are not allowed with the *svy* prefix; see [SVY] *svy*.

fweights, *iweights*, and *pweights* are allowed; see [U] 11.1.6 *weight*. Only one type of weight may be specified.

Weights are not supported under the Laplacian approximation or for crossed models.

startvalues(), *startgrid*, *noestimate*, *dnumerical*, and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

ll[(*varname*|#)] and *ul*[(*varname*|#)] indicate the censoring points. You may specify one or both. *ll*() indicates the lower limit for left-censoring. *ul*() indicates the upper limit for right-censoring. Observations with *depvar* ≤ *ll*() are left-censored, observations with *depvar* ≥ *ul*() are right-censored, and remaining observations are not censored.

offset(*varname*) specifies that *varname* be included in the fixed-effects portion of the model with the coefficient constrained to be 1.

covariance(*vartype*) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, *unstructured*, *fixed*(*matname*), or *pattern*(*matname*).

covariance(*independent*) covariance structure allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0. The default is *covariance*(*independent*) unless a crossed random-effects model is fit, in which case the default is *covariance*(*identity*).

covariance(*exchangeable*) structure specifies one common variance for all random effects and one common pairwise covariance.

covariance(*identity*) is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

covariance(*unstructured*) allows for all variances and covariances to be distinct. If an equation consists of *p* random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(*fixed*(*matname*)) and *covariance*(*pattern*(*matname*)) covariance structures provide a convenient way to impose constraints on variances and covariances of random effects. Each specification requires a *matname* that defines the restrictions placed on variances and

covariances. Only elements in the lower triangle of *matname* are used, and row and column names of *matname* are ignored. A missing value in *matname* means that a given element is unrestricted. In a `fixed(matname)` covariance structure, (co)variance (i, j) is constrained to equal the value specified in the i, j th entry of *matname*. In a `pattern(matname)` covariance structure, (co)variances (i, j) and (k, l) are constrained to be equal if $matname[i, j] = matname[k, l]$.

`fweight(varname)` specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, `[fw=fwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `fweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [fw = wt1] || school: ... , fweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) frequency weights, and `wt2` would hold the second-level (the school-level) frequency weights.

`iweight(varname)` specifies importance weights at higher levels in a multilevel model, whereas importance weights at the first level (the observation level) are specified in the usual manner, for example, `[iw=iwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `iweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [iw = wt1] || school: ... , iweight(wt2) ...
```

the variable `wt1` would hold the first-level (the observation-level) importance weights, and `wt2` would hold the second-level (the school-level) importance weights.

`pweight(varname)` specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, `[pw=pwtvar1]`. *varname* can be any valid Stata variable name, and you can specify `pweight()` at levels two and higher of a multilevel model. For example, in the two-level model

```
. mecmd fixed_portion [pw = wt1] || school: ... , pweight(wt2) ...
```

variable `wt1` would hold the first-level (the observation-level) sampling weights, and `wt2` would hold the second-level (the school-level) sampling weights.

`constraints(constraints)`, `collinear`; see [\[R\] estimation options](#).

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [\[R\] vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

Reporting

`level(#)`, `nocnsreport`; see [\[R\] estimation options](#).

`notable` suppresses the estimation table, either at estimation or upon replay.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [\[R\] estimation options](#).

Integration

`intmethod(intmethod)` specifies the integration method to be used for the random-effects model. `mvaghermite` performs mean–variance adaptive Gauss–Hermite quadrature; `mcaghermite` performs mode-curvature adaptive Gauss–Hermite quadrature; `ghermite` performs nonadaptive Gauss–Hermite quadrature; and `laplace` performs the Laplacian approximation, equivalent to mode-curvature adaptive Gaussian quadrature with one integration point.

The default integration method is `mvaghermite` unless a crossed random-effects model is fit, in which case the default integration method is `laplace`. The Laplacian approximation has been known to produce biased parameter estimates; however, the bias tends to be more prominent in the estimates of the variance components rather than in the estimates of the fixed effects.

For crossed random-effects models, estimation with more than one quadrature point may be prohibitively intensive even for a small number of levels. For this reason, the integration method defaults to the Laplacian approximation. You may override this behavior by specifying a different integration method.

`intpoints(#)` sets the number of integration points for quadrature. The default is `intpoints(7)`, which means that seven quadrature points are used for each level of random effects. This option is not allowed with `intmethod(laplace)`.

The more integration points, the more accurate the approximation to the log likelihood. However, computation time increases as a function of the number of quadrature points raised to a power equaling the dimension of the random-effects specification. In crossed random-effects models and in models with many levels or many random coefficients, this increase can be substantial.

Maximization

maximize_options: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, `nonrtolerance`, and `from(init_specs)`; see [R] [maximize](#). Those that require special mention for `metobit` are listed below.

`from()` accepts a properly labeled vector of initial values or a list of coefficient names with values. A list of values is not allowed.

The following options are available with `metobit` but are not shown in the dialog box:

`startvalues(svmethod)`, `startgrid[(gridspec)]`, `noestimate`, and `dnumerical`; see [ME] [meglm](#).

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Mixed-effects tobit regression is tobit regression containing both fixed effects and random effects. In longitudinal data and panel data, random effects are useful for modeling intracluster correlation; that is, observations in the same cluster are correlated because they share common cluster-level random effects.

In a mixed-effects tobit regression, the values of the outcome variable may be observed, unobserved but known to fall below a given limit (left-censored data), or unobserved but known to fall above a given limit (right-censored data). That is, the observed data, y_{ij}^* , represent possibly censored versions of y_{ij} for the i th observation within the j th cluster.

The observed outcome is therefore defined as

$$y_{ij}^* = \begin{cases} y_{ij} & \text{if } a < y_{ij} < b \\ a & \text{if } y_{ij} \leq a \\ b & \text{if } y_{ij} \geq b \end{cases}$$

where a is the lower-censoring limit and b is the upper-censoring limit. If the data are uncensored, $y_{ij}^* = y_{ij}$, and the value is determined by the value of the outcome variable. If they are left-censored, all that is known is that $y_{ij} \leq a$ and y_{ij}^* is determined by `ll()`. If they are right-censored, all that is known is that $y_{ij} \geq b$ and y_{ij}^* is determined by `ul()`. The censoring limits specified in `ll()` and `ul()` can be the same for all observations or can vary from observation to observation.

Regardless of the type of censoring, the expected value of the underlying dependent variable—say, $\mathbf{y}$ —is modeled using the following linear prediction:

$$E(\mathbf{y}|\mathbf{X}, \mathbf{u}) = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} \quad (1)$$

$\mathbf{X}$ is an $n \times p$ design/covariate matrix, analogous to the covariates you would find in a standard linear regression model, with regression coefficients (fixed effects) $\boldsymbol{\beta}$. $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. This linear prediction also contains the offset when `offset()` is specified.

The columns of matrix $\mathbf{Z}$ are the covariates corresponding to the random effects and can be used to represent both random intercepts and random coefficients. For example, in a random-intercepts model, $\mathbf{Z}$ is simply the scalar 1. The random effects $\mathbf{u}$ are realizations from a multivariate normal distribution with mean $\mathbf{0}$ and $q \times q$ variance matrix $\boldsymbol{\Sigma}$. The random effects are not directly estimated as model parameters but are instead summarized according to the unique elements of $\boldsymbol{\Sigma}$, known as variance components. One special case of (1) places $\mathbf{Z} = \mathbf{X}$ so that all covariate effects are essentially random and distributed as multivariate normal with mean $\boldsymbol{\beta}$ and variance $\boldsymbol{\Sigma}$.

Below we present a short example of mixed-effects tobit regression; refer to [ME] **me** and [ME] **meglm** for additional examples of random-effects models. A two-level tobit model can also be fit using `xttobit`; see [XT] **xttobit**. In the absence of random effects, mixed-effects tobit regression reduces to standard tobit regression; see [R] **tobit**.

► Example 1: Random-intercept model

We have wage data on young women who were between ages 14 and 26 in 1968 and who were surveyed over the period 1968–1988; see [XT] **xt** for a more detailed discussion of the data. We are interested in the effect of completed years of schooling, current age, union membership, and residence in the South on wages.

```
. use http://www.stata-press.com/data/r15/nlswork
(National Longitudinal Survey. Young Women 14-26 years of age in 1968)
```

We fit a mixed-effects tobit model of the log of inflation-adjusted wages (`ln_wage`). For illustration purposes, we use the `ul()` option to impose an artificial upper limit at 1.96, the 75th percentile of the recorded log wages.

```
. metobit ln_wage i.union age south##c.grade || idcode:, ul(1.96)
```

Fitting fixed-effects model:

```
Iteration 0:  log likelihood = -11628.188
Iteration 1:  log likelihood = -10617.455
Iteration 2:  log likelihood = -10555.304
Iteration 3:  log likelihood = -10554.78
Iteration 4:  log likelihood = -10554.78
```

Refining starting values:

```
Grid node 0:  log likelihood = -10225.917
```

Fitting full model:

```
Iteration 0:  log likelihood = -10225.917 (not concave)
Iteration 1:  log likelihood = -8728.9674 (not concave)
Iteration 2:  log likelihood = -7827.6894 (not concave)
Iteration 3:  log likelihood = -7112.0272
Iteration 4:  log likelihood = -6894.0253
Iteration 5:  log likelihood = -6821.7055
Iteration 6:  log likelihood = -6818.5592
Iteration 7:  log likelihood = -6818.5512
Iteration 8:  log likelihood = -6818.5512
```

Mixed-effects tobit regression

Number of obs	=	19,224
Uncensored	=	13,188
Left-censored	=	0
Right-censored	=	6,036

Limits: lower = -inf
upper = 1.96

Group variable: idcode

Number of groups	=	4,148
Obs per group:		
min	=	1
avg	=	4.6
max	=	12

Integration method: mvaghermite

Integration pts.	=	7
Wald chi2(5)	=	2812.43
Prob > chi2	=	0.0000

Log likelihood = -6818.5512

ln_wage	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
1.union	.1418088	.0068398	20.73	0.000	.1284029 .1552146
age	.0107585	.0004068	26.45	0.000	.0099612 .0115559
1.south	-.2373995	.048346	-4.91	0.000	-.3321559 -.1426431
grade	.0763865	.0029104	26.25	0.000	.0706822 .0820909
south# c.grade					
1	.0099306	.0037452	2.65	0.008	.0025902 .0172709
_cons	.4146363	.0396691	10.45	0.000	.3368864 .4923862
idcode var(_cons)	.0985482	.003018			.0928071 .1046444
var(e.ln_w~e)	.0619327	.000876			.0602394 .0636736

LR test vs. tobit model: chibar2(01) = 7472.46 Prob >= chibar2 = 0.0000

The estimation table reports the fixed effects, which are interpreted just as you would the output from `tobit`, and the estimated variance components. Because the dependent variable is log transformed, the fixed-effects coefficients can be interpreted in terms of a percent change. For example, we see that on average, union members make 14.2% more than nonunion members and that each additional year of age is associated with a 1.1% increase in wages.

The random-effects equation is labeled `idcode`. The estimated variance of the subject-specific random intercept is 0.098 with standard error 0.003. A likelihood-ratio test comparing the model with a tobit model without random effects is provided under the table and indicates that the two-level tobit model is preferred.



Stored results

`metobit` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(N_unc)</code>	number of uncensored observations
<code>e(N_lrc)</code>	number of left-censored observations
<code>e(N_rlc)</code>	number of right-censored observations
<code>e(k)</code>	number of parameters
<code>e(k_dv)</code>	number of dependent variables
<code>e(k_eq)</code>	number of equations in <code>e(b)</code>
<code>e(k_eq_model)</code>	number of equations in overall model test
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(df_m)</code>	model degrees of freedom
<code>e(ll)</code>	log likelihood
<code>e(N_clust)</code>	number of clusters
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(ic)</code>	number of iterations
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>meglm</code>
<code>e(cmd2)</code>	<code>metobit</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	names of dependent variables
<code>e(llopt)</code>	contents of <code>ll()</code> , if specified
<code>e(ulopt)</code>	contents of <code>ul()</code> , if specified
<code>e(wtype)</code>	weight type
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(iweightk)</code>	<code>iweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(covariates)</code>	list of covariates
<code>e(ivars)</code>	grouping variables
<code>e(model)</code>	<code>tobit</code>
<code>e(title)</code>	title in estimation output
<code>e(link)</code>	<code>identity</code>
<code>e(family)</code>	<code>gaussian</code>
<code>e(clustvar)</code>	name of cluster variable
<code>e(offset)</code>	<code>offset</code>
<code>e(intmethod)</code>	integration method
<code>e(n_quad)</code>	number of integration points
<code>e(chi2type)</code>	Wald; type of model χ^2

<code>e(vce)</code>	<code>vcetype</code> specified in <code>vce()</code>
<code>e(vcetype)</code>	title used to label Std. Err.
<code>e(opt)</code>	type of optimization
<code>e(which)</code>	max or min; whether optimizer is to perform maximization or minimization
<code>e(ml_method)</code>	type of ml method
<code>e(user)</code>	name of likelihood-evaluator program
<code>e(technique)</code>	maximization technique
<code>e(datasignature)</code>	the checksum
<code>e(datasignaturevars)</code>	variables used in calculation of checksum
<code>e(properties)</code>	b V
<code>e(estat_cmd)</code>	program used to implement <code>estat</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(marginsnotok)</code>	predictions disallowed by <code>margins</code>
<code>e(marginswtype)</code>	weight type for <code>margins</code>
<code>e(marginswexp)</code>	weight expression for <code>margins</code>
<code>e(footnote)</code>	program used to implement the footnote display
<code>e(asbalanced)</code>	factor variables <code>fvset</code> as <code>asbalanced</code>
<code>e(asobserved)</code>	factor variables <code>fvset</code> as <code>asobserved</code>

Matrices

<code>e(b)</code>	coefficient vector
<code>e(Cns)</code>	constraints matrix
<code>e(ilog)</code>	iteration log (up to 20 iterations)
<code>e(gradient)</code>	gradient vector
<code>e(N_g)</code>	group counts
<code>e(g_min)</code>	group-size minimums
<code>e(g_avg)</code>	group-size averages
<code>e(g_max)</code>	group-size maximums
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(V_modelbased)</code>	model-based variance

Functions

<code>e(sample)</code>	marks estimation sample

Methods and formulas

Without a loss of generality, consider a two-level regression model

$$E(\mathbf{y}_j | \mathbf{X}_j, \mathbf{u}_j) = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j \quad \mathbf{y} \sim \text{normal}$$

for $j = 1, \dots, M$ clusters, with the j th cluster consisting of n_j observations, where, for the j th cluster, $\mathbf{y}_j$ is the $n_j \times 1$ censored response vector, $\mathbf{X}_j$ is the $n_j \times p$ matrix of fixed predictors, $\mathbf{Z}_j$ is the $n_j \times q$ matrix of random predictors, $\mathbf{u}_j$ is the $q \times 1$ vector of random effects, and $\boldsymbol{\beta}$ is the $p \times 1$ vector of regression coefficients on the fixed predictors. The random effects, $\mathbf{u}_j$, are assumed to be multivariate normal with mean $\mathbf{0}$ and variance $\boldsymbol{\Sigma}$.

Let η_j be the linear predictor, $\eta_j = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{Z}_j \mathbf{u}_j$, that also includes the offset variable when `offset()` is specified. y_{ij} and η_{ij} are the i th individual elements of $\mathbf{y}_j$ and $\boldsymbol{\eta}_j$, $i = 1, \dots, n_j$. a_{ij} refers to the lower limit for observation ij , and b_{ij} refers to the upper limit for observation ij . The conditional density function for the response at observation ij is then

$$f(y_{ij}^*|\eta_{ij}) = \begin{cases} (\sqrt{2\pi}\sigma_\epsilon)^{-1} \exp^{-(y_{ij}-\eta_{ij})^2/(2\sigma_\epsilon^2)} & \text{if } y_{ij} = y_{ij}^* \\ \Phi\left(\frac{a_{ij}-\eta_{ij}}{\sigma_\epsilon}\right) & \text{if } y_{ij} \leq y_{ij}^* \\ 1 - \Phi\left(\frac{b_{ij}-\eta_{ij}}{\sigma_\epsilon}\right) & \text{if } y_{ij} \geq y_{ij}^* \end{cases}$$

where $\Phi(\cdot)$ is the cumulative normal distribution.

Because the observations are assumed to be conditionally independent, the conditional log density function for cluster j is

$$\log f(\mathbf{y}_j^*|\boldsymbol{\eta}_j) = \sum_{i=1}^{n_i} \log f(y_{ij}^*|\eta_{ij})$$

and the likelihood function for cluster j is given by

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int_{\Re^q} f(\mathbf{y}_j^*|\boldsymbol{\eta}_j) \exp\left(-\frac{1}{2} \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j\right) d\mathbf{u}_j \\ &= (2\pi)^{-q/2} |\boldsymbol{\Sigma}|^{-1/2} \int_{\Re^q} \exp\left\{\log f(\mathbf{y}_j^*|\boldsymbol{\eta}_j) - \frac{1}{2} \mathbf{u}_j' \boldsymbol{\Sigma}^{-1} \mathbf{u}_j\right\} d\mathbf{u}_j \end{aligned} \tag{1}$$

where $\Re$ denotes the set of values on the real line and $\Re^q$ is the analog in q -dimensional space.

The integration in (2) has no closed form and thus must be approximated; see [Methods and formulas](#) in [ME] **meglm** for details.

metobit supports multilevel weights and survey data; see [Methods and formulas](#) in [ME] **meglm** for details.

Also see

[ME] **metobit postestimation** — Postestimation tools for metobit

[ME] **meintreg** — Multilevel mixed-effects interval regression

[ME] **me** — Introduction to multilevel mixed-effects models

[BAYES] **bayes: metobit** — Bayesian multilevel tobit regression

[R] **tobit** — Tobit regression

[SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)

[SVY] **svy estimation** — Estimation commands for survey data

[XT] **xttobit** — Random-effects tobit models

[U] **20 Estimation and postestimation commands**

Postestimation commands
Remarks and examples

predict
Methods and formulas

margins
Also see

Postestimation commands

The following postestimation commands are of special interest after `metobit`:

Command	Description
<code>estat group</code>	summarize the composition of the nested groups
<code>estat icc</code>	estimate intraclass correlations
<code>estat sd</code>	display variance components as standard deviations and correlations

The following standard postestimation commands are also available:

Command	Description
<code>contrast</code>	contrasts and ANOVA-style joint tests of estimates
<code>estat ic</code>	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
<code>estat summarize</code>	summary statistics for the estimation sample
<code>estat vce</code>	variance–covariance matrix of the estimators (VCE)
<code>estat (svy)</code>	postestimation statistics for survey data
<code>estimates</code>	cataloging estimation results
* <code>hausman</code>	Hausman's specification test
<code>lincom</code>	point estimates, standard errors, testing, and inference for linear combinations of coefficients
* <code>lrtest</code>	likelihood-ratio test
<code>margins</code>	marginal means, predictive margins, marginal effects, and average marginal effects
<code>marginsplot</code>	graph the results from margins (profile plots, interaction plots, etc.)
<code>nlcom</code>	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
<code>predict</code>	predictions, residuals, influence statistics, and other diagnostic measures
<code>predictnl</code>	point estimates, standard errors, testing, and inference for generalized predictions
<code>pwcompare</code>	pairwise comparisons of estimates
<code>test</code>	Wald tests of simple and composite linear hypotheses
<code>testnl</code>	Wald tests of nonlinear hypotheses

* `hausman` and `lrtest` are not appropriate with `svy` estimation results.

predict

Description for predict

`predict` creates a new variable containing predictions such as linear predictions, standard errors, probabilities, and expected values.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining predictions of the outcome and other statistics

```
predict [type] newvarsspec [if] [in] [, statistic options]
```

Syntax for obtaining estimated random effects and their standard errors

```
predict [type] newvarsspec [if] [in], reffects [re_options]
```

Syntax for obtaining ML scores

```
predict [type] newvarsspec [if] [in], scores
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
eta	fitted linear predictor; the default
xb	linear predictor for the fixed portion of the model only
stdp	standard error of the fixed-portion linear prediction
pr(<i>a</i> , <i>b</i>)	$\Pr(a < y < b)$
e(<i>a</i> , <i>b</i>)	$E(y \mid a < y < b)$
ystar(<i>a</i> , <i>b</i>)	$E(y^*), y^* = \max\{a, \min(y, b)\}$

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

where *a* and *b* may be numbers or variables; *a* missing (*a* ≥ .) means $-\infty$, and *b* missing (*b* ≥ .) means $+\infty$; see [U] [12.2.1 Missing values](#).

<i>options</i>	Description
Main	
<u>conditional</u> (<i>ctype</i>)	compute <i>statistic</i> conditional on estimated random effects; default is <code>conditional(ebmeans)</code>
<u>marginal</u>	compute <i>statistic</i> marginally with respect to the random effects
<u>nooffset</u>	make calculation ignoring offset or exposure
Integration	
<u>int_options</u>	integration options
<i>ctype</i>	Description
<u>ebmeans</u>	empirical Bayes means of random effects; the default
<u>ebmodes</u>	empirical Bayes modes of random effects
<u>fixedonly</u>	prediction for the fixed portion of the model only
<i>re_options</i>	Description
Main	
<u>ebmeans</u>	use empirical Bayes means of random effects; the default
<u>ebmodes</u>	use empirical Bayes modes of random effects
<u>reses</u> (<i>stub*</i> <i>newvarlist</i>)	calculate standard errors of empirical Bayes estimates
Integration	
<u>int_options</u>	integration options
<i>int_options</i>	Description
<u>intpoints</u> (#)	use # quadrature points to compute marginal predictions and empirical Bayes means
<u>iterate</u> (#)	set maximum number of iterations in computing statistics involving empirical Bayes estimators
<u>tolerance</u> (#)	set convergence tolerance for computing statistics involving empirical Bayes estimators

Options for predict

Main

`eta`, the default, calculates the fitted linear prediction.

`pr(a,b)` calculates estimates of $\Pr(a < y < b)$, which is the probability that y would be observed in the interval (a, b) .

a and b may be specified as numbers or variable names; lb and ub are variable names; `pr(20,30)` calculates $\Pr(20 < y < 30)$; `pr(lb,ub)` calculates $\Pr(lb < y < ub)$; and `pr(20,ub)` calculates $\Pr(20 < y < ub)$.

a missing ($a \geq .$) means $-\infty$; `pr(.,30)` calculates $\Pr(-\infty < y < 30)$; `pr(lb,30)` calculates $\Pr(-\infty < y < 30)$ in observations for which $lb \geq .$ (and calculates $\Pr(lb < y < 30)$ elsewhere).

`b` missing ($b \geq .$) means $+\infty$; `pr(20,.)` calculates $\Pr(+\infty > y > 20)$;
`pr(20,ub)` calculates $\Pr(+\infty > y > 20)$ in observations for which $ub \geq .$
(and calculates $\Pr(20 < y < ub)$ elsewhere).
`e(a,b)` calculates estimates of $E(y \mid a < y < b)$, which is the expected value of y conditional on y being in the interval (a,b) , meaning that y is truncated. a and b are specified as they are for `pr()`.
`ystar(a,b)` calculates estimates of $E(y^*)$, where $y^* = a$ if $y \leq a$, $y^* = b$ if $y \geq b$, and $y^* = y$ otherwise, meaning that y^* is the censored version of y . a and b are specified as they are for `pr()`.
`xb`, `stdp`, `scores`, `conditional()`, `marginal`, and `nooffset`; see [ME] [meglm postestimation](#).
`reffects`, `ebmeans`, `ebmodes`, and `reses()`; see [ME] [meglm postestimation](#).

Integration

`intpoints()`, `iterate()`, `tolerance()`; see [ME] [meglm postestimation](#).

margins

Description for margins

`margins` estimates margins of response for linear predictions, probabilities, and expected values.

Menu for margins

Statistics > Postestimation

Syntax for margins

<code>margins</code>	<code>[marginlist]</code>	<code>[, options]</code>
<code>margins</code>	<code>[marginlist]</code>	<code>, predict(statistic ...)</code> <code>[predict(statistic ...) ...]</code> <code>[options]</code>
<i>statistic</i>	Description	
<code>eta</code>	fitted linear predictor; the default	
<code>xb</code>	linear predictor for the fixed portion of the model only	
<code>pr(a,b)</code>	$\Pr(a < y < b)$	
<code>e(a,b)</code>	$E(y \mid a < y < b)$	
<code>ystar(a,b)</code>	$E(y^*)$, $y^* = \max\{a, \min(y, b)\}$	
<code>stdp</code>	not allowed with <code>margins</code>	

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.
For the full syntax, see [R] [margins](#).

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed-effects tobit model with `metobit`.

The `predict` command allows us to compute marginal and conditional predictions. Unless stated differently, we use the word “conditional” to mean “conditional on the empirical Bayes predictions of the random effects.” The default prediction is the linear prediction, `eta`, which is the expected value of the unobserved censored variable. Predictions of expected values for censored and truncated versions of the response are also available.

➤ Example 1: Predicting censored and uncensored means

In [example 1](#) of [\[ME\] metobit](#), we analyzed wages for a subpopulation from the National Longitudinal Survey. The dependent variable is the logarithm of wage, and we fit a model that assumes that the data are right-censored at 1.9.


```
. use http://www.stata-press.com/data/r15/nlswork3
(National Longitudinal Survey. Young Women 14-26 years of age in 1968)
. metobit ln_wage union age south##c.grade || idcode:, ul(1.9)
(output omitted)
```

Below, we use `predict` to predict both the mean for the (unobserved) uncensored variable and the (censored) observed values. We also manually generate the censored version of `ln_wage`.

```
. predict uncens_pred, eta marginal
(9310 missing values generated)
. predict cens_pred, ystar(.,1.9) marginal
. generate double ln_wage_cens = min(ln_wage,1.9)
```

To see how the two predictions differ, we can plot them side by side against the censored wage (`ln_wage_cens`).

```
. scatter uncens_pred ln_wage_cens, name(gr1) xsize(4) ysize(4)
. scatter cens_pred ln_wage_cens, name(gr2) xsize(4) ysize(4)
. graph combine gr1 gr2, ycommon
```



We see that many of the predictions for the uncensored variable exceed the censoring point, while the predictions for the censored variable never fall above the upper-censoring limit.



Methods and formulas

Methods and formulas are presented under the following headings:

Introduction
Conditional predictions
Marginal predictions
Marginal variance of the linear predictor

Introduction

This postestimation entry presents the methods and formulas used to calculate the `pr()`, `e()`, and `ystar()` statistics. See [Methods and formulas](#) of [ME] **estat icc** for a discussion of intraclass correlations. See [Methods and formulas](#) of [ME] **meglm postestimation** for a discussion of the remaining postestimation features.

Recall that in a two-level model, the linear predictor for any i th observation in the j th cluster is defined as $\eta_{ij} = \mathbf{x}_{ij}\beta + \mathbf{z}_{ij}\mathbf{u}_j$. Let ℓ_{ij} represent a lower bound for y_{ij} and ul_{ij} represent an upper bound.

Conditional predictions

The probability that $y_{ij}|\hat{\eta}_{ij}$ is observed in the interval (ℓ_{ij}, ul_{ij}) —the `pr(a,b)` option—is calculated as

$$\text{pr}(\ell_{ij}, ul_{ij}) = \Pr(\ell_{ij} < \hat{\eta}_{ij} + \epsilon_{ij} < ul_{ij}) = \Phi\left(\frac{ul_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right) - \Phi\left(\frac{\ell_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right)$$

where $\hat{\sigma}_\epsilon$ is the estimated residual standard deviation.

The `e(a,b)` option computes the expected value of $y_{ij}|\hat{\eta}_{ij}$ conditional on $y_{ij}|\hat{\eta}_{ij}$ being in the interval (ℓ_{ij}, ul_{ij}) , that is, when $y_{ij}|\hat{\eta}_{ij}$ is truncated. The expected value is calculated as

$$\begin{aligned} \text{e}(\ell_{ij}, ul_{ij}) &= E(\hat{\eta}_{ij} + \epsilon_{ij} \mid \ell_{ij} < \hat{\eta}_{ij} + \epsilon_{ij} < ul_{ij}) \\ &= \hat{\eta}_{ij} - \hat{\sigma}_\epsilon \frac{\phi\left(\frac{ul_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right) - \phi\left(\frac{\ell_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right)}{\Phi\left(\frac{ul_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right) - \Phi\left(\frac{\ell_{ij} - \hat{\eta}_{ij}}{\hat{\sigma}_\epsilon}\right)} \end{aligned}$$

where ϕ is the normal density and Φ is the cumulative normal distribution.

You can also compute `ystar(a,b)`—the expected value of $y_{ij}|\hat{\eta}_{ij}$, where y_{ij} is assumed censored at ℓ_{ij} and ul_{ij} :

$$y_{ij}^* = \begin{cases} \ell_{ij} & \text{if } y_{ij} \leq \ell_{ij} \\ \eta_{ij} + \epsilon_{ij} & \text{if } \ell_{ij} < y_{ij} < ul_{ij} \\ ul_{ij} & \text{if } y_{ij} \geq ul_{ij} \end{cases}$$

This computation can be expressed in several ways, but the most intuitive formulation involves a combination of the two statistics just defined:

$$E(y_{ij}^*) = \text{pr}(-\infty, \ell_{ij})\ell_{ij} + \text{pr}(\ell_{ij}, ul_{ij})\text{e}(\ell_{ij}, ul_{ij}) + \text{pr}(ul_{ij}, +\infty)ul_{ij}$$

Marginal predictions

When the `marginal` option is specified, the `pr()` statistic is calculated as

$$\text{pr}(\ell\ell_{ij}, ul_{ij}) = \Phi\left(\frac{ul_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right) - \Phi\left(\frac{\ell\ell_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right)$$

where $\hat{s}_{ij}$ is the square root of the estimated marginal variance of the linear predictor, defined in detail below.

The marginal `e()` statistic is calculated as

$$e(\ell\ell_{ij}, ul_{ij}) = \mathbf{x}_{ij}\hat{\boldsymbol{\beta}} - \hat{s}_{ij} \frac{\phi\left(\frac{ul_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right) - \phi\left(\frac{\ell\ell_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right)}{\Phi\left(\frac{ul_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right) - \Phi\left(\frac{\ell\ell_{ij} - \mathbf{x}_{ij}\hat{\boldsymbol{\beta}}}{\hat{s}_{ij}}\right)}$$

and the marginal `ystar()` statistic is calculated as above with marginal predictions used in place of the conditional ones.

Marginal variance of the linear predictor

In a two-level model, the marginal variance for observation ij is given by

$$\sigma_{ij}^2 = \sigma_{\epsilon}^2 + \mathbf{z}_{ij}\boldsymbol{\Sigma}_2\mathbf{z}_{ij}'$$

where σ_{ϵ}^2 is the residual variance at level 1 and $\boldsymbol{\Sigma}_2$ is the variance matrix of the random effects at level 2. The marginal standard deviation is $s_{ij} = \sqrt{\sigma_{ij}^2}$.

In general, for a G -level random-effects model, the marginal variance for one observation is given by

$$\sigma^2 = \sigma_{\epsilon}^2 + \sum_{g=2}^G \mathbf{z}_g \boldsymbol{\Sigma}_g \mathbf{z}_g'$$

where $\mathbf{z}_g$ is a row vector of the covariates at level g for that observation and $\boldsymbol{\Sigma}_g$ is the variance matrix of the random effects at level g .

Also see

[ME] [metobit](#) — Multilevel mixed-effects tobit regression

[ME] [meglm postestimation](#) — Postestimation tools for meglm

[U] [20 Estimation and postestimation commands](#)

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
Acknowledgments	References	Also see	

Description

`mixed` fits linear mixed-effects models. These models are also known as multilevel models or hierarchical linear models. The overall error distribution of the linear mixed-effects model is assumed to be Gaussian, and heteroskedasticity and correlations within lowest-level groups also may be modeled.

Quick start

Linear mixed-effects model of `y` on `x` with random intercepts by `lev2`

```
mixed y x || lev2:
```

As above, but perform restricted maximum-likelihood (REML) estimation instead of the default maximum likelihood (ML) estimation

```
mixed y x || lev2:, reml
```

As above, but perform small-sample inference on `x` using the Kenward–Roger degrees of freedom (DF) method

```
mixed y x || lev2:, reml dfmethod(kroger)
```

Add random coefficients on `x`

```
mixed y x || lev2: x
```

As above, but allow correlation between the random slopes and intercepts

```
mixed y x || lev2: x, covariance(unstructured)
```

Three-level model with random intercepts by `lev2` and `lev3` for `lev2` nested within `lev3`

```
mixed y x || lev3: || lev2:
```

Crossed-effects model with two-way crossed effects by factors `a` and `b`

```
mixed y x || _all:R.a || b:
```

Menu

Statistics > Multilevel mixed-effects models > Linear regression

Syntax

```
mixed depvar fe_equation [|| re_equation] [|| re_equation ...] [, options]
```

where the syntax of *fe_equation* is

```
[indepvars] [if] [in] [weight] [, fe_options]
```

and the syntax of *re_equation* is one of the following:

for random coefficients and intercepts

```
levelvar: [varlist] [, re_options]
```

for random effects among the values of a factor variable

```
levelvar: R.varname [, re_options]
```

levelvar is a variable identifying the group structure for the random effects at that level or is `_all` representing one group comprising all observations.

<i>fe_options</i>	Description
Model	
<u>noconstant</u>	suppress constant term from the fixed-effects equation

<i>re_options</i>	Description
Model	
<u>covariance</u> (<i>vartype</i>)	variance–covariance structure of the random effects
<u>noconstant</u>	suppress constant term from the random-effects equation
<u>collinear</u>	keep collinear variables
<u>fweight</u> (<i>exp</i>)	frequency weights at higher levels
<u>pweight</u> (<i>exp</i>)	sampling weights at higher levels

<i>options</i>	Description
Model	
<u>mle</u>	fit model via maximum likelihood; the default
<u>reml</u>	fit model via restricted maximum likelihood
<u>dfmethod</u> (<i>df_method</i>)	specify method for computing DF of a <i>t</i> distribution
<u>pwscale</u> (<i>scale_method</i>)	control scaling of sampling weights in two-level models
<u>residuals</u> (<i>rspec</i>)	structure of residual errors
SE/Robust	
<u>vce</u> (<i>vcetype</i>)	<i>vcetype</i> may be <code>oim</code> , <code>robust</code> , or <code>cluster clustvar</code> ; types other than <code>oim</code> may not be combined with <code>dfmethod()</code>
Reporting	
<u>level</u> (#)	set confidence level; default is <code>level(95)</code>
<u>variance</u>	show random-effects and residual-error parameter estimates as variances and covariances; the default
<u>stddeviations</u>	show random-effects and residual-error parameter estimates as standard deviations and correlations
<u>dftable</u> (<i>dftable</i>)	specify contents of fixed-effects table; requires <code>dfmethod()</code> at estimation
<u>norettable</u>	suppress random-effects table
<u>nofetable</u>	suppress fixed-effects table
<u>estmetric</u>	show parameter estimates in the estimation metric
<u>noheader</u>	suppress output header
<u>nogroup</u>	suppress table summarizing groups
<u>nostderr</u>	do not estimate standard errors of random-effects parameters
<u>display_options</u>	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
EM options	
<u>emiterate</u> (#)	number of EM iterations; default is <code>emiterate(20)</code>
<u>emtolerance</u> (#)	EM convergence tolerance; default is <code>emtolerance(1e-10)</code>
<u>emonly</u>	fit model exclusively using EM
<u>emlog</u>	show EM iteration log
<u>emdots</u>	show EM iterations as dots
Maximization	
<u>maximize_options</u>	control the maximization process; seldom used
<u>matsqrt</u>	parameterize variance components using matrix square roots; the default
<u>matlog</u>	parameterize variance components using matrix logarithms
<u>small</u>	replay small-sample inference results
<u>coeflegend</u>	display legend instead of statistics

<i>vartype</i>	Description
<u>independent</u>	one unique variance parameter per random effect, all covariances 0; the default unless the R. notation is used
<u>exchangeable</u>	equal variances for random effects, and one common pairwise covariance
<u>identity</u>	equal variances for random effects, all covariances 0; the default if the R. notation is used
<u>unstructured</u>	all variances and covariances to be distinctly estimated

<i>df_method</i>	Description
<u>residual</u>	residual degrees of freedom, $n - \text{rank}(X)$
<u>repeated</u>	repeated-measures ANOVA
<u>anova</u>	ANOVA
<u>satterthwaite</u> $\left[, \text{dfopts} \right]$	generalized Satterthwaite approximation; REML estimation only
<u>kröger</u> $\left[, \text{dfopts} \right]$	Kenward–Roger; REML estimation only

<i>df_table</i>	Description
<u>default</u>	test statistics, p -values, and confidence intervals; the default
<u>ci</u>	DFS and confidence intervals
<u>pvalue</u>	DFS, test statistics, and p -values

indepvars may contain factor variables; see [U] 11.4.3 Factor variables.

devar, *indepvars*, and *varlist* may contain time-series operators; see [U] 11.4.4 Time-series varlists.

bayes, *bootstrap*, *by*, *jackknife*, *mi estimate*, *rolling*, and *statsby* are allowed; see [U] 11.1.10 Prefix commands. For more details, see [BAYES] bayes: mixed.

mi estimate is not allowed if *dfmethod()* is specified.

Weights are not allowed with the *bootstrap* prefix; see [R] bootstrap.

pweights and *fweights* are allowed; see [U] 11.1.6 weight. However, no weights are allowed if either option *reml* or option *dfmethod()* is specified.

small and *coeflegend* do not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

noconstant suppresses the constant (intercept) term and may be specified for the fixed-effects equation and for any of or all the random-effects equations.

covariance(vartype) specifies the structure of the covariance matrix for the random effects and may be specified for each random-effects equation. *vartype* is one of the following: *independent*, *exchangeable*, *identity*, or *unstructured*.

independent allows for a distinct variance for each random effect within a random-effects equation and assumes that all covariances are 0.

exchangeable specifies one common variance for all random effects and one common pairwise covariance.

identity is short for “multiple of the identity”; that is, all variances are equal and all covariances are 0.

unstructured allows for all variances and covariances to be distinct. If an equation consists of p random-effects terms, the unstructured covariance matrix will have $p(p + 1)/2$ unique parameters.

covariance(independent) is the default, except when the R. notation is used, in which case **covariance(identity)** is the default and only **covariance(identity)** and **covariance(exchangeable)** are allowed.

collinear specifies that **mixed** not omit collinear variables from the random-effects equation. Usually, there is no reason to leave collinear variables in place; in fact, doing so usually causes the estimation to fail because of the matrix singularity caused by the collinearity. However, with certain models (for example, a random-effects model with a full set of contrasts), the variables may be collinear, yet the model is fully identified because of restrictions on the random-effects covariance structure. In such cases, using the **collinear** option allows the estimation to take place with the random-effects equation intact.

fweight(*exp*) specifies frequency weights at higher levels in a multilevel model, whereas frequency weights at the first level (the observation level) are specified in the usual manner, for example, [**fw=fwtvar1**]. *exp* can be any valid Stata variable, and you can specify **fweight()** at levels two and higher of a multilevel model. For example, in the two-level model

```
. mixed fixed_portion [fw = wt1] || school: ..., fweight(wt2) ...
```

the variable **wt1** would hold the first-level (the observation-level) frequency weights, and **wt2** would hold the second-level (the school-level) frequency weights.

pweight(*exp*) specifies sampling weights at higher levels in a multilevel model, whereas sampling weights at the first level (the observation level) are specified in the usual manner, for example, [**pw=pwtvar1**]. *exp* can be any valid Stata variable, and you can specify **pweight()** at levels two and higher of a multilevel model. For example, in the two-level model

```
. mixed fixed_portion [pw = wt1] || school: ..., pweight(wt2) ...
```

variable **wt1** would hold the first-level (the observation-level) sampling weights, and **wt2** would hold the second-level (the school-level) sampling weights.

See [Survey data](#) in *Remarks and examples* below for more information regarding the use of sampling weights in multilevel models.

mle and **reml** specify the statistical method for fitting the model.

mle, the default, specifies that the model be fit using ML. Options **dfmethod(satterthwaite)** and **dfmethod(kroger)** are not supported under ML estimation.

reml specifies that the model be fit using REML, also known as residual maximum likelihood.

dfmethod(*df_method*) requests that reported hypothesis tests for the fixed effects (coefficients) use a small-sample adjustment. By default, inference is based on a large-sample approximation of the sampling distributions of the test statistics by normal and χ^2 distributions. Caution should be exercised when choosing a DF method; see [Small-sample inference for fixed effects](#) in *Remarks and examples* for details.

When `dfmethod(df_method)` is specified, the sampling distributions of the test statistics are approximated by a t distribution, according to the requested method for computing the DF. `df_method` is one of the following: `residual`, `repeated`, `anova`, `satterthwaite`, or `kroger`.

`residual` uses the residual degrees of freedom, $n - \text{rank}(X)$, as the DF for all tests of fixed effects. For a linear model without random effects with independent and identically distributed (i.i.d.) errors, the distributions of the test statistics for fixed effects are t distributions with the residual DF. For other mixed-effects models, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

`repeated` uses the repeated-measures ANOVA method for computing the DF. It is used with balanced repeated-measures designs with spherical correlation error structures. It partitions the residual degrees of freedom into the between-subject degrees of freedom and the within-subject degrees of freedom. `repeated` is supported only with two-level models. For more complex mixed-effects models or with unbalanced data, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

`anova` uses the traditional ANOVA method for computing the DF. According to this method, the DF for a test of a fixed effect of a given variable depends on whether that variable is also included in any of the random-effects equations. For traditional ANOVA models with balanced designs, this method provides exact sampling distributions of the test statistics. For more complex mixed-effects models or with unbalanced data, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

`satterthwaite[, dfopts]` implements a generalization of the [Satterthwaite \(1946\)](#) approximation of the unknown sampling distributions of test statistics for complex linear mixed-effect models. This method is supported only with REML estimation.

`kroger[, dfopts]` implements the [Kenward and Roger \(1997\)](#) method, which is designed to approximate unknown sampling distributions of test statistics for complex linear mixed-effects models. This method is supported only with REML estimation.

`dfopts` is either `eim` or `oim`.

`eim` specifies that the expected information matrix be used to compute Satterthwaite or Kenward–Roger degrees of freedom. This is the default.

`oim` specifies that the observed information matrix be used to compute Satterthwaite or Kenward–Roger degrees of freedom.

Residual, repeated, and ANOVA methods are suitable only when the sampling distributions of the test statistics are known to be t or F . This is usually only known for certain classes of linear mixed-effects models with simple covariance structures and when data are balanced. These methods are available with both ML and REML estimation.

For unbalanced data or balanced data with complicated covariance structures, the sampling distributions of the test statistics are unknown and can only be approximated. The Satterthwaite and Kenward–Roger methods provide approximations to the distributions in these cases. According to [Schaalje, McBride, and Fellingham \(2002\)](#), the Kenward–Roger method should, in general, be preferred to the Satterthwaite method. However, there are situations in which the two methods are expected to perform similarly, such as with compound symmetry covariance structures. The Kenward–Roger method is more computationally demanding than the Satterthwaite method. Both methods are available only with REML estimation. See [Small-sample inference for fixed effects in Remarks and examples](#) for examples and more detailed descriptions of the DF methods.

`dfmethod()` may not be combined with weighted estimation, the `mi estimate` prefix, or `vce()`, unless it is the default `vce(oim)`.

`pwscale(scale_method)` controls how sampling weights (if specified) are scaled in two-level models. `scale_method` is one of the following: `size`, `effective`, or `gk`.

`size` specifies that first-level (observation-level) weights be scaled so that they sum to the sample size of their corresponding second-level cluster. Second-level sampling weights are left unchanged.

`effective` specifies that first-level weights be scaled so that they sum to the effective sample size of their corresponding second-level cluster. Second-level sampling weights are left unchanged.

`gk` specifies the [Graubard and Korn \(1996\)](#) method. Under this method, second-level weights are set to the cluster averages of the products of the weights at both levels, and first-level weights are then set equal to 1.

`pwscale()` is supported only with two-level models. See [Survey data](#) in *Remarks and examples* below for more details on using `pwscale()`. `pwscale()` may not be combined with the `dfmethod()` option.

`residuals(rspec)` specifies the structure of the residual errors within the lowest-level groups (the second level of a multilevel model with the observations comprising the first level) of the linear mixed model. For example, if you are modeling random effects for classes nested within schools, then `residuals()` refers to the residual variance–covariance structure of the observations within classes, the lowest-level groups. `rspec` has the following syntax:

$$\text{restyle} \left[, \text{residual\_options} \right]$$

`restyle` is one of the following: `independent`, `exchangeable`, `ar #`, `ma #`, `unstructured banded #`, `toeplitz #`, or `exponential`.

`independent`, the default, specifies that all residuals be i.i.d. Gaussian with one common variance. When combined with `by(varname)`, independence is still assumed, but you estimate a distinct variance for each level of `varname`. Unlike with the structures described below, `varname` does not need to be constant within groups.

`exchangeable` estimates two parameters, one common within-group variance and one common pairwise covariance. When combined with `by(varname)`, these two parameters are distinctly estimated for each level of `varname`. Because you are modeling a within-group covariance, `varname` must be constant within lowest-level groups.

`ar #` assumes that within-group errors have an autoregressive (AR) structure of order `#`; `ar 1` is the default. The `t(varname)` option is required, where `varname` is an integer-valued time variable used to order the observations within groups and to determine the lags between successive observations. Any nonconsecutive time values will be treated as gaps. For this structure, `# + 1` parameters are estimated (`#` AR coefficients and one overall error variance). `restyle ar` may be combined with `by(varname)`, but `varname` must be constant within groups.

`ma #` assumes that within-group errors have a moving-average (MA) structure of order `#`; `ma 1` is the default. The `t(varname)` option is required, where `varname` is an integer-valued time variable used to order the observations within groups and to determine the lags between successive observations. Any nonconsecutive time values will be treated as gaps. For this structure, `# + 1` parameters are estimated (`#` MA coefficients and one overall error variance). `restyle ma` may be combined with `by(varname)`, but `varname` must be constant within groups.

`unstructured` is the most general structure; it estimates distinct variances for each within-group error and distinct covariances for each within-group error pair. The `t(varname)` option is required, where `varname` is a nonnegative-integer-valued variable that identifies the observations within each group. The groups may be unbalanced in that not all levels of `t()` need to be observed within every group, but you may not have repeated `t()` values within any particular group. When you have p levels of `t()`, then $p(p + 1)/2$ parameters are estimated. `restyle`

`unstructured` may be combined with `by(varname)`, but `varname` must be constant within groups.

`banded #` is a special case of `unstructured` that restricts estimation to the covariances within the first `#` off-diagonals and sets the covariances outside this band to 0. The `t(varname)` option is required, where `varname` is a nonnegative-integer-valued variable that identifies the observations within each group. `#` is an integer between 0 and $p - 1$, where p is the number of levels of `t()`. By default, `#` is $p - 1$; that is, all elements of the covariance matrix are estimated. When `#` is 0, only the diagonal elements of the covariance matrix are estimated. `restype banded` may be combined with `by(varname)`, but `varname` must be constant within groups.

`toeplitz #` assumes that within-group errors have Toeplitz structure of order `#`, for which correlations are constant with respect to time lags less than or equal to `#` and are 0 for lags greater than `#`. The `t(varname)` option is required, where `varname` is an integer-valued time variable used to order the observations within groups and to determine the lags between successive observations. `#` is an integer between 1 and the maximum observed lag (the default). Any nonconsecutive time values will be treated as gaps. For this structure, `# + 1` parameters are estimated (`#` correlations and one overall error variance). `restype toeplitz` may be combined with `by(varname)`, but `varname` must be constant within groups.

`exponential` is a generalization of the AR covariance model that allows for unequally spaced and noninteger time values. The `t(varname)` option is required, where `varname` is real-valued. For the exponential covariance model, the correlation between two errors is the parameter ρ , raised to a power equal to the absolute value of the difference between the `t()` values for those errors. For this structure, two parameters are estimated (the correlation parameter ρ and one overall error variance). `restype exponential` may be combined with `by(varname)`, but `varname` must be constant within groups.

`residual_options` are `by(varname)` and `t(varname)`.

`by(varname)` is for use within the `residuals()` option and specifies that a set of distinct residual-error parameters be estimated for each level of `varname`. In other words, you use `by()` to model heteroskedasticity.

`t(varname)` is for use within the `residuals()` option to specify a time variable for the `ar`, `ma`, `toeplitz`, and `exponential` structures, or to identify the observations when `restype` is `unstructured` or `banded`.

SE/Robust

`vce(vcetype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`oim`), that are robust to some kinds of misspecification (`robust`), and that allow for intragroup correlation (`cluster clustvar`); see [R] [vce_option](#). If `vce(robust)` is specified, robust variances are clustered at the highest level in the multilevel model.

`vce(robust)` and `vce(cluster clustvar)` are not supported with REML estimation. Only `vce(oim)` is allowed in combination with `dfmethod()`.

Reporting

`level(#)`; see [R] [estimation options](#).

`variance`, the default, displays the random-effects and residual-error parameter estimates as variances and covariances.

`stddeviations` displays the random-effects and residual-error parameter estimates as standard deviations and correlations.

`dftable(dftable)` specifies the contents of the fixed-effects table for small-sample inference when `dfmethod()` is used during estimation. *dftable* is one of the following: `default`, `ci`, or `pvalue`.

`default` displays the default standard fixed-effects table that contains test statistics, *p*-values, and confidence intervals.

`ci` displays the fixed-effects table in which the columns containing statistics and *p*-values are replaced with a column containing coefficient-specific DFs. Confidence intervals are also displayed.

`pvalue` displays the fixed-effects table that includes a column containing DFs with the standard columns containing test statistics and *p*-values. Confidence intervals are not displayed.

`norettable` suppresses the random-effects table from the output.

`nofetable` suppresses the fixed-effects table from the output.

`estmetric` displays all parameter estimates in the estimation metric. Fixed-effects estimates are unchanged from those normally displayed, but random-effects parameter estimates are displayed as log-standard deviations and hyperbolic arctangents of correlations, with equation names that organize them by model level. Residual-variance parameter estimates are also displayed in their original estimation metric.

`noheader` suppresses the output header, either at estimation or upon replay.

`nogroup` suppresses the display of group summary information (number of groups, average group size, minimum, and maximum) from the output header.

`nostderr` prevents `mixed` from calculating standard errors for the estimated random-effects parameters, although standard errors are still provided for the fixed-effects parameters. Specifying this option will speed up computation times. `nostderr` is available only when residuals are modeled as independent with constant variance.

display_options: `nocl`, `nopvalues`, `noomitted`, `vsquish`, `noemptycells`, `baselevels`, `allbaselevels`, `nofvlabel`, `fvwrap(#)`, `fvwrapon(style)`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nolstretch`; see [R] [estimation options](#).

EM options

These options control the expectation-maximization (EM) iterations that take place before estimation switches to a gradient-based method. When residuals are modeled as independent with constant variance, EM will either converge to the solution or bring parameter estimates close to the solution. For other residual structures or for weighted estimation, EM is used to obtain starting values.

`emiterate(#)` specifies the number of EM iterations to perform. The default is `emiterate(20)`.

`emtolerance(#)` specifies the convergence tolerance for the EM algorithm. The default is `emtolerance(1e-10)`. EM iterations will be halted once the log (restricted) likelihood changes by a relative amount less than `#`. At that point, optimization switches to a gradient-based method, unless `emonly` is specified, in which case maximization stops.

`emonly` specifies that the likelihood be maximized exclusively using EM. The advantage of specifying `emonly` is that EM iterations are typically much faster than those for gradient-based methods. The disadvantages are that EM iterations can be slow to converge (if at all) and that EM provides no facility for estimating standard errors for the random-effects parameters. `emonly` is available only with unweighted estimation and when residuals are modeled as independent with constant variance.

`emlog` specifies that the EM iteration log be shown. The EM iteration log is, by default, not displayed unless the `emonly` option is specified.

`emdots` specifies that the EM iterations be shown as dots. This option can be convenient because the EM algorithm may require many iterations to converge.

Maximization

`maximize_options`: `difficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, and `nonrtolerance`; see [R] [maximize](#). Those that require special mention for `mixed` are listed below.

For the `technique()` option, the default is `technique(nr)`. The `bhhh` algorithm may not be specified.

`matsqrt` (the default), during optimization, parameterizes variance components by using the matrix square roots of the variance–covariance matrices formed by these components at each model level.

`matlog`, during optimization, parameterizes variance components by using the matrix logarithms of the variance–covariance matrices formed by these components at each model level.

The `matsqrt` parameterization ensures that variance–covariance matrices are positive semidefinite, while `matlog` ensures matrices that are positive definite. For most problems, the matrix square root is more stable near the boundary of the parameter space. However, if convergence is problematic, one option may be to try the alternate `matlog` parameterization. When convergence is not an issue, both parameterizations yield equivalent results.

The following options are available with `mixed` but are not shown in the dialog box:

`small` replays previously obtained small-sample results. This option is available only upon replay and requires that the `dfmethod()` option be used during estimation. `small` is equivalent to `dftable(default)` upon replay.

`coeflegend`; see [R] [estimation options](#).

Remarks and examples

Remarks are presented under the following headings:

- [Introduction](#)
- [Two-level models](#)
- [Covariance structures](#)
- [Likelihood versus restricted likelihood](#)
- [Three-level models](#)
- [Blocked-diagonal covariance structures](#)
- [Heteroskedastic random effects](#)
- [Heteroskedastic residual errors](#)
- [Other residual-error structures](#)
- [Crossed-effects models](#)
- [Diagnosing convergence problems](#)
- [Survey data](#)
- [Small-sample inference for fixed effects](#)

Introduction

Linear mixed models are models containing both fixed effects and random effects. They are a generalization of linear regression allowing for the inclusion of random deviations (effects) other than those associated with the overall error term. In matrix notation,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon} \quad (1)$$

where $\mathbf{y}$ is the $n \times 1$ vector of responses, $\mathbf{X}$ is an $n \times p$ design/covariate matrix for the fixed effects $\boldsymbol{\beta}$, and $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. The $n \times 1$ vector of errors $\boldsymbol{\epsilon}$ is assumed to be multivariate normal with mean 0 and variance matrix $\sigma_\epsilon^2 \mathbf{R}$.

The fixed portion of (1), $\mathbf{X}\boldsymbol{\beta}$, is analogous to the linear predictor from a standard OLS regression model with $\boldsymbol{\beta}$ being the regression coefficients to be estimated. For the random portion of (1), $\mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}$, we assume that $\mathbf{u}$ has variance–covariance matrix $\mathbf{G}$ and that $\mathbf{u}$ is orthogonal to $\boldsymbol{\epsilon}$ so that

$$\text{Var} \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\epsilon} \end{bmatrix} = \begin{bmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \sigma_\epsilon^2 \mathbf{R} \end{bmatrix}$$

The random effects $\mathbf{u}$ are not directly estimated (although they may be predicted), but instead are characterized by the elements of $\mathbf{G}$, known as variance components, that are estimated along with the overall residual variance σ_ϵ^2 and the residual-variance parameters that are contained within $\mathbf{R}$.

The general forms of the design matrices $\mathbf{X}$ and $\mathbf{Z}$ allow estimation for a broad class of linear models: blocked designs, split-plot designs, growth curves, multilevel or hierarchical designs, etc. They also allow a flexible method of modeling within-cluster correlation. Subjects within the same cluster can be correlated as a result of a shared random intercept, or through a shared random slope on (say) age, or both. The general specification of $\mathbf{G}$ also provides additional flexibility—the random intercept and random slope could themselves be modeled as independent, or correlated, or independent with equal variances, and so forth. The general structure of $\mathbf{R}$ also allows for residual errors to be heteroskedastic and correlated, and allows flexibility in exactly how these characteristics can be modeled.

Comprehensive treatments of mixed models are provided by, among others, [Searle, Casella, and McCulloch \(1992\)](#); [McCulloch, Searle, and Neuhaus \(2008\)](#); [Verbeke and Molenberghs \(2000\)](#); [Raudenbush and Bryk \(2002\)](#); and [Pinheiro and Bates \(2000\)](#). In particular, chapter 2 of [Searle, Casella, and McCulloch \(1992\)](#) provides an excellent history.

The key to fitting mixed models lies in estimating the variance components, and for that there exist many methods. Most of the early literature in mixed models dealt with estimating variance components in ANOVA models. For simple models with balanced data, estimating variance components amounts to solving a system of equations obtained by setting expected mean-squares expressions equal to their observed counterparts. Much of the work in extending the ANOVA method to unbalanced data for general ANOVA designs is due to [Henderson \(1953\)](#).

The ANOVA method, however, has its shortcomings. Among these is a lack of uniqueness in that alternative, unbiased estimates of variance components could be derived using other quadratic forms of the data in place of observed mean squares ([Searle, Casella, and McCulloch 1992](#), 38–39). As a result, ANOVA methods gave way to more modern methods, such as minimum norm quadratic unbiased estimation (MINQUE) and minimum variance quadratic unbiased estimation (MIVQUE); see [Rao \(1973\)](#) for MINQUE and [LaMotte \(1973\)](#) for MIVQUE. Both methods involve finding optimal quadratic forms of the data that are unbiased for the variance components.

The most popular methods, however, are ML and REML, and these are the two methods that are supported by `mixed`. The ML estimates are based on the usual application of likelihood theory, given

the distributional assumptions of the model. The basic idea behind REML (Thompson 1962) is that you can form a set of linear contrasts of the response that do not depend on the fixed effects β , but instead depend only on the variance components to be estimated. You then apply ML methods by using the distribution of the linear contrasts to form the likelihood.

Returning to (1): in clustered-data situations, it is convenient not to consider all n observations at once but instead to organize the mixed model as a series of M independent groups or clusters

$$\mathbf{y}_j = \mathbf{X}_j\beta + \mathbf{Z}_j\mathbf{u}_j + \epsilon_j \quad (2)$$

for $j = 1, \dots, M$, with cluster j consisting of n_j observations. The response $\mathbf{y}_j$ comprises the rows of $\mathbf{y}$ corresponding with the j th cluster, with $\mathbf{X}_j$ and ϵ_j defined analogously. The random effects $\mathbf{u}_j$ can now be thought of as M realizations of a $q \times 1$ vector that is normally distributed with mean $\mathbf{0}$ and $q \times q$ variance matrix Σ . The matrix $\mathbf{Z}_i$ is the $n_j \times q$ design matrix for the j th cluster random effects. Relating this to (1), note that

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{Z}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{Z}_M \end{bmatrix}; \quad \mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_M \end{bmatrix}; \quad \mathbf{G} = \mathbf{I}_M \otimes \Sigma; \quad \mathbf{R} = \mathbf{I}_M \otimes \Lambda \quad (3)$$

The mixed-model formulation (2) is from Laird and Ware (1982) and offers two key advantages. First, it makes specifications of random-effects terms easier. If the clusters are schools, you can simply specify a random effect at the school level, as opposed to thinking of what a school-level random effect would mean when all the data are considered as a whole (if it helps, think Kronecker products). Second, representing a mixed-model with (2) generalizes easily to more than one set of random effects. For example, if classes are nested within schools, then (2) can be generalized to allow random effects at both the school and the class-within-school levels. This we demonstrate later.

In the sections that follow, we assume that residuals are independent with constant variance; that is, in (3) we treat Λ equal to the identity matrix and limit ourselves to estimating one overall residual variance, σ_ϵ^2 . Beginning in *Heteroskedastic residual errors*, we relax this assumption.


Two-level models

We begin with a simple application of (2) as a two-level model, because a one-level linear model, by our terminology, is just standard OLS regression.

► Example 1: Two-level random intercept model

Consider a longitudinal dataset, used by both Ruppert, Wand, and Carroll (2003) and Diggle et al. (2002), consisting of weight measurements of 48 pigs on 9 successive weeks. Pigs are identified by the variable `id`. Below is a plot of the growth curves for the first 10 pigs.

```
. use http://www.stata-press.com/data/r15/pig
(Longitudinal analysis of pig weights)
. twoway connected weight week if id<=10, connect(L)
```



It seems clear that each pig experiences a linear trend in growth and that overall weight measurements vary from pig to pig. Because we are not really interested in these particular 48 pigs per se, we instead treat them as a random sample from a larger population and model the between-pig variability as a random effect, or in the terminology of (2), as a random-intercept term at the pig level. We thus wish to fit the model

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_j + \epsilon_{ij} \tag{4}$$

for $i = 1, \dots, 9$ weeks and $j = 1, \dots, 48$ pigs. The fixed portion of the model, $\beta_0 + \beta_1 \text{week}_{ij}$, simply states that we want one overall regression line representing the population average. The random effect u_j serves to shift this regression line up or down according to each pig. Because the random effects occur at the pig level (id), we fit the model by typing

```
. mixed weight week || id:
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -1014.9268
Iteration 1:   log likelihood = -1014.9268
Computing standard errors:
Mixed-effects ML regression              Number of obs      =      432
Group variable: id                      Number of groups   =      48
                                         Obs per group:
                                             min =          9
                                             avg =         9.0
                                             max =          9
                                         Wald chi2(1)       =    25337.49
                                         Prob > chi2        =      0.0000
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0390124	159.18	0.000	6.133433	6.286359
_cons	19.35561	.5974059	32.40	0.000	18.18472	20.52651

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(_cons)	14.81751	3.124226	9.801716	22.40002
	var(Residual)	4.383264	.3163348	3.805112	5.04926

```
LR test vs. linear model: chibar2(01) = 472.65          Prob >= chibar2 = 0.0000
```

Notes:

1. By typing `weight week`, we specified the response, `weight`, and the fixed portion of the model in the same way that we would if we were using `regress` or any other estimation command. Our fixed effects are a coefficient on `week` and a constant term.
2. When we added `|| id:`, we specified random effects at the level identified by the group variable `id`, that is, the pig level (level two). Because we wanted only a random intercept, that is all we had to type.
3. The estimation log consists of three parts:
 - a. A set of EM iterations used to refine starting values. By default, the iterations themselves are not displayed, but you can display them with the `emlog` option.
 - b. A set of gradient-based iterations. By default, these are Newton–Raphson iterations, but other methods are available by specifying the appropriate `maximize` options; see [\[R\] maximize](#).
 - c. The message “Computing standard errors”. This is just to inform you that `mixed` has finished its iterative maximization and is now reparameterizing from a matrix-based parameterization (see [Methods and formulas](#)) to the natural metric of variance components and their estimated standard errors.
4. The output title, “Mixed-effects ML regression”, informs us that our model was fit using ML, the default. For REML estimates, use the `reml` option.

Because this model is a simple random-intercept model fit by ML, it would be equivalent to using `xtreg` with its `mle` option.

5. The first estimation table reports the fixed effects. We estimate $\beta_0 = 19.36$ and $\beta_1 = 6.21$.

6. The second estimation table shows the estimated variance components. The first section of the table is labeled `id: Identity`, meaning that these are random effects at the `id` (pig) level and that their variance–covariance matrix is a multiple of the identity matrix; that is, $\Sigma = \sigma_u^2 \mathbf{I}$. Because we have only one random effect at this level, `mixed` knew that `Identity` is the only possible covariance structure. In any case, the variance of the level-two errors, σ_u^2 , is estimated as 14.82 with standard error 3.12.
7. The row labeled `var(Residual)` displays the estimated variance of the overall error term; that is, $\hat{\sigma}_\epsilon^2 = 4.38$. This is the variance of the level-one errors, that is, the residuals.
8. Finally, a likelihood-ratio test comparing the model with one-level ordinary linear regression, model (4) without u_j , is provided and is highly significant for these data.

We now store our estimates for later use:

```
. estimates store randint
```

► Example 2: Two-level random slope model

Extending (4) to allow for a random slope on `week` yields the model

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_{0j} + u_{1j} \text{week}_{ij} + \epsilon_{ij} \tag{5}$$

and we fit this with `mixed`:

```
. mixed weight week || id: week
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -869.03825
Iteration 1:   log likelihood = -869.03825
Computing standard errors:
Mixed-effects ML regression              Number of obs      =      432
Group variable: id                      Number of groups   =      48
                                         Obs per group:
                                         min =              9
                                         avg =             9.0
                                         max =              9
                                         Wald chi2(1)      =    4689.51
                                         Prob > chi2       =      0.0000

Log likelihood = -869.03825
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0906819	68.48	0.000	6.032163	6.387629
_cons	19.35561	.3979159	48.64	0.000	18.57571	20.13551

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Independent	var(week)	.3680668	.0801181	.2402389	.5639103
	var(_cons)	6.756364	1.543503	4.317721	10.57235
	var(Residual)	1.598811	.1233988	1.374358	1.85992

```
LR test vs. linear model: chi2(2) = 764.42          Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
. estimates store randslope
```

Because we did not specify a covariance structure for the random effects $(u_{0j}, u_{1j})'$, `mixed` used the default `Independent` structure; that is,

$$\Sigma = \text{Var} \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} = \begin{bmatrix} \sigma_{u0}^2 & 0 \\ 0 & \sigma_{u1}^2 \end{bmatrix} \quad (6)$$

with $\hat{\sigma}_{u0}^2 = 6.76$ and $\hat{\sigma}_{u1}^2 = 0.37$. Our point estimates of the fixed effects are essentially identical to those from model (4), but note that this does not hold generally. Given the 95% confidence interval for $\hat{\sigma}_{u1}^2$, it would seem that the random slope is significant, and we can use `lrtest` and our two stored estimation results to verify this fact:

```
. lrtest randslope randint
Likelihood-ratio test                                LR chi2(1) =    291.78
(Assumption: randint nested in randslope)            Prob > chi2 =     0.0000
Note: The reported degrees of freedom assumes the null hypothesis is not on
      the boundary of the parameter space.  If this is not true, then the
      reported test is conservative.
```

The near-zero significance level favors the model that allows for a random pig-specific regression line over the model that allows only for a pig-specific shift.

◀

Covariance structures

In [example 2](#), we fit a model with the default `Independent` covariance given in (6). Within any random-effects level specification, we can override this default by specifying an alternative covariance structure via the `covariance()` option.

► Example 3: Two-level model with correlated random effects

We generalize (6) to allow u_{0j} and u_{1j} to be correlated; that is,

$$\Sigma = \text{Var} \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} = \begin{bmatrix} \sigma_{u0}^2 & \sigma_{01} \\ \sigma_{01} & \sigma_{u1}^2 \end{bmatrix}$$


```
. mixed weight week || id: week, covariance(unstructured)
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:    log likelihood = -868.96185
Iteration 1:    log likelihood = -868.96185
Computing standard errors:
Mixed-effects ML regression              Number of obs      =      432
Group variable: id                      Number of groups   =      48
                                         Obs per group:
                                         min =              9
                                         avg =             9.0
                                         max =              9
                                         Wald chi2(1)      =    4649.17
                                         Prob > chi2       =      0.0000

Log likelihood = -868.96185
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0910745	68.18	0.000	6.031393	6.388399
_cons	19.35561	.3996387	48.43	0.000	18.57234	20.13889

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Unstructured					
	var(week)	.3715251	.0812958	.2419532	.570486
	var(_cons)	6.823363	1.566194	4.351297	10.69986
	cov(week,_cons)	-.0984378	.2545767	-.5973991	.4005234
	var(Residual)	1.596829	.123198	1.372735	1.857505

LR test vs. linear model: chi2(3) = 764.58 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

But we do not find the correlation to be at all significant.

```
. lrtest . randslope
Likelihood-ratio test                      LR chi2(1) =      0.15
(Assumption: randslope nested in .)       Prob > chi2 =    0.6959
```



Instead, we could have also specified `covariance(identity)`, restricting u_{0j} and u_{1j} to not only be independent but also to have common variance, or we could have specified `covariance(exchangeable)`, which imposes a common variance but allows for a nonzero correlation.

Likelihood versus restricted likelihood

Thus far, all our examples have used ML to estimate variance components. We could have just as easily asked for REML estimates. Refitting the model in [example 2](#) by REML, we get

```

. mixed weight week || id: week, reml
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log restricted-likelihood = -870.51473
Iteration 1:   log restricted-likelihood = -870.51473
Computing standard errors:
Mixed-effects REML regression          Number of obs      =      432
Group variable: id                    Number of groups   =      48
                                      Obs per group:
                                          min =          9
                                          avg =         9.0
                                          max =          9
                                      Wald chi2(1)         =    4592.10
Log restricted-likelihood = -870.51473   Prob > chi2         =      0.0000

```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0916387	67.77	0.000	6.030287	6.389504
_cons	19.35561	.4021144	48.13	0.000	18.56748	20.14374

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Independent	var(week)	.3764405	.0827027	.2447317	.5790317
	var(_cons)	6.917604	1.593247	4.404624	10.86432
	var(Residual)	1.598784	.1234011	1.374328	1.859898

```

LR test vs. linear model: chi2(2) = 765.92          Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

```

Although ML estimators are based on the usual likelihood theory, the idea behind REML is to transform the response into a set of linear contrasts whose distribution is free of the fixed effects β . The restricted likelihood is then formed by considering the distribution of the linear contrasts. Not only does this make the maximization problem free of β , it also incorporates the degrees of freedom used to estimate β into the estimation of the variance components. This follows because, by necessity, the rank of the linear contrasts must be less than the number of observations.

As a simple example, consider a constant-only regression where $y_i \sim N(\mu, \sigma^2)$ for $i = 1, \dots, n$. The ML estimate of σ^2 can be derived theoretically as the n -divided sample variance. The REML estimate can be derived by considering the first $n - 1$ error contrasts, $y_i - \bar{y}$, whose joint distribution is free of μ . Applying maximum likelihood to this distribution results in an estimate of σ^2 , that is, the $(n - 1)$ -divided sample variance, which is unbiased for σ^2 .

The unbiasedness property of REML extends to all mixed models when the data are balanced, and thus REML would seem the clear choice in balanced-data problems, although in large samples the difference between ML and REML is negligible. One disadvantage of REML is that likelihood-ratio (LR) tests based on REML are inappropriate for comparing models with different fixed-effects specifications. ML is appropriate for such LR tests and has the advantage of being easy to explain and being the method of choice for other estimators.

Another factor to consider is that ML estimation under `mixed` is more feature-rich, allowing for weighted estimation and robust variance-covariance matrices, features not supported under REML. In the end, which method to use should be based both on your needs and on personal taste.

Examining the REML output, we find that the estimates of the variance components are slightly larger than the ML estimates. This is typical, because ML estimates, which do not incorporate the degrees of freedom used to estimate the fixed effects, tend to be biased downward.

Three-level models

The clustered-data representation of the mixed model given in (2) can be extended to two nested levels of clustering, creating a three-level model once the observations are considered. Formally,

$$\mathbf{y}_{jk} = \mathbf{X}_{jk}\boldsymbol{\beta} + \mathbf{Z}_{jk}^{(3)}\mathbf{u}_k^{(3)} + \mathbf{Z}_{jk}^{(2)}\mathbf{u}_{jk}^{(2)} + \boldsymbol{\epsilon}_{jk} \quad (7)$$

for $i = 1, \dots, n_{jk}$ first-level observations nested within $j = 1, \dots, M_k$ second-level groups, which are nested within $k = 1, \dots, M$ third-level groups. Group j, k consists of n_{jk} observations, so $\mathbf{y}_{jk}$, $\mathbf{X}_{jk}$, and $\boldsymbol{\epsilon}_{jk}$ each have row dimension n_{jk} . $\mathbf{Z}_{jk}^{(3)}$ is the $n_{jk} \times q_3$ design matrix for the third-level random effects $\mathbf{u}_k^{(3)}$, and $\mathbf{Z}_{jk}^{(2)}$ is the $n_{jk} \times q_2$ design matrix for the second-level random effects $\mathbf{u}_{jk}^{(2)}$. Furthermore, assume that

$$\mathbf{u}_k^{(3)} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_3); \quad \mathbf{u}_{jk}^{(2)} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_2); \quad \boldsymbol{\epsilon}_{jk} \sim N(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I})$$

and that $\mathbf{u}_k^{(3)}$, $\mathbf{u}_{jk}^{(2)}$, and $\boldsymbol{\epsilon}_{jk}$ are independent.

Fitting a three-level model requires you to specify two random-effects equations: one for level three and then one for level two. The variable list for the first equation represents $\mathbf{Z}_{jk}^{(3)}$ and for the second equation represents $\mathbf{Z}_{jk}^{(2)}$; that is, you specify the levels top to bottom in `mixed`.

► Example 4: Three-level model with random intercepts

Baltagi, Song, and Jung (2001) estimate a Cobb–Douglas production function examining the productivity of public capital in each state’s private output. Originally provided by Munnell (1990), the data were recorded over 1970–1986 for 48 states grouped into nine regions.

```
. use http://www.stata-press.com/data/r15/productivity
(Public Capital Productivity)

. describe

Contains data from http://www.stata-press.com/data/r15/productivity.dta
   obs:           816                Public Capital Productivity
  vars:           11                29 Mar 2016 10:57
 size:        29,376                (_dta has notes)
```

variable name	storage type	display format	value label
state	byte	%9.0g	states 1-48
region	byte	%9.0g	regions 1-9
year	int	%9.0g	years 1970-1986
public	float	%9.0g	public capital stock
hwy	float	%9.0g	log(highway component of public)
water	float	%9.0g	log(water component of public)
other	float	%9.0g	log(bldg/other component of public)
private	float	%9.0g	log(private capital stock)
gsp	float	%9.0g	log(gross state product)
emp	float	%9.0g	log(non-agriculture payrolls)
unemp	float	%9.0g	state unemployment rate

Sorted by:

Because the states are nested within regions, we fit a three-level mixed model with random intercepts at both the region and the state-within-region levels. That is, we use (7) with both $\mathbf{Z}_{jk}^{(3)}$ and $\mathbf{Z}_{jk}^{(2)}$ set to the $n_{jk} \times 1$ column of ones, and $\Sigma_3 = \sigma_3^2$ and $\Sigma_2 = \sigma_2^2$ are both scalars.

```
. mixed gsp private emp hwy water other unemp || region: || state:
(output omitted)
```

Mixed-effects ML regression Number of obs = 816

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
region	9	51	90.7	136
state	48	17	17.0	17

Log likelihood = 1430.5017	Wald chi2(6) = 18829.06
	Prob > chi2 = 0.0000

gsp	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
private	.2671484	.0212591	12.57	0.000	.2254814	.3088154
emp	.754072	.0261868	28.80	0.000	.7027468	.8053973
hwy	.0709767	.023041	3.08	0.002	.0258172	.1161363
water	.0761187	.0139248	5.47	0.000	.0488266	.1034109
other	-.0999955	.0169366	-5.90	0.000	-.1331906	-.0668004
unemp	-.0058983	.0009031	-6.53	0.000	-.0076684	-.0041282
_cons	2.128823	.1543854	13.79	0.000	1.826233	2.431413

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
region: Identity				
var(_cons)	.0014506	.0012995	.0002506	.0083957
state: Identity				
var(_cons)	.0062757	.0014871	.0039442	.0099855
var(Residual)	.0013461	.0000689	.0012176	.0014882

LR test vs. linear model: chi2(2) = 1154.73 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

Notes:

1. Our model now has two random-effects equations, separated by ||. The first is a random intercept (constant only) at the `region` level (level three), and the second is a random intercept at the `state` level (level two). The order in which these are specified (from left to right) is significant—`mixed` assumes that `state` is nested within `region`.
2. The information on groups is now displayed as a table, with one row for each grouping. You can suppress this table with the `nogroup` or the `noheader` option, which will suppress the rest of the header, as well.
3. The variance-component estimates are now organized and labeled according to level.

After adjusting for the nested-level error structure, we find that the highway and water components of public capital had significant positive effects on private output, whereas the other public buildings component had a negative effect.



❏ Technical note

In the previous example, the states are coded 1–48 and are nested within nine regions. `mixed` treated the states as nested within regions, regardless of whether the codes for each state were unique between regions. That is, even if codes for states were duplicated between regions, `mixed` would have enforced the nesting and produced the same results.

The group information at the top of the `mixed` output and that produced by the postestimation command `estat group` (see [ME] [estat group](#)) take the nesting into account. The statistics are thus not necessarily what you would get if you instead `tabulated` each group variable individually.



Model (7) extends in a straightforward manner to more than three levels, as does the specification of such models in `mixed`.

Blocked-diagonal covariance structures

Covariance matrices of random effects within an equation can be modeled either as a multiple of the identity matrix, as diagonal (that is, `Independent`), as exchangeable, or as general symmetric (`Unstructured`). These may also be combined to produce more complex block-diagonal covariance structures, effectively placing constraints on the variance components.

► Example 5: Using repeated levels to induce blocked-diagonal covariance structures

Returning to our productivity data, we now add random coefficients on `hwy` and `unemp` at the `region` level. This only slightly changes the estimates of the fixed effects, so we focus our attention on the variance components:

```
. mixed gsp private emp hwy water other unemp || region: hwy unemp || state:,
> nolog nogroup nofetable
Mixed-effects ML regression                Number of obs    =          816
                                           Wald chi2(6)       =    17137.94
Log likelihood = 1447.6787                Prob > chi2        =     0.0000
```

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
region: Independent				
var(hwy)	.0000209	.0001103	6.71e-10	.6506957
var(unemp)	.0000238	.0000135	7.84e-06	.0000722
var(_cons)	.0030349	.0086684	.0000112	.8191291
state: Identity				
var(_cons)	.0063658	.0015611	.0039365	.0102943
var(Residual)	.0012469	.0000643	.001127	.0013795

LR test vs. linear model: chi2(4) = 1189.08 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

```
. estimates store prodc
```

This model is the same as that fit in [example 4](#) except that $\mathbf{Z}_{jk}^{(3)}$ is now the $n_{jk} \times 3$ matrix with columns determined by the values of `hwy`, `unemp`, and an intercept term (one), in that order, and (because we used the default `Independent` structure) Σ_3 is

$$\Sigma_3 = \begin{pmatrix} \text{hwy} & \text{unemp} & \text{\_cons} \\ \sigma_a^2 & 0 & 0 \\ 0 & \sigma_b^2 & 0 \\ 0 & 0 & \sigma_c^2 \end{pmatrix}$$

The random-effects specification at the state level remains unchanged; that is, Σ_2 is still treated as the scalar variance of the random intercepts at the state level.

An LR test comparing this model with that from [example 4](#) favors the inclusion of the two random coefficients, a fact we leave to the interested reader to verify.

The estimated variance components, upon examination, reveal that the variances of the random coefficients on `hwy` and `unemp` could be treated as equal. That is,

$$\Sigma_3 = \begin{pmatrix} \text{hwy} & \text{unemp} & \text{\_cons} \\ \sigma_a^2 & 0 & 0 \\ 0 & \sigma_a^2 & 0 \\ 0 & 0 & \sigma_c^2 \end{pmatrix}$$

looks plausible. We can impose this equality constraint by treating Σ_3 as block diagonal: the first block is a 2×2 multiple of the identity matrix, that is, $\sigma_a^2 \mathbf{I}_2$; the second is a scalar, equivalently, a 1×1 multiple of the identity.

We construct block-diagonal covariances by repeating level specifications:

```
. mixed gsp private emp hwy water other unemp || region: hwy unemp,
> cov(identity) || region: || state:, nolog nogroup nofetable
Mixed-effects ML regression      Number of obs      =      816
                                Wald chi2(6)         =    17136.65
Log likelihood = 1447.6784        Prob > chi2      =     0.0000
```

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
region: Identity var(hwy unemp)	.0000238	.0000134	7.89e-06	.0000719
region: Identity var(_cons)	.0028191	.0030429	.0003399	.023383
state: Identity var(_cons)	.006358	.0015309	.0039661	.0101925
var(Residual)	.0012469	.0000643	.001127	.0013795

LR test vs. linear model: chi2(3) = 1189.08 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

We specified two equations for the `region` level: the first for the random coefficients on `hwy` and `unemp` with covariance set to `Identity` and the second for the random intercept `_cons`, whose covariance defaults to `Identity` because it is of dimension 1. `mixed` labeled the estimate of σ_a^2 as `var(hwy unemp)` to designate that it is common to the random coefficients on both `hwy` and `unemp`.

An LR test shows that the constrained model fits equally well.

```
. lrtest . prodr
Likelihood-ratio test      LR chi2(1) =      0.00
(Assumption: . nested in prodr)  Prob > chi2 =    0.9784
Note: The reported degrees of freedom assumes the null hypothesis is not on
the boundary of the parameter space. If this is not true, then the
reported test is conservative.
```



Because the null hypothesis for this test is one of equality ($H_0 : \sigma_a^2 = \sigma_b^2$), it is not on the boundary of the parameter space. As such, we can take the reported significance as precise rather than a conservative estimate.

You can repeat level specifications as often as you like, defining successive blocks of a block-diagonal covariance matrix. However, repeated-level equations must be listed consecutively; otherwise, `mixed` will give an error.

❑ Technical note

In the previous estimation output, there was no constant term included in the first `region` equation, even though we did not use the `noconstant` option. When you specify repeated-level equations, `mixed` knows not to put constant terms in each equation because such a model would be unidentified. By default, it places the constant in the last repeated-level equation, but you can use `noconstant` creatively to override this.



Linear mixed-effects models can also be fit using `meglm` with the default `gaussian` family. `meglm` provides two more covariance structures through which you can impose constraints on variance components; see [ME] [meglm](#) for details.

Heteroskedastic random effects

Blocked-diagonal covariance structures and repeated-level specifications of random effects can also be used to model heteroskedasticity among random effects at a given level.

► Example 6: Using repeated levels to model heteroskedasticity

Following [Rabe-Hesketh and Skrondal \(2012, sec. 7.2\)](#), we analyze data from Asian children in a British community who were weighed up to four times, roughly between the ages of 6 weeks and 27 months. The dataset is a random sample of data previously analyzed by [Goldstein \(1986\)](#) and [Prosser, Rasbash, and Goldstein \(1991\)](#).

```
. use http://www.stata-press.com/data/r15/childweight
(Weight data on Asian children)


. describe

Contains data from http://www.stata-press.com/data/r15/childweight.dta
   obs:                198                Weight data on Asian children
  vars:                  5                23 May 2016 15:12
 size:                 3,168              (_dta has notes)
```

variable name	storage type	display format	value label	variable label
id	int	%8.0g		child identifier
age	float	%8.0g		age in years
weight	float	%8.0g		weight in Kg
brthwt	int	%8.0g		Birth weight in g
girl	float	%9.0g	bg	gender

Sorted by: id age

```
. graph twoway (line weight age, connect(ascending)), by(girl)
> xtitle(Age in years) ytitle(Weight in kg)
```



Ignoring gender effects for the moment, we begin with the following model for the i th measurement on the j th child:

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{age}_{ij} + \beta_2 \text{age}_{ij}^2 + u_{j0} + u_{j1} \text{age}_{ij} + \epsilon_{ij}$$


```
. mixed weight i.girl i.girl#c.age c.age#c.age || id: age, nolog
```

Mixed-effects ML regression	Number of obs	=	198
Group variable: id	Number of groups	=	68
	Obs per group:		
	min	=	1
	avg	=	2.9
	max	=	5
Log likelihood =	Wald chi2(4)	=	1942.30
	Prob > chi2	=	0.0000

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
girl						
girl	-.5104676	.2145529	-2.38	0.017	-.9309835	-.0899516
girl#c.age						
boy	7.806765	.2524583	30.92	0.000	7.311956	8.301574
girl	7.577296	.2531318	29.93	0.000	7.081166	8.073425
c.age#c.age	-1.654323	.0871752	-18.98	0.000	-1.825183	-1.483463
_cons	3.754275	.1726404	21.75	0.000	3.415906	4.092644

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Independent					
	var(age)	.2772846	.0769233	.1609861	.4775987
	var(_cons)	.4076892	.12386	.2247635	.7394906
	var(Residual)	.3131704	.047684	.2323672	.422072

LR test vs. linear model: chi2(2) = 104.39 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

. estimates store homoskedastic

The main gender effect is significant at the 5% level, but the gender–age interaction is not:

```
. test 0.girl#c.age = 1.girl#c.age
```

(1)	[weight]0b.girl#c.age - [weight]1.girl#c.age = 0
	chi2(1) = 1.66
	Prob > chi2 = 0.1978

On average, boys are heavier than girls, but their average linear growth rates are not significantly different.

In the above model, we introduced a gender effect on average growth, but we still assumed that the variability in child-specific deviations from this average was the same for boys and girls. To check this assumption, we introduce gender into the random component of the model. Because support for factor-variable notation is limited in specifications of random effects (see [Crossed-effects models](#) below), we need to generate the interactions ourselves.

```
. generate boy = !girl
. generate boyXage = boy*age
. generate girlXage = girl*age
. mixed weight i.girl i.girl#c.age c.age#c.age || id: boy boyXage, noconstant
> || id: girl girlXage, noconstant nolog nofetable

Mixed-effects ML regression      Number of obs      =      198
Group variable: id              Number of groups   =      68

                                Obs per group:
                                    min =          1
                                    avg  =         2.9
                                    max  =          5

                                Wald chi2(4)          =    2358.11
                                Prob > chi2           =      0.0000

Log likelihood = -248.94752
```

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: Independent				
var(boy)	.3161091	.1557911	.1203181	.8305061
var(boyXage)	.4734482	.1574626	.2467028	.9085962
id: Independent				
var(girl)	.5798676	.1959725	.2989896	1.124609
var(girlXage)	.0664634	.0553274	.0130017	.3397538
var(Residual)	.3078826	.046484	.2290188	.4139037

```
LR test vs. linear model: chi2(4) = 112.86          Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.
. estimates store heteroskedastic
```

In the above, we suppress displaying the fixed portion of the model (the `nofetable` option) because it does not differ much from that of the previous model.

Our previous model had the random-effects specification

```
|| id: age
```

which we have replaced with the dual repeated-level specification

```
|| id: boy boyXage, noconstant || id: girl girlXage, noconstant
```

The former models a random intercept and random slope on age, and does so treating all children as a random sample from one population. The latter also specifies a random intercept and random slope on age, but allows for the variability of the random intercepts and slopes to differ between boys and girls. In other words, it allows for heteroskedasticity in random effects due to gender. We use the `noconstant` option so that we can separate the overall random intercept (automatically provided by the former syntax) into one specific to boys and one specific to girls.

There seems to be a large gender effect in the variability of linear growth rates. We can compare both models with an LR test, recalling that we stored the previous estimation results under the name `homoskedastic`:

```
. lrtest homoskedastic heteroskedastic

Likelihood-ratio test          LR chi2(2) =      8.47
(Assumption: homoskedastic nested in heteroskedastic) Prob > chi2 =    0.0145

Note: The reported degrees of freedom assumes the null hypothesis is not on
the boundary of the parameter space. If this is not true, then the
reported test is conservative.
```

Because the null hypothesis here is one of equality of variances and not that variances are 0, the above does not test on the boundary; thus we can treat the significance level as precise and not conservative. Either way, the results favor the new model with heteroskedastic random effects.

Heteroskedastic residual errors

Up to this point, we have assumed that the level-one residual errors—the ϵ 's in the stated models—have been i.i.d. Gaussian with variance σ_ϵ^2 . This is demonstrated in `mixed` output in the random-effects table, where up until now we have estimated a single residual-error variance, labeled as `var(Residual)`.

To relax the assumptions of homoskedasticity or independence of residual errors, use the `residuals()` option.

► Example 7: Independent residual variance structure

West, Welch, and Galecki (2015, chap. 7) analyze data studying the effect of ceramic dental veneer placement on gingival (gum) health. Data on 55 teeth located in the maxillary arches of 12 patients were considered.

```
. use http://www.stata-press.com/data/r15/veneer, clear
(Dental veneer data)

. describe
Contains data from http://www.stata-press.com/data/r15/veneer.dta
obs:                110          Dental veneer data
vars:                7           24 May 2016 12:11
size:               1,100        (_dta has notes)
```

variable name	storage type	display format	value label	variable label
patient	byte	%8.0g		Patient ID
tooth	byte	%8.0g		Tooth number with patient
gcf	byte	%8.0g		Gingival crevicular fluid (GCF)
age	byte	%8.0g		Patient age
base_gcf	byte	%8.0g		Baseline GCF
cda	float	%9.0g		Average contour difference after veneer placement
followup	byte	%9.0g	t	Follow-up time: 3 or 6 months

Sorted by:

Veneers were placed to match the original contour of the tooth as closely as possible, and researchers were interested in how contour differences (variable `cda`) impacted gingival health. Gingival health was measured as the amount of gingival crevicular fluid (GCF) at each tooth, measured at baseline (variable `base_gcf`) and at two posttreatment follow-ups at 3 and 6 months. The variable `gcf` records GCF at follow-up, and the variable `followup` records the follow-up time.

Because two measurements were taken for each tooth and there exist multiple teeth per patient, we fit a three-level model with the following random effects: a random intercept and random slope on follow-up time at the patient level, and a random intercept at the tooth level. For the i th measurement of the j th tooth from the k th patient, we have

$$\text{gcf}_{ijk} = \beta_0 + \beta_1 \text{followup}_{ijk} + \beta_2 \text{base\_gcf}_{ijk} + \beta_3 \text{cda}_{ijk} + \beta_4 \text{age}_{ijk} + u_{0k} + u_{1k} \text{followup}_{ijk} + v_{0jk} + \epsilon_{ijk}$$

which we can fit using `mixed`:

```
. mixed gcf followup base_gcf cda age || patient: followup, cov(un) || tooth:,  
> reml nolog
```

Mixed-effects REML regression Number of obs = 110

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
patient	12	2	9.2	12
tooth	55	2	2.0	2

	Wald chi2(4)	=	7.48
Log restricted-likelihood = -420.92761	Prob > chi2	=	0.1128

gcf	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
followup	.3009815	1.936863	0.16	0.877	-3.4952	4.097163
base_gcf	-.0183127	.1433094	-0.13	0.898	-.299194	.2625685
cda	-.329303	.5292525	-0.62	0.534	-1.366619	.7080128
age	-.5773932	.2139656	-2.70	0.007	-.9967582	-.1580283
_cons	45.73862	12.55497	3.64	0.000	21.13133	70.34591

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
patient: Unstructured				
var(followup)	41.88772	18.79997	17.38009	100.9535
var(_cons)	524.9851	253.0205	204.1287	1350.175
cov(followup,_cons)	-140.4229	66.57623	-270.9099	-9.935907
tooth: Identity				
var(_cons)	47.45738	16.63034	23.8792	94.3165
var(Residual)	48.86704	10.50523	32.06479	74.47382

LR test vs. linear model: $\chi^2(4) = 91.12$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

We used REML estimation for no other reason than variety.

Among the other features of the model fit, we note that the residual variance σ_{ϵ}^2 was estimated as 48.87 and that our model assumed that the residuals were independent with constant variance (homoskedastic). Because it may be the case that the precision of gcf measurements could change over time, we modify the above to estimate two distinct error variances: one for the 3-month follow-up and one for the 6-month follow-up.

To fit this model, we add the `residuals(independent, by(followup))` option, which maintains independence of residual errors but allows for heteroskedasticity with respect to follow-up time.

The **exchangeable** structure assumes one overall variance and one common pairwise covariance; that is,

$$\text{Var}(\epsilon_j) = \text{Var} \begin{bmatrix} \epsilon_{j1} \\ \epsilon_{j2} \\ \vdots \\ \epsilon_{jn_j} \end{bmatrix} = \begin{bmatrix} \sigma_\epsilon^2 & \sigma_1 & \cdots & \sigma_1 \\ \sigma_1 & \sigma_\epsilon^2 & \cdots & \sigma_1 \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_1 & \sigma_1 & \sigma_1 & \sigma_\epsilon^2 \end{bmatrix}$$

By default, **mixed** will report estimates of the two parameters as estimates of the common variance σ_ϵ^2 and of the covariance σ_1 . When the **by(varname)** option is also specified, these two parameters are estimated for each level *varname*.

The **ar p** structure assumes that the errors have an AR structure of order *p*. That is,

$$\epsilon_{ij} = \phi_1 \epsilon_{i-1,j} + \cdots + \phi_p \epsilon_{i-p,j} + u_{ij}$$

where u_{ij} are i.i.d. Gaussian with mean 0 and variance σ_u^2 . **mixed** reports estimates of $\phi_1, \dots, \phi_p$ and the overall error variance σ_ϵ^2 , which can be derived from the above expression. The **t(varname)** option is required, where *varname* is a time variable used to order the observations within lowest-level groups and to determine any gaps between observations. When the **by(varname)** option is also specified, the set of $p + 1$ parameters is estimated for each level of *varname*. If $p = 1$, then the estimate of ϕ_1 is reported as **rho**, because in this case it represents the correlation between successive error terms.

The **ma q** structure assumes that the errors are an MA process of order *q*. That is,

$$\epsilon_{ij} = u_{ij} + \theta_1 u_{i-1,j} + \cdots + \theta_q u_{i-q,j}$$

where u_{ij} are i.i.d. Gaussian with mean 0 and variance σ_u^2 . **mixed** reports estimates of $\theta_1, \dots, \theta_q$ and the overall error variance σ_ϵ^2 , which can be derived from the above expression. The **t(varname)** option is required, where *varname* is a time variable used to order the observations within lowest-level groups and to determine any gaps between observations. When the **by(varname)** option is also specified, the set of $q + 1$ parameters is estimated for each level of *varname*.

The **unstructured** structure is the most general and estimates unique variances and unique pairwise covariances for all residuals within the lowest-level grouping. Because the data may be unbalanced and the ordering of the observations is arbitrary, the **t(varname)** option is required, where *varname* is an identification variable that matches error terms in different groups. If *varname* has *n* distinct levels, then $n(n + 1)/2$ parameters are estimated. Not all *n* levels need to be observed within each group, but duplicated levels of *varname* within a given group are not allowed because they would cause a singularity in the estimated error-variance matrix for that group. When the **by(varname)** option is also specified, the set of $n(n + 1)/2$ parameters is estimated for each level of *varname*.

The **banded q** structure is a special case of **unstructured** that confines estimation to within the first *q* off-diagonal elements of the residual variance–covariance matrix and sets the covariances outside this band to 0. As is the case with **unstructured**, the **t(varname)** option is required, where *varname* is an identification variable that matches error terms in different groups. However, with **banded** variance structures, the ordering of the values in *varname* is significant because it determines which covariances are to be estimated and which are to be set to 0. For example, if *varname* has $n = 5$ distinct values $t = 1, 2, 3, 4, 5$, then a banded variance–covariance structure of order $q = 2$ would estimate the following:

$$\text{Var}(\epsilon_j) = \text{Var} \begin{bmatrix} \epsilon_{1j} \\ \epsilon_{2j} \\ \epsilon_{3j} \\ \epsilon_{4j} \\ \epsilon_{5j} \end{bmatrix} = \begin{bmatrix} \sigma_1^2 & \sigma_{12} & \sigma_{13} & 0 & 0 \\ \sigma_{12} & \sigma_2^2 & \sigma_{23} & \sigma_{24} & 0 \\ \sigma_{13} & \sigma_{23} & \sigma_3^2 & \sigma_{34} & \sigma_{35} \\ 0 & \sigma_{24} & \sigma_{34} & \sigma_4^2 & \sigma_{45} \\ 0 & 0 & \sigma_{35} & \sigma_{45} & \sigma_5^2 \end{bmatrix}$$

In other words, you would have an unstructured variance matrix that constrains $\sigma_{14} = \sigma_{15} = \sigma_{25} = 0$. If *varname* has n distinct levels, then $(q+1)(2n-q)/2$ parameters are estimated. Not all n levels need to be observed within each group, but duplicated levels of *varname* within a given group are not allowed because they would cause a singularity in the estimated error-variance matrix for that group. When the `by(varname)` option is also specified, the set of parameters is estimated for each level of *varname*. If q is left unspecified, then `banded` is equivalent to `unstructured`; that is, all variances and covariances are estimated. When $q = 0$, $\text{Var}(\epsilon_j)$ is treated as diagonal and can thus be used to model uncorrelated yet heteroskedastic residual errors.

The `toeplitz q` structure assumes that the residual errors are homoskedastic and that the correlation between two errors is determined by the time lag between the two. That is, $\text{Var}(\epsilon_{ij}) = \sigma_\epsilon^2$ and

$$\text{Corr}(\epsilon_{ij}, \epsilon_{i+k,j}) = \rho_k$$

If the lag k is less than or equal to q , then the pairwise correlation ρ_k is estimated; if the lag is greater than q , then ρ_k is assumed to be 0. If q is left unspecified, then ρ_k is estimated for each observed lag k . The `t(varname)` option is required, where *varname* is a time variable t used to determine the lags between pairs of residual errors. As such, `t()` must be integer-valued. $q+1$ parameters are estimated: one overall variance σ_ϵ^2 and q correlations. When the `by(varname)` option is also specified, the set of $q+1$ parameters is estimated for each level of *varname*.

The `exponential` structure is a generalization of the AR structure that allows for noninteger and irregularly spaced time lags. That is, $\text{Var}(\epsilon_{ij}) = \sigma_\epsilon^2$ and

$$\text{Corr}(\epsilon_{ij}, \epsilon_{kj}) = \rho^{|i-k|}$$

for $0 \leq \rho \leq 1$, with i and k not required to be integers. The `t(varname)` option is required, where *varname* is a time variable used to determine i and k for each residual-error pair. `t()` is real-valued. `mixed` reports estimates of σ_ϵ^2 and ρ . When the `by(varname)` option is also specified, these two parameters are estimated for each level of *varname*.

► Example 8: Autoregressive residual variance structure

Pinheiro and Bates (2000, chap. 5) analyze data from a study of the estrus cycles of mares. Originally analyzed in Pierson and Ginther (1987), the data record the number of ovarian follicles larger than 10mm, daily over a period ranging from three days before ovulation to three days after the subsequent ovulation.


```
. use http://www.stata-press.com/data/r15/ovary
(Ovarian follicles in mares)

. describe

Contains data from http://www.stata-press.com/data/r15/ovary.dta
  obs:          308              Ovarian follicles in mares
  vars:           6              20 May 2016 13:49
  size:         5,544            (_dta has notes)
```

variable name	storage type	display format	value label	variable label
mare	byte	%9.0g		mare ID
stime	float	%9.0g		Scaled time
follicles	byte	%9.0g		Number of ovarian follicles > 10 mm in diameter
sin1	float	%9.0g		sine(2*pi*stime)
cos1	float	%9.0g		cosine(2*pi*stime)
time	float	%9.0g		time order within mare

Sorted by: mare stime

The `stime` variable is time that has been scaled so that ovulation occurs at scaled times 0 and 1, and the `time` variable records the time ordering within mares. Because graphical evidence suggests a periodic behavior, the analysis includes the `sin1` and `cos1` variables, which are sine and cosine transformations of scaled time, respectively.

We consider the following model for the i th measurement on the j th mare:

$$\text{follicles}_{ij} = \beta_0 + \beta_1 \text{sin1}_{ij} + \beta_2 \text{cos1}_{ij} + u_j + \epsilon_{ij}$$

The above model incorporates the cyclical nature of the data as affecting the overall average number of follicles and includes mare-specific random effects u_j . Because we believe successive measurements within each mare are probably correlated (even after controlling for the periodicity in the average), we also model the within-mare errors as being AR of order 2.

```
. mixed follicles sin1 cos1 || mare:, residuals(ar 2, t(time)) reml nolog

Mixed-effects REML regression              Number of obs   =          308
Group variable: mare                      Number of groups  =           11
                                         Obs per group:
                                             min =           25
                                             avg =          28.0
                                             max =           31

                                         Wald chi2(2)      =          34.72
                                         Prob > chi2       =          0.0000

Log restricted-likelihood = -772.59855
```

follicles	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
sin1	-2.899227	.5110786	-5.67	0.000	-3.900923	-1.897532
cos1	-.8652936	.5432925	-1.59	0.111	-1.930127	.1995402
_cons	12.14455	.9473712	12.82	0.000	10.28774	14.00136

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
mare: Identity	var(_cons)	7.092607	4.402031	2.101392	23.93892
Residual: AR(2)					
	phi1	.5386104	.0624897	.4161329	.661088
	phi2	.1446712	.0632039	.0207939	.2685486
	var(e)	14.25104	2.435233	10.19512	19.92052

LR test vs. linear model: $\chi^2(3) = 251.67$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

We picked an order of 2 as a guess, but we could have used LR tests of competing AR models to determine the optimal order, because models of smaller order are nested within those of larger order.

◀

► Example 9: Unstructured residual variance structure

Fitzmaurice, Laird, and Ware (2011, chap. 7) analyzed data on 37 subjects who participated in an exercise therapy trial.

```
. use http://www.stata-press.com/data/r15/exercise
(Exercise Therapy Trial)
. describe
Contains data from http://www.stata-press.com/data/r15/exercise.dta
   obs:      259               Exercise Therapy Trial
   vars:      4               24 Jun 2016 18:35
   size:    1,036             (_dta has notes)
```

variable name	storage type	display format	value label	variable label
id	byte	%9.0g		Person ID
day	byte	%9.0g		Day of measurement
program	byte	%9.0g		1 = reps increase; 2 = weights increase
strength	byte	%9.0g		Strength measurement

Sorted by: id day

Subjects (variable `id`) were placed on either an increased-repetition regimen (`program==1`) or a program that kept the repetitions constant but increased weight (`program==2`). Muscle-strength measurements (variable `strength`) were taken at baseline (`day==0`) and then every two days over the next twelve days.

Following Fitzmaurice, Laird, and Ware (2011, chap. 7), and to demonstrate fitting residual-error structures to data collected at uneven time points, we confine our analysis to those data collected at baseline and at days 4, 6, 8, and 12. We fit a full two-way factorial model of `strength` on `program` and `day`, with an unstructured residual-error covariance matrix over those repeated measurements taken on the same subject:

```
. keep if inlist(day, 0, 4, 6, 8, 12)
(74 observations deleted)

. mixed strength i.program##i.day || id:,
> noconstant residuals(unstructured, t(day)) nolog

Mixed-effects ML regression      Number of obs      =      173
Group variable: id              Number of groups   =      37
                                Obs per group:
                                min =      3
                                avg =      4.7
                                max =      5
                                Wald chi2(9)          =      45.85
                                Prob > chi2           =      0.0000

Log likelihood = -296.58215
```

strength	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
2.program	1.360119	1.003549	1.36	0.175	-.6068016	3.32704
day						
4	1.125	.3322583	3.39	0.001	.4737858	1.776214
6	1.360127	.3766894	3.61	0.000	.6218298	2.098425
8	1.583563	.4905876	3.23	0.001	.6220287	2.545097
12	1.623576	.5372947	3.02	0.003	.5704977	2.676654
program#day						
2 4	-.169034	.4423472	-0.38	0.702	-1.036019	.6979506
2 6	.2113012	.4982385	0.42	0.671	-.7652283	1.187831
2 8	-.1299763	.6524813	-0.20	0.842	-1.408816	1.148864
2 12	.3212829	.7306782	0.44	0.660	-1.11082	1.753386
_cons	79.6875	.7560448	105.40	0.000	78.20568	81.16932

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: (empty)				
Residual: Unstructured				
var(e0)	9.14566	2.126233	5.798599	14.42471
var(e4)	11.87114	2.761187	7.524987	18.72747
var(e6)	10.06571	2.348835	6.371125	15.90275
var(e8)	13.22464	3.113885	8.336026	20.98014
var(e12)	13.16909	3.167316	8.219245	21.09985
cov(e0,e4)	9.625235	2.331946	5.054705	14.19577
cov(e0,e6)	8.489042	2.106352	4.360668	12.61742
cov(e0,e8)	9.280413	2.369524	4.636232	13.92459
cov(e0,e12)	8.898006	2.348212	4.295594	13.50042
cov(e4,e6)	10.49184	2.492498	5.606639	15.37705
cov(e4,e8)	11.89787	2.848714	6.314492	17.48125
cov(e4,e12)	11.28344	2.804991	5.78576	16.78112
cov(e6,e8)	11.0507	2.646955	5.862762	16.23863
cov(e6,e12)	10.5006	2.590246	5.423812	15.57739
cov(e8,e12)	12.4091	3.010761	6.508121	18.31009

LR test vs. linear model: chi2(14) = 314.67 Prob > chi2 = 0.0000

Note: The reported degrees of freedom assumes the null hypothesis is not on the boundary of the parameter space. If this is not true, then the reported test is conservative.

Because we are using the variable id only to group the repeated measurements and not to introduce random effects at the subject level, we use the noconstant option to omit any subject-level effects.

The unstructured covariance matrix is the most general and contains many parameters. In this example, we estimate a distinct residual variance for each day and a distinct covariance for each pair of days.

That there is positive covariance between all pairs of measurements is evident, but what is not as evident is whether the covariances may be more parsimoniously represented. One option would be to explore whether the correlation diminishes as the time gap between strength measurements increases and whether it diminishes systematically. Given the irregularity of the time intervals, an exponential structure would be more appropriate than, say, an AR or MA structure.

```
. estimates store unstructured
. mixed strength i.program##i.day || id:, noconstant
> residuals(exponential, t(day)) nolog nofetable
```

Mixed-effects ML regression	Number of obs	=	173
Group variable: id	Number of groups	=	37
	Obs per group:		
	min	=	3
	avg	=	4.7
	max	=	5
	Wald chi2(9)	=	36.77
	Prob > chi2	=	0.0000

Log likelihood = -307.83324

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: (empty)				
Residual: Exponential				
rho	.9786462	.0051238	.9659207	.9866854
var(e)	11.22349	2.338371	7.460762	16.88389

LR test vs. linear model: chi2(1) = 292.17 Prob > chi2 = 0.0000

Note: The reported degrees of freedom assumes the null hypothesis is not on the boundary of the parameter space. If this is not true, then the reported test is conservative.

In the above example, we suppressed displaying the main regression parameters because they did not differ much from those of the previous model. While the unstructured model estimated 15 variance–covariance parameters, the exponential model claims to get the job done with just 2, a fact that is not disputed by an LR test comparing the two nested models (at least not at the 0.01 level).

```
. lrtest unstructured .
```

Likelihood-ratio test	LR chi2(13)	=	22.50
(Assumption: . nested in unstructured)	Prob > chi2	=	0.0481

Note: The reported degrees of freedom assumes the null hypothesis is not on the boundary of the parameter space. If this is not true, then the reported test is conservative.

Crossed-effects models

Not all mixed models contain nested levels of random effects.

► Example 10: Crossed-effects model

Returning to our longitudinal analysis of pig weights, suppose that instead of (5) we wish to fit

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_i + v_j + \epsilon_{ij} \quad (8)$$

for the $i = 1, \dots, 9$ weeks and $j = 1, \dots, 48$ pigs and

$$u_i \sim N(0, \sigma_u^2); \quad v_j \sim N(0, \sigma_v^2); \quad \epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$$

all independently. Both (5) and (8) assume an overall population-average growth curve $\beta_0 + \beta_1 \text{week}$ and a random pig-specific shift.

The models differ in how `week` enters into the random part of the model. In (5), we assume that the effect due to `week` is linear and pig specific (a random slope); in (8), we assume that the effect due to `week`, u_i , is systematic to that week and common to all pigs. The rationale behind (8) could be that, assuming that the pigs were measured contemporaneously, we might be concerned that week-specific random factors such as weather and feeding patterns had significant systematic effects on all pigs.

Model (8) is an example of a two-way crossed-effects model, with the pig effects v_j being crossed with the week effects u_i . One way to fit such models is to consider all the data as one big cluster, and treat the u_i and v_j as a series of $9 + 48 = 57$ random coefficients on indicator variables for `week` and `pig`. In the notation of (2),

$$\mathbf{u} = \begin{bmatrix} u_1 \\ \vdots \\ u_9 \\ v_1 \\ \vdots \\ v_{48} \end{bmatrix} \sim N(\mathbf{0}, \mathbf{G}); \quad \mathbf{G} = \begin{bmatrix} \sigma_u^2 \mathbf{I}_9 & \mathbf{0} \\ \mathbf{0} & \sigma_v^2 \mathbf{I}_{48} \end{bmatrix}$$

Because $\mathbf{G}$ is block diagonal, it can be represented in `mixed` as repeated-level equations. All we need is an identification variable to identify all the observations as one big group and a way to tell `mixed` to treat `week` and `pig` as factor variables (or equivalently, as two sets of overparameterized indicator variables identifying weeks and pigs, respectively). `mixed` supports the special group designation `_all` for the former and the `R.varname` notation for the latter.

```

. use http://www.stata-press.com/data/r15/pig, clear
(Longitudinal analysis of pig weights)
. mixed weight week || _all: R.week || _all: R.id
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:  log likelihood = -1013.824
Iteration 1:  log likelihood = -1013.824
Computing standard errors:
Mixed-effects ML regression              Number of obs      =       432
Group variable: _all                    Number of groups   =        1
                                         Obs per group:
                                         min =       432
                                         avg =     432.0
                                         max =       432
                                         Wald chi2(1)      =    13258.28
                                         Prob > chi2       =     0.0000
Log likelihood = -1013.824

```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0539313	115.14	0.000	6.104192	6.315599
_cons	19.35561	.6333982	30.56	0.000	18.11418	20.59705

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
_all: Identity	var(R.week)	.0849874	.0868856	.0114588	.6303302
_all: Identity	var(R.id)	14.83623	3.126142	9.816733	22.42231
	var(Residual)	4.297328	.3134404	3.724888	4.957741

LR test vs. linear model: chi2(2) = 474.85 Prob > chi2 = 0.0000

Note: LR test is conservative and provided only for reference.

. estimates store crossed

Thus we estimate $\hat{\sigma}_u^2 = 0.08$ and $\hat{\sigma}_v^2 = 14.84$. Both (5) and (8) estimate a total of five parameters: two fixed effects and three variance components. The models, however, are not nested within each other, which precludes the use of an LR test to compare both models. Refitting model (5) and looking at the Akaike information criteria values by using `estimates stats`,

```

. quietly mixed weight week || id:week
. estimates stats crossed

```

Akaike's information criterion and Bayesian information criterion

Model	Obs	ll(null)	ll(model)	df	AIC	BIC
crossed	432	.	-1013.824	5	2037.648	2057.99
.	432	.	-869.0383	5	1748.077	1768.419

Note: N=Obs used in calculating BIC; see [\[R\] BIC note](#).

definitely favors model (5). This finding is not surprising given that our rationale behind (8) was somewhat fictitious. In our `estimates stats` output, the values of `ll(null)` are missing. `mixed` does not fit a constant-only model as part of its usual estimation of the full model, but you can use `mixed` to fit a constant-only model directly, if you wish.

The `R.varname` notation is equivalent to giving a list of overparameterized (none dropped) indicator variables for use in a random-effects specification. When you specify `R.varname`, `mixed` handles the calculations internally rather than creating the indicators in the data. Because the set of indicators is overparameterized, `R.varname` implies `noconstant`. You can include factor variables in the fixed-effects specification by using standard methods; see [U] 11.4.3 **Factor variables**. However, random-effects equations support only the `R.varname` factor specification. For more complex factor specifications (such as interactions) in random-effects equations, use `generate` to form the variables manually, as we demonstrated in [example 6](#).

□ Technical note

Although we were able to fit the crossed-effects model (8), it came at the expense of increasing the column dimension of our random-effects design from 2 in model (5) to 57 in model (8). Computation time and memory requirements grow (roughly) quadratically with the dimension of the random effects. As a result, fitting such crossed-effects models is feasible only when the total column dimension is small to moderate.

Reexamining model (8), we note that if we drop u_i , we end up with a model equivalent to (4), meaning that we could have fit (4) by typing

```
. mixed weight week || _all: R.id
```

instead of

```
. mixed weight week || id:
```

as we did when we originally fit the model. The results of both estimations are identical, but the latter specification, organized at the cluster (pig) level with random-effects dimension 1 (a random intercept) is much more computationally efficient. Whereas with the first form we are limited in how many pigs we can analyze, there is no such limitation with the second form.

Furthermore, we fit model (8) by using

```
. mixed weight week || _all: R.week || _all: R.id
```

as a direct way to demonstrate the `R.` notation. However, we can technically treat pigs as nested within the `_all` group, yielding the equivalent and more efficient (total column dimension 10) way to fit (8):

```
. mixed weight week || _all: R.week || id:
```

We leave it to you to verify that both produce identical results. See [Rabe-Hesketh and Skrondal \(2012\)](#) for additional techniques to make calculations more efficient in more complex models.

□

▷ Example 11: Three-level model expressed in terms of a two-level model

As another example of how the same model may be fit in different ways by using `mixed` (and as a way to demonstrate `covariance(exchangeable)`), consider the three-level model used in [example 4](#):

$$y_{jk} = \mathbf{X}_{jk}\beta + u_k^{(3)} + u_{jk}^{(2)} + \epsilon_{jk}$$

where $\mathbf{y}_{jk}$ represents the logarithms of gross state products for the $n_{jk} = 17$ observations from state j in region k , $\mathbf{X}_{jk}$ is a set of regressors, $u_k^{(3)}$ is a random intercept at the region level, and $u_{jk}^{(2)}$ is a random intercept at the state (nested within region) level. We assume that $u_k^{(3)} \sim N(0, \sigma_3^2)$ and $u_{jk}^{(2)} \sim N(0, \sigma_2^2)$ independently. Define

$$\mathbf{v}_k = \begin{bmatrix} u_k^{(3)} + u_{1k}^{(2)} \\ u_k^{(3)} + u_{2k}^{(2)} \\ \vdots \\ u_k^{(3)} + u_{M_k, k}^{(2)} \end{bmatrix}$$

where M_k is the number of states in region k . Making this substitution, we can stack the observations for all the states within region k to get

$$\mathbf{y}_k = \mathbf{X}_k \boldsymbol{\beta} + \mathbf{Z}_k \mathbf{v}_k + \boldsymbol{\epsilon}_k$$

where $\mathbf{Z}_k$ is a set of indicators identifying the states within each region; that is,

$$\mathbf{Z}_k = \mathbf{I}_{M_k} \otimes \mathbf{J}_{17}$$

for a k -column vector of 1s $\mathbf{J}_k$, and

$$\boldsymbol{\Sigma} = \text{Var}(\mathbf{v}_k) = \begin{bmatrix} \sigma_3^2 + \sigma_2^2 & \sigma_3^2 & \cdots & \sigma_3^2 \\ \sigma_3^2 & \sigma_3^2 + \sigma_2^2 & \cdots & \sigma_3^2 \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_3^2 & \sigma_3^2 & \sigma_3^2 & \sigma_3^2 + \sigma_2^2 \end{bmatrix}_{M_k \times M_k}$$

Because $\boldsymbol{\Sigma}$ is an exchangeable matrix, we can fit this alternative form of the model by specifying the **exchangeable** covariance structure.


```
. use http://www.stata-press.com/data/r15/productivity
(Public Capital Productivity)
. mixed gsp private emp hwy water other unemp || region: R.state,
> cov(exchangeable)
(output omitted)

Mixed-effects ML regression                Number of obs    =      816
Group variable: region                    Number of groups   =       9
                                           Obs per group:
                                           min =          51
                                           avg  =         90.7
                                           max  =         136

                                           Wald chi2(6)      =   18829.06
                                           Prob > chi2       =    0.0000

Log likelihood = 1430.5017
```

gsp	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
private	.2671484	.0212591	12.57	0.000	.2254813	.3088154
emp	.7540721	.0261868	28.80	0.000	.7027468	.8053973
hwy	.0709767	.023041	3.08	0.002	.0258172	.1161363
water	.0761187	.0139248	5.47	0.000	.0488266	.1034109
other	-.0999955	.0169366	-5.90	0.000	-.1331907	-.0668004
unemp	-.0058983	.0009031	-6.53	0.000	-.0076684	-.0041282
_cons	2.128823	.1543855	13.79	0.000	1.826233	2.431413

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
region: Exchangeable				
var(R.state)	.0077263	.0017926	.0049032	.0121749
cov(R.state)	.0014506	.0012995	-.0010963	.0039975
var(Residual)	.0013461	.0000689	.0012176	.0014882

LR test vs. linear model: chi2(2) = 1154.73 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

The estimates of the fixed effects and their standard errors are equivalent to those from [example 4](#), and remapping the variance components from $(\sigma_3^2 + \sigma_2^2, \sigma_3^2, \sigma_\epsilon^2)$, as displayed here, to $(\sigma_3^2, \sigma_2^2, \sigma_\epsilon^2)$, as displayed in [example 4](#), will show that they are equivalent as well.

Of course, given the discussion in the previous technical note, it is more efficient to fit this model as we did originally, as a three-level model.



Diagnosing convergence problems

Given the flexibility of mixed-effects models, you will find that some models fail to converge when used with your data; see [Diagnosing convergence problems](#) in [\[ME\] me](#) for advice applicable to mixed-effects models in general.

In unweighted LME models with independent and homoskedastic residuals, one useful way to diagnose problems of nonconvergence is to rely on the EM algorithm ([Dempster, Laird, and Rubin 1977](#)), normally used by `mixed` only as a means of refining starting values. The advantages of EM are that it does not require a Hessian calculation, each successive EM iteration will result in a larger likelihood, iterations can be calculated quickly, and iterations will quickly bring parameter estimates into a

neighborhood of the solution. The disadvantages of EM are that, once in a neighborhood of the solution, it can be slow to converge, if at all, and EM provides no facility for estimating standard errors of the estimated variance components. One useful property of EM is that it is always willing to provide a solution if you allow it to iterate enough times, if you are satisfied with being in a neighborhood of the optimum rather than right on the optimum, and if standard errors of variance components are not crucial to your analysis.

If you encounter a nonconvergent model, try using the `emonly` option to bypass gradient-based optimization. Use `emiterate(#)` to specify the maximum number of EM iterations, which you will usually want to set much higher than the default of 20. If your EM solution shows an estimated variance component that is near 0, a ridge is formed by an interval of values near 0, which produces the same likelihood and looks equally good to the optimizer. In this case, the solution is to drop the offending variance component from the model.

Survey data

Multilevel modeling of survey data is a little different from standard modeling in that weighted sampling can take place at multiple levels in the model, resulting in multiple sampling weights. Most survey datasets, regardless of the design, contain one overall inclusion weight for each observation in the data. This weight reflects the inverse of the probability of ultimate selection, and by “ultimate” we mean that it factors in all levels of clustered sampling, corrections for noninclusion and oversampling, poststratification, etc.

For simplicity, in what follows assume a simple two-stage sampling design where groups are randomly sampled and then individuals within groups are sampled. Also assume that no additional weight corrections are performed; that is, sampling weights are simply the inverse of the probability of selection. The sampling weight for observation i in cluster j in our two-level sample is then $w_{ij} = 1/\pi_{ij}$, where π_{ij} is the probability that observation i, j is selected. If you were performing a standard analysis such as OLS regression with `regress`, you would simply use a variable holding w_{ij} as your `pweight` variable, and the fact that it came from two levels of sampling would not concern you. Perhaps you would type `vce(cluster groupvar)` where `groupvar` identifies the top-level groups to get standard errors that control for correlation within these groups, but you would still use only a single weight variable.

Now take these same data and fit a two-level model with `mixed`. As seen in (14) in *Methods and formulas* later in this entry, it is not sufficient to use the single sampling weight w_{ij} , because weights enter into the log likelihood at both the group level and the individual level. Instead, what is required for a two-level model under this sampling design is w_j , the inverse of the probability that group j is selected in the first stage, and $w_{i|j}$, the inverse of the probability that individual i from group j is selected at the second stage *conditional on group j already being selected*. It simply will not do to just use w_{ij} without making any assumptions about w_j .

Given the rules of conditional probability, $w_{ij} = w_j w_{i|j}$. If your dataset has only w_{ij} , then you will need to either assume equal probability sampling at the first stage ($w_j = 1$ for all j) or find some way to recover w_j from other variables in your data; see [Rabe-Hesketh and Skrondal \(2006\)](#) and the references therein for some suggestions on how to do this, but realize that there is little yet known about how well these approximations perform in practice.

What you really need to fit your two-level model are data that contain w_j in addition to either w_{ij} or $w_{i|j}$. If you have w_{ij} —that is, the unconditional inclusion weight for observation i, j —then you need to either divide w_{ij} by w_j to obtain $w_{i|j}$ or rescale w_{ij} so that its dependence on w_j disappears. If you already have $w_{i|j}$, then rescaling becomes optional (but still an important decision to make).

Weight rescaling is not an exact science, because the scale of the level-one weights is at issue regardless of whether they represent w_{ij} or $w_{i|j}$: because w_{ij} is unique to group j , the group-to-group magnitudes of these weights need to be normalized so that they are “consistent” from group to group. This is in stark contrast to a standard analysis, where the scale of sampling weights does not factor into estimation, instead only affecting the estimate of the total population size.

`mixed` offers three methods for standardizing weights in a two-level model, and you can specify which method you want via the `pwscale()` option. If you specify `pwscale(size)`, then the $w_{i|j}$ (or w_{ij} , it does not matter) are scaled to sum to the cluster size n_j . Method `pwscale(effective)` adds in a dependence on the sum of the squared weights so that level-one weights sum to the “effective” sample size. Just like `pwscale(size)`, `pwscale(effective)` also behaves the same whether you have $w_{i|j}$ or w_{ij} , and so it can be used with either.

Although both `pwscale(size)` and `pwscale(effective)` leave w_j untouched, the `pwscale(gk)` method is a little different in that 1) it changes the weights at both levels and 2) it does assume you have $w_{i|j}$ for level-one weights and not w_{ij} (if you have the latter, then first divide by w_j). Using the method of [Graubard and Korn \(1996\)](#), it sets the weights at the group level (level two) to the cluster averages of the products of both level weights (this product being w_{ij}). It then sets the individual weights to 1 everywhere; see [Methods and formulas](#) for the computational details of all three methods.

Determining which method is “best” is a tough call and depends on cluster size (the smaller the clusters, the greater the sensitivity to scale), whether the sampling is informative (that is, the sampling weights are correlated with the residuals), whether you are interested primarily in regression coefficients or in variance components, whether you have a simple random-intercept model or a more complex random-coefficients model, and other factors; see [Rabe-Hesketh and Skrondal \(2006\)](#), [Carle \(2009\)](#), and [Pfeffermann et al. \(1998\)](#) for some detailed advice. At the very least, you want to compare estimates across all three scaling methods (four, if you add no scaling) and perform a sensitivity analysis.

If you choose to rescale level-one weights, it does not matter whether you have $w_{i|j}$ or w_{ij} . For the `pwscale(size)` and `pwscale(effective)` methods, you get identical results, and even though `pwscale(gk)` assumes $w_{i|j}$, you can obtain this as $w_{i|j} = w_{ij}/w_j$ before proceeding.

If you do not specify `pwscale()`, then no scaling takes place, and thus at a minimum, you need to make sure you have $w_{i|j}$ in your data and not w_{ij} .

► Example 12: Mixed-effect models with survey data

[Rabe-Hesketh and Skrondal \(2006\)](#) analyzed data from the 2000 Programme for International Student Assessment (PISA) study on reading proficiency among 15-year-old American students, as performed by the Organisation for Economic Co-operation and Development (OECD). The original study was a three-stage cluster sample, where geographic areas were sampled at the first stage, schools at the second, and students at the third. Our version of the data does not contain the geographic-areas variable, so we treat this as a two-stage sample where schools are sampled at the first stage and students at the second.

```
. use http://www.stata-press.com/data/r15/pisa2000
(Programme for International Student Assessment (PISA) 2000 data)

. describe
Contains data from http://www.stata-press.com/data/r15/pisa2000.dta
   obs:                2,069                Programme for International
                                                Student Assessment (PISA) 2000
                                                data
   vars:                 11                12 Jun 2016 10:08
   size:                37,242            (_dta has notes)
```

variable name	storage type	display format	value label	variable label
female	byte	%8.0g		1 if female
isei	byte	%8.0g		International socio-economic index
w_fstuw	float	%9.0g		Student-level weight
wrschbw	float	%9.0g		School-level weight
high_school	byte	%8.0g		1 if highest level by either parent is high school
college	byte	%8.0g		1 if highest level by either parent is college
one_for	byte	%8.0g		1 if one parent foreign born
both_for	byte	%8.0g		1 if both parents are foreign born
test_lang	byte	%8.0g		1 if English (the test language) is spoken at home
pass_read	byte	%8.0g		1 if passed reading proficiency threshold
id_school	int	%8.0g		School ID

Sorted by:

For student i in school j , where the variable `id_school` identifies the schools, the variable `w_fstuw` is a student-level overall inclusion weight (w_{ij} , not $w_{i|j}$) adjusted for noninclusion and nonparticipation of students, and the variable `wrschbw` is the school-level weight w_j adjusted for oversampling of schools with more minority students. The weight adjustments do not interfere with the methods prescribed above, and thus we can treat the weight variables simply as w_{ij} and w_j , respectively.

Rabe-Hesketh and Skrondal (2006) fit a two-level logistic model for passing a reading proficiency threshold. We fit a two-level linear random-intercept model for socioeconomic index. Because we have w_{ij} and not $w_{i|j}$, we rescale using `pwscale(size)` and thus obtain results as if we had $w_{i|j}$.

```
. mixed isei female high_school college one_for both_for test_lang
> [pw=w_fstuw] || id_school: , pweight(wnrschbw) pwscale(size)
(output omitted)

Mixed-effects regression                Number of obs    =      2,069
Group variable: id_school              Number of groups =      148

Obs per group:
    min =          1
    avg =         14.0
    max =          28

Wald chi2(6) =      187.23
Prob > chi2   =      0.0000

Log pseudolikelihood = -1443093.9
                                (Std. Err. adjusted for 148 clusters in id_school)
```

isei	Coef.	Robust Std. Err.	z	P> z	[95% Conf. Interval]	
female	.59379	.8732886	0.68	0.497	-1.117824	2.305404
high_school	6.410618	1.500337	4.27	0.000	3.470011	9.351224
college	19.39494	2.121145	9.14	0.000	15.23757	23.55231
one_for	-.9584613	1.789947	-0.54	0.592	-4.466692	2.54977
both_for	-.2021101	2.32633	-0.09	0.931	-4.761633	4.357413
test_lang	2.519539	2.393165	1.05	0.292	-2.170978	7.210056
_cons	28.10788	2.435712	11.54	0.000	23.33397	32.88179

Random-effects Parameters	Estimate	Robust Std. Err.	[95% Conf. Interval]	
id_school: Identity var(_cons)	34.69374	8.574865	21.37318	56.31617
var(Residual)	218.7382	11.22111	197.8147	241.8748

Notes:

1. We specified the level-one weights using standard Stata weight syntax, that is, [pw=w_fstuw].
2. We specified the level-two weights via the pweight(wnrschbw) option as part of the random-effects specification for the id_school level. As such, it is treated as a school-level weight. Accordingly, wnrschbw needs to be constant within schools, and mixed did check for that before estimating.
3. Because our level-one weights are unconditional, we specified pwscale(size) to rescale them.
4. As is the case with other estimation commands in Stata, standard errors in the presence of sampling weights are robust.
5. Robust standard errors are clustered at the top level of the model, and this will always be true unless you specify vce(cluster clustvar), where clustvar identifies an even higher level of grouping.

As a form of sensitivity analysis, we compare the above with scaling via pwscale(gk). Because pwscale(gk) assumes w_{ij} , you want to first divide w_{ij} by w_j . But you can handle that within the weight specification itself.

```
. mixed isei female high_school college one_for both_for test_lang
> [pw=wfstwtw/wnrschbw] || id_school:, pweight(wnrschbw) pwscale(gk)
(output omitted)
Mixed-effects regression                Number of obs      =      2,069
Group variable: id_school              Number of groups   =      148
                                      Obs per group:
                                      min =          1
                                      avg =         14.0
                                      max =          28
                                      Wald chi2(6)      =      291.37
                                      Prob > chi2       =      0.0000
Log pseudolikelihood = -7270505.6      Prob > chi2       =      0.0000
                                      (Std. Err. adjusted for 148 clusters in id_school)
```

isei	Coef.	Robust Std. Err.	z	P> z	[95% Conf. Interval]	
female	-.3519458	.7436334	-0.47	0.636	-1.80944	1.105549
high_school	7.074911	1.139777	6.21	0.000	4.84099	9.308833
college	19.27285	1.286029	14.99	0.000	16.75228	21.79342
one_for	-.9142879	1.783091	-0.51	0.608	-4.409082	2.580506
both_for	1.214151	1.611885	0.75	0.451	-1.945085	4.373388
test_lang	2.661866	1.556491	1.71	0.087	-.3887996	5.712532
_cons	31.20145	1.907413	16.36	0.000	27.46299	34.93991

Random-effects Parameters	Estimate	Robust Std. Err.	[95% Conf. Interval]	
id_school: Identity				
var(_cons)	31.67522	6.792239	20.80622	48.22209
var(Residual)	226.2429	8.150714	210.8188	242.7955

The results are somewhat similar to before, which is good news from a sensitivity standpoint. Note that we specified `[pw=wfstwtw/wnrschbw]` and thus did the conversion from w_{ij} to $w_{i|j}$ within our call to `mixed`.

◀

We close this section with a bit of bad news. Although weight rescaling and the issues that arise have been well studied for two-level models, as pointed out by [Carle \(2009\)](#), “... a best practice for scaling weights across multiple levels has yet to be advanced.” As such, `pwscale()` is currently supported only for two-level models. If you are fitting a higher-level model with survey data, you need to make sure your sampling weights are conditional on selection at the previous stage and not overall inclusion weights, because there is currently no rescaling option to fall back on if you do not.

Small-sample inference for fixed effects

Researchers are often interested in making inferences about fixed effects in a linear mixed-effects model. In the special case where the data are balanced and the mixed-effects model has a simple covariance structure, the sampling distributions of the statistics for testing hypotheses about fixed effects are known to follow an F distribution with specific denominator degrees of freedom (DDF) under the null hypothesis. For example, the test statistics for testing hypotheses about fixed effects in balanced split-plot designs and balanced repeated-measures designs have exact t or F distributions. In general, however, the null sampling distributions of test statistics for fixed effects are not known and can only be approximated in more complicated mixed-effects models.

For a large sample, the null sampling distributions of the test statistics can be approximated by a normal distribution for a one-hypothesis test and a χ^2 distribution for a multiple-hypotheses test. This is the default behavior of `mixed`. However, these large-sample approximations may not be appropriate in small samples, and t and F distributions may provide better approximations.

You can specify the `dfmethod()` option to request small-sample inference for fixed effects. `mixed` with the `dfmethod()` option uses a t distribution for one-hypothesis tests and an F distribution for multiple-hypotheses tests for inference about fixed effects. We use `DF` to refer to degrees of freedom of a t distribution, and we use `DDF` to refer to denominator degrees of freedom of an F distribution.

Researchers have proposed various approximations that use t and F distributions but differ in how respective `DF` and `DDF` are computed (for example, Khuri, Mathew, and Sinha [1998]; Brown and Prescott [2015]; Schluchter and Elashoff [1990]; Elston [1998]; Kackar and Harville [1984]; Giesbrecht and Burns [1985]; Fai and Cornelius [1996]; and Kenward and Roger [1997, 2009]). `mixed` provides five methods with the `dfmethod()` option for calculating the `DF` of a t distribution: `residual`, `repeated`, `anova`, `satterthwaite`, and `kröger`.

Residual DDF (DF). This method uses the residual degrees of freedom, $n - p$, as the `DDF` for all tests of fixed effects. For a linear model without random effects and with i.i.d errors, the distributions of the test statistics for testing the fixed effects are exact t or F distributions with the residual `DF`.

Repeated DDF (DF). This method partitions the residual degrees of freedom into the between-subject degrees of freedom and the within-subject degrees of freedom. This partitioning of the degrees of freedom arises from balanced repeated-measures ANOVA analysis. If levels of a fixed effect change within a subject, then the within-subject degrees of freedom is assigned to the fixed effect of interest; otherwise, the between-subject degrees of freedom is assigned to that fixed effect. Winer, Brown, and Michels (1991) showed that this method is appropriate only when the data are balanced and the correlation structure is assumed to be spherical. The repeated `DDF` method is supported only with two-level models. For `DDF` methods accounting for unbalanced repeated measures, see, for example, Schluchter and Elashoff (1990).

ANOVA DDF (DF). This method mimics the traditional ANOVA method. It determines the `DDF` for a fixed effect depending on whether the corresponding covariate is contained in any of the random-effects equations. If the covariate is contained in a random-effects equation, the `DDF` for the fixed effect is computed as the number of levels of the level variable from that equation minus one. If the covariate is specified in more than one random-effects equation, the `DDF` for the fixed effect is computed as the smallest number of levels of the level variables from those equations minus one and is a conservative estimate of the true `DDF`. If the covariate is specified only in the fixed-effects equation, the `DDF` is computed as $\nu_{\text{ddf}} = n - \text{rank}(\mathbf{X}, \mathbf{Z})$. This method leads to an exact sampling distribution of the test statistics only when random effects are balanced and the residuals are i.i.d; see, for example, chapter 1.6 in Brown and Prescott (2015) for details.

Satterthwaite DDF (DF). This method performs a generalization of the Satterthwaite approximation based on Kackar and Harville (1984), Giesbrecht and Burns (1985), and Fai and Cornelius (1996). Giesbrecht and Burns (1985) developed a method of computing the `DDF` for a single-hypothesis test that is analogous to Satterthwaite's approximation of the degrees of freedom of a linear combination of ANOVA mean squares. For a multiple-hypotheses test, Fai and Cornelius (1996) proposed an extension of the Giesbrecht–Burns single-degree-of-freedom method. This method involves the spectral decomposition of the contrast matrix of the hypothesis test and repeated application of the single-degree-of-freedom t test. See *Denominator degrees of freedom* in *Methods and formulas* for more computational details.

Kenward–Roger DDF (DF). This method, developed by Kenward and Roger (1997), was designed to provide an approximation that improves the performance of hypothesis tests about fixed effects in small samples for complicated mixed-effects models and reproduces the exact inference available

for simpler mixed-effects models. It provides adjusted test statistics, more appropriate DDFs for the approximate F distributions when exact inference is not available, and yields the exact t and F distributions when exact inference is available. This method first accounts for the small-sample bias and the variability of the estimated random effects to obtain an adjusted estimator of the fixed-effects covariance matrix. Then, it proposes an approximate F test based on a scaled Wald test statistic that uses the adjusted variance–covariance estimator. See *Denominator degrees of freedom* in *Methods and formulas* for more computational details.

Residual, repeated, and ANOVA are known as “exact” methods in the literature. These methods are suitable only when the sampling distributions of the test statistics are known to be t or F . This is usually only known for certain classes of linear mixed-effects models with simple covariance structures and when data are balanced. These methods are available with both ML and REML estimation.

Satterthwaite and Kenward–Roger are known as “approximation” methods in the literature. These methods are for unbalanced data and complicated covariance structures where the sampling distributions of test statistics are unknown and can only be approximated. Both methods are available only with REML estimation. For single-hypothesis tests, DDFs calculated with the Kenward–Roger method are the same as those calculated with the Satterthwaite method, but they differ for multiple-hypotheses tests. Although DDFs of the two methods are the same for single-hypothesis tests, the inference is not the same because the Kenward–Roger method uses bias-adjusted standard errors.

Except for the special cases for which the sampling distributions are known, there is no definitive recommendation for which approximation performs best. [Schaalje, McBride, and Fellingham \(2002\)](#) compared the Satterthwaite method with the Kenward–Roger method via simulation using different covariance structures and various sample sizes. They concluded that the Kenward–Roger method outperforms the Satterthwaite method in most situations. They recommend using the Satterthwaite method only when the covariance structure of the data is compound symmetry and the sample size is moderately large. The Kenward–Roger method, however, is not guaranteed to work well in all situations. For example, for more complicated covariance structures and very small-sample sizes, the Kenward–Roger method may produce inflated type I error rates. In conclusion, you should choose your DDF method carefully. See, for example, [Schaalje, McBride, and Fellingham \(2002\)](#), [Chen and Wei \(2003\)](#), [Vallejo et al. \(2004\)](#), and [West, Welch, and Galecki \(2015\)](#) for a comparison of different approximations.

Both types of methods, exact and approximation, are available for single-hypothesis tests. For multiple-hypotheses tests, exact methods are available only if DDFs associated with fixed effects are the same for all tested covariates. See *Denominator degrees of freedom* in *Methods and formulas* for details.

► Example 13: Small-sample inference with a balanced repeated-measures design

Consider an example from [Winer, Brown, and Michels \(1991\)](#), table 4.3), also analyzed in [example 15](#) of [\[R\] anova](#), which reports the reaction time for five subjects who were tested with four drugs. The reaction time was recorded in the variable `score`. Assume that `person` is random (that is, we wish to infer to the larger population of possible subjects) and `drug` is fixed (that is, only four drugs are of interest). This is an example of a mixed-effects model with a simple covariance structure—a balanced repeated-measures design. The dataset contains only 20 observations, so we would like to account for the small sample in our analysis. Because this is a balanced repeated-measures design, we can use the repeated method to obtain small-sample inference for fixed effects. We specify the `dfmethod(repeated)` option with `mixed`. We also request REML estimates by specifying the `reml` option to account for the small number of groups.


```
. use http://www.stata-press.com/data/r15/t43
(T4.3 -- Winer, Brown, Michels)

. mixed score i.drug || person:, reml dfmethod(repeated)
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log restricted-likelihood = -49.640099
Iteration 1:   log restricted-likelihood = -49.640099

Computing standard errors:
Computing degrees of freedom:
Mixed-effects REML regression
Group variable: person
Number of obs      =      20
Number of groups   =       5
Obs per group:
    min =      4
    avg =     4.0
    max =      4
DF method: Repeated
DF:
    min =     4.00
    avg =    10.00
    max =    12.00
Log restricted-likelihood = -49.640099
F(3, 12.00) = 24.76
Prob > F = 0.0000
```

score	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
drug						
2	-.8	1.939072	-0.41	0.687	-5.024874	3.424874
3	-10.8	1.939072	-5.57	0.000	-15.02487	-6.575126
4	5.6	1.939072	2.89	0.014	1.375126	9.824874
_cons	26.4	3.149604	8.38	0.001	17.6553	35.1447

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
person: Identity				
var(_cons)	40.20004	30.10272	9.264606	174.4319
var(Residual)	9.399997	3.837532	4.22305	20.92325

```
LR test vs. linear model: chibar2(01) = 15.03      Prob >= chibar2 = 0.0001
```

In the table for fixed effects, *t* statistics are reported instead of the default *z* statistics. We can compare our small-sample inference with the corresponding large-sample inference for fixed effects. We do not need to rerun the estimation command, because we can obtain large-sample results upon replay by default.

```
. mixed
Mixed-effects REML regression      Number of obs      =      20
Group variable: person             Number of groups   =       5
                                   Obs per group:
                                   min =       4
                                   avg =      4.0
                                   max =       4
                                   Wald chi2(3)      =     74.28
                                   Prob > chi2       =     0.0000
Log restricted-likelihood = -49.640099
```

score	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
drug						
2	-.8	1.939072	-0.41	0.680	-4.600511	3.000511
3	-10.8	1.939072	-5.57	0.000	-14.60051	-6.999489
4	5.6	1.939072	2.89	0.004	1.799489	9.400511
_cons	26.4	3.149604	8.38	0.000	20.22689	32.57311

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
person: Identity					
	var(_cons)	40.20004	30.10272	9.264606	174.4319
	var(Residual)	9.399997	3.837532	4.22305	20.92325

```
LR test vs. linear model: chibar2(01) = 15.03      Prob >= chibar2 = 0.0001
```

Comparing the above large-sample inference for fixed effects of `drug` with the small-sample inference, we see that the p -value for the level 4 of `drug` changes from 0.014 to 0.004.

If we wanted to replay our small-sample estimation results, we would type

```
. mixed, small
(output omitted)
```

The specified DF method and summaries of the coefficient-specific DFs are reported in the output header. We can use the `dftable()` option to display a fixed-effects table containing coefficient-specific DFs. `dftable(pvalue)` reports the fixed-effects table containing DFs, t statistics, and p -values, and `dftable(ci)` reports the fixed-effects table containing DFs and confidence intervals.

```
. mixed, dftable(pvalue) norettable
Mixed-effects REML regression
Group variable: person
Number of obs      =      20
Number of groups   =       5
Obs per group:
    min =      4
    avg =     4.0
    max =      4
DF method: Repeated
DF:
    min =     4.00
    avg =    10.00
    max =    12.00
F(3,    12.00) =    24.76
Prob > F      =    0.0000
Log restricted-likelihood = -49.640099
```

score	Coef.	Std. Err.	DF	t	P> t
drug					
2	-.8	1.939072	12.0	-0.41	0.687
3	-10.8	1.939072	12.0	-5.57	0.000
4	5.6	1.939072	12.0	2.89	0.014
_cons	26.4	3.149604	4.0	8.38	0.001

Because levels of drug vary within person, the within-subject degrees of freedom, 12, are assigned to the coefficients for the levels of drug. The DF for the constant term is always the between-subject degrees of freedom, 4 in this example, because it is constant within random-effects levels.

The model *F* test is reported in the output header instead of the default χ^2 test. The *F* statistic for testing drug = 0 is 24.76 with DDF = 12, which agrees with the results of `anova, repeated()`:

```
. anova score person drug, repeated(drug)
Number of obs =      20
R-squared     = 0.9244
Root MSE     = 3.06594
Adj R-squared = 0.8803
Source | Partial SS      df      MS      F      Prob>F
-----+-----
Model |      1379         7       197    20.96  0.0000
person |      680.8         4     170.2   18.11  0.0001
drug   |      698.2         3    232.73333  24.76  0.0000
Residual |      112.8        12       9.4
Total |      1491.8       19    78.515789
Between-subjects error term: person
Levels: 5 (4 df)
Lowest b.s.e. variable: person
Repeated variable: drug
Huynh-Feldt epsilon = 1.0789
*Huynh-Feldt epsilon reset to 1.0000
Greenhouse-Geisser epsilon = 0.6049
Box's conservative epsilon = 0.3333
```

Source	df	F	Prob > F			
			Regular	H-F	G-G	Box
drug	3	24.76	0.0000	0.0000	0.0006	0.0076
Residual	12					

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
patient: Unstructured				
var(followup)	41.88772	18.79997	17.38009	100.9535
var(_cons)	524.9851	253.0205	204.1287	1350.175
cov(followup,_cons)	-140.4229	66.57623	-270.9099	-9.935907
tooth: Identity				
var(_cons)	47.45738	16.63034	23.8792	94.3165
var(Residual)	48.86704	10.50523	32.06479	74.47382

LR test vs. linear model: $\chi^2(4) = 91.12$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

Compared with the p -values of the large-sample results from [example 7](#), the p -values for `age` and `_cons` change substantially from 0.007 and 0.000 to 0.033 and 0.002, respectively. Note that for the Kenward–Roger method, not only the p -values and confidence intervals differ from those of the large-sample results but also the standard errors for the fixed effects differ. The standard errors differ because this method uses a bias-adjusted estimator of the variance–covariance matrix of fixed effects.

Now, let's fit the model using the Satterthwaite approximation:

```
. mixed gcf followup base_gcf cda age || patient: followup, cov(un)
> || tooth:, reml nolog dfmethod(satterthwaite)
```

Mixed-effects REML regression Number of obs = 110

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
patient	12	2	9.2	12
tooth	55	2	2.0	2

DF method: Satterthwaite	DF:	min =	10.41
		avg =	28.96
		max =	50.71

	F(4, 16.49)	=	1.87
Log restricted-likelihood = -420.92761	Prob > F	=	0.1638

gcf	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
followup	.3009815	1.936863	0.16	0.879	-3.963754	4.565717
base_gcf	-.0183127	.1433094	-0.13	0.899	-.3065704	.269945
cda	-.329303	.5292525	-0.62	0.537	-1.39197	.7333636
age	-.5773932	.2139656	-2.70	0.022	-1.051598	-.1031885
_cons	45.73862	12.55497	3.64	0.001	19.90352	71.57372

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
patient: Unstructured				
var(followup)	41.88772	18.79997	17.38009	100.9535
var(_cons)	524.9851	253.0205	204.1287	1350.175
cov(followup,_cons)	-140.4229	66.57623	-270.9099	-9.935907
tooth: Identity				
var(_cons)	47.45738	16.63034	23.8792	94.3165
var(Residual)	48.86704	10.50523	32.06479	74.47382

LR test vs. linear model: $\chi^2(4) = 91.12$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

Using the Satterthwaite method, we see that the p -value for `age` is 0.022 and for `_cons` is 0.001 and that these are again substantially different from their large-sample counterparts. On the other hand, unlike the standard errors for the Kenward–Roger method, those for the Satterthwaite method are the same as the standard errors from the large-sample results.

Looking at the DF summaries in the output header of the two methods, we notice that they are exactly the same. This is because DFs for fixed effects obtained using the Kenward–Roger and Satterthwaite methods are the same for single-hypothesis tests. (You can verify this by specifying, for example, `dftable(pvalue)` with the above commands or by using `estat df`; see [ME] [estat df](#).) The DDFs differ, however, for multiple-hypotheses tests. For example, DDF computed for the overall model test using `dfmethod(satterthwaite)` (16.49) is smaller than that computed using `dfmethod(kroger)` (27.96).

There are no general guidelines to which method should be preferred, but according to [Schaalje, McBride, and Fellingham \(2002\)](#), the Kenward–Roger method outperforms the Satterthwaite method when the variance–covariance structure of the random effects is unstructured, which is the case in our example.

◀

Determining which DDF method is best is a difficult task and may often need simulation. The choice of the method depends on the specified covariance structure, sample size, and imbalance of the data. No method applies to all situations; thus you should use caution when choosing among methods.

Stored results

`mixed` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_f)</code>	number of fixed-effects parameters
<code>e(k_r)</code>	number of random-effects parameters
<code>e(k_rs)</code>	number of variances
<code>e(k_rc)</code>	number of covariances
<code>e(k_res)</code>	number of residual-error parameters
<code>e(N_clust)</code>	number of clusters
<code>e(nrgroups)</code>	number of residual-error <code>by()</code> groups
<code>e(ar_p)</code>	AR order of residual errors, if specified
<code>e(ma_q)</code>	MA order of residual errors, if specified
<code>e(res_order)</code>	order of residual-error structure, if appropriate
<code>e(df_m)</code>	model degrees of freedom
<code>e(small)</code>	1 if <code>dfmethod()</code> option specified, 0 otherwise
<code>e(F)</code>	overall <i>F</i> test statistic when <code>dfmethod()</code> is specified
<code>e(ddf_m)</code>	model DDF
<code>e(df_max)</code>	maximum DF
<code>e(df_avg)</code>	average DF
<code>e(df_min)</code>	minimum DF
<code>e(ll)</code>	log (restricted) likelihood
<code>e(chi2)</code>	χ^2
<code>e(p)</code>	significance
<code>e(ll_c)</code>	log likelihood, comparison model
<code>e(chi2_c)</code>	χ^2 , comparison test
<code>e(df_c)</code>	degrees of freedom, comparison test
<code>e(p_c)</code>	significance, comparison test
<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(rc)</code>	return code
<code>e(converged)</code>	1 if converged, 0 otherwise

Macros

<code>e(cmd)</code>	<code>mixed</code>
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	name of dependent variable
<code>e(wtype)</code>	weight type (first-level weights)
<code>e(wexp)</code>	weight expression (first-level weights)
<code>e(fweightk)</code>	<code>fweight</code> variable for k th highest level, if specified
<code>e(pweightk)</code>	<code>pweight</code> variable for k th highest level, if specified
<code>e(ivars)</code>	grouping variables
<code>e(title)</code>	title in estimation output
<code>e(redim)</code>	random-effects dimensions
<code>e(vartypes)</code>	variance-structure types
<code>e(revars)</code>	random-effects covariates
<code>e(resopt)</code>	<code>residuals()</code> specification, as typed
<code>e(rstructure)</code>	residual-error structure
<code>e(rstructlab)</code>	residual-error structure output label
<code>e(rbyvar)</code>	residual-error <code>by()</code> variable, if specified
<code>e(rglabels)</code>	residual-error <code>by()</code> groups labels
<code>e(pwscale)</code>	sampling-weight scaling method
<code>e(timevar)</code>	residual-error <code>t()</code> variable, if specified
<code>e(dfmethod)</code>	DF method specified in <code>dfmethod()</code>
<code>e(dftitle)</code>	title for DF method
<code>e(chi2type)</code>	Wald; type of model χ^2 test
<code>e(clustvar)</code>	name of cluster variable
<code>e(vce)</code>	<code>vctype</code> specified in <code>vce()</code>
<code>e(vctype)</code>	title used to label Std. Err.
<code>e(method)</code>	ML or REML
<code>e(opt)</code>	type of optimization
<code>e(optmetric)</code>	<code>matsqrt</code> or <code>matlog</code> ; random-effects matrix parameterization

e(emonly)	emonly, if specified
e(ml_method)	type of ml method
e(technique)	maximization technique
e(datasignature)	the checksum
e(datasignaturevars)	variables used in calculation of checksum
e(properties)	b V
e(estat_cmd)	program used to implement estat
e(predict)	program used to implement predict
e(marginswtype)	weight type for margins
e(marginswexp)	weight expression for margins
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved

Matrices

e(b)	coefficient vector
e(N_g)	group counts
e(g_min)	group-size minimums
e(g_avg)	group-size averages
e(g_max)	group-size maximums
e(tmap)	ID mapping for unstructured residual errors
e(V)	variance–covariance matrix of the estimators
e(V_modelbased)	model-based variance
e(df)	parameter-specific DF for fixed effects
e(V_df)	variance–covariance matrix of the estimators when dfmethod(kroger) is specified

Functions

e(sample)	marks estimation sample

Methods and formulas

As given by (1), in the absence of weights we have the linear mixed model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}$$

where $\mathbf{y}$ is the $n \times 1$ vector of responses, $\mathbf{X}$ is an $n \times p$ design/covariate matrix for the fixed effects $\boldsymbol{\beta}$, and $\mathbf{Z}$ is the $n \times q$ design/covariate matrix for the random effects $\mathbf{u}$. The $n \times 1$ vector of errors $\boldsymbol{\epsilon}$ is for now assumed to be multivariate normal with mean 0 and variance matrix $\sigma_{\epsilon}^2 \mathbf{I}_n$. We also assume that $\mathbf{u}$ has variance–covariance matrix $\mathbf{G}$ and that $\mathbf{u}$ is orthogonal to $\boldsymbol{\epsilon}$ so that

$$\text{Var} \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\epsilon} \end{bmatrix} = \begin{bmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \sigma_{\epsilon}^2 \mathbf{I}_n \end{bmatrix}$$

Considering the combined error term $\mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}$, we see that $\mathbf{y}$ is multivariate normal with mean $\mathbf{X}\boldsymbol{\beta}$ and $n \times n$ variance–covariance matrix

$$\mathbf{V} = \mathbf{Z}\mathbf{G}\mathbf{Z}' + \sigma_{\epsilon}^2 \mathbf{I}_n$$

Defining $\boldsymbol{\theta}$ as the vector of unique elements of $\mathbf{G}$ results in the log likelihood

$$L(\boldsymbol{\beta}, \boldsymbol{\theta}, \sigma_{\epsilon}^2) = -\frac{1}{2} \{n \log(2\pi) + \log |\mathbf{V}| + (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \mathbf{V}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})\} \quad (9)$$

which is maximized as a function of $\boldsymbol{\beta}$, $\boldsymbol{\theta}$, and σ_{ϵ}^2 . As explained in chapter 6 of [Searle, Casella, and McCulloch \(1992\)](#), considering instead the likelihood of a set of linear contrasts $\mathbf{K}\mathbf{y}$ that do not depend on $\boldsymbol{\beta}$ results in the restricted log likelihood

$$L_R(\boldsymbol{\beta}, \boldsymbol{\theta}, \sigma_{\epsilon}^2) = L(\boldsymbol{\beta}, \boldsymbol{\theta}, \sigma_{\epsilon}^2) - \frac{1}{2} \log |\mathbf{X}' \mathbf{V}^{-1} \mathbf{X}| \quad (10)$$

Given the high dimension of $\mathbf{V}$, however, the log-likelihood and restricted log-likelihood criteria are not usually computed by brute-force application of the above expressions. Instead, you can simplify the problem by subdividing the data into independent clusters (and subclusters if possible) and using matrix decomposition methods on the smaller matrices that result from treating each cluster one at a time.

Consider the two-level model described previously in (2),

$$\mathbf{y}_j = \mathbf{X}_j\boldsymbol{\beta} + \mathbf{Z}_j\mathbf{u}_j + \boldsymbol{\epsilon}_j$$

for $j = 1, \dots, M$ clusters with cluster j containing n_j observations, with $\text{Var}(\mathbf{u}_j) = \boldsymbol{\Sigma}$, a $q \times q$ matrix.

Efficient methods for computing (9) and (10) are given in chapter 2 of [Pinheiro and Bates \(2000\)](#). Namely, for the two-level model, define $\boldsymbol{\Delta}$ to be the Cholesky factor of $\sigma_\epsilon^2 \boldsymbol{\Sigma}^{-1}$, such that $\sigma_\epsilon^2 \boldsymbol{\Sigma}^{-1} = \boldsymbol{\Delta}'\boldsymbol{\Delta}$. For $j = 1, \dots, M$, decompose

$$\begin{bmatrix} \mathbf{Z}_j \\ \boldsymbol{\Delta} \end{bmatrix} = \mathbf{Q}_j \begin{bmatrix} \mathbf{R}_{11j} \\ \mathbf{0} \end{bmatrix}$$

by using an orthogonal-triangular (QR) decomposition, with $\mathbf{Q}_j$ a $(n_j + q)$ -square matrix and $\mathbf{R}_{11j}$ a q -square matrix. We then apply $\mathbf{Q}_j$ as follows:

$$\begin{bmatrix} \mathbf{R}_{10j} \\ \mathbf{R}_{00j} \end{bmatrix} = \mathbf{Q}_j' \begin{bmatrix} \mathbf{X}_j \\ \mathbf{0} \end{bmatrix}; \quad \begin{bmatrix} \mathbf{c}_{1j} \\ \mathbf{c}_{0j} \end{bmatrix} = \mathbf{Q}_j' \begin{bmatrix} \mathbf{y}_j \\ \mathbf{0} \end{bmatrix}$$

Stack the $\mathbf{R}_{00j}$ and $\mathbf{c}_{0j}$ matrices, and perform the additional QR decomposition

$$\begin{bmatrix} \mathbf{R}_{001} & \mathbf{c}_{01} \\ \vdots & \vdots \\ \mathbf{R}_{00M} & \mathbf{c}_{0M} \end{bmatrix} = \mathbf{Q}_0 \begin{bmatrix} \mathbf{R}_{00} & \mathbf{c}_0 \\ \mathbf{0} & \mathbf{c}_1 \end{bmatrix}$$

[Pinheiro and Bates \(2000\)](#) show that ML estimates of $\boldsymbol{\beta}$, σ_ϵ^2 , and $\boldsymbol{\Delta}$ (the unique elements of $\boldsymbol{\Delta}$, that is) are obtained by maximizing the profile log likelihood (profiled in $\boldsymbol{\Delta}$)

$$L(\boldsymbol{\Delta}) = \frac{n}{2} \{ \log n - \log(2\pi) - 1 \} - n \log \|\mathbf{c}_1\| + \sum_{j=1}^M \log \left| \frac{\det(\boldsymbol{\Delta})}{\det(\mathbf{R}_{11j})} \right| \quad (11)$$

where $\|\cdot\|$ denotes the 2-norm. Following this maximization with

$$\hat{\boldsymbol{\beta}} = \mathbf{R}_{00}^{-1} \mathbf{c}_0; \quad \hat{\sigma}_\epsilon^2 = n^{-1} \|\mathbf{c}_1\|^2 \quad (12)$$

REML estimates are obtained by maximizing

$$L_R(\Delta) = \frac{n-p}{2} \{ \log(n-p) - \log(2\pi) - 1 \} - (n-p) \log \|\mathbf{c}_1\| - \log |\det(\mathbf{R}_{00})| + \sum_{j=1}^M \log \left| \frac{\det(\Delta)}{\det(\mathbf{R}_{11j})} \right| \quad (13)$$

followed by

$$\hat{\beta} = \mathbf{R}_{00}^{-1} \mathbf{c}_0; \quad \hat{\sigma}_\epsilon^2 = (n-p)^{-1} \|\mathbf{c}_1\|^2$$

For numerical stability, maximization of (11) and (13) is not performed with respect to the unique elements of Δ but instead with respect to the unique elements of the matrix square root (or matrix logarithm if the `matlog` option is specified) of Σ/σ_ϵ^2 ; define γ to be the vector containing these elements.

Once maximization with respect to γ is completed, $(\gamma, \sigma_\epsilon^2)$ is reparameterized to $\{\alpha, \log(\sigma_\epsilon)\}$, where α is a vector containing the unique elements of Σ , expressed as logarithms of standard deviations for the diagonal elements and hyperbolic arctangents of the correlations for off-diagonal elements. This last step is necessary 1) to obtain a joint variance–covariance estimate of the elements of Σ and σ_ϵ^2 ; 2) to obtain a parameterization under which parameter estimates can be interpreted individually, rather than as elements of a matrix square root (or logarithm); and 3) to parameterize these elements such that their ranges each encompass the entire real line.

Obtaining a joint variance–covariance matrix for the estimated $\{\alpha, \log(\sigma_\epsilon)\}$ requires the evaluation of the log likelihood (or log-restricted likelihood) with only β profiled out. For ML, we have

$$\begin{aligned} L^* \{\alpha, \log(\sigma_\epsilon)\} &= L\{\Delta(\alpha, \sigma_\epsilon^2), \sigma_\epsilon^2\} \\ &= -\frac{n}{2} \log(2\pi\sigma_\epsilon^2) - \frac{\|\mathbf{c}_1\|^2}{2\sigma_\epsilon^2} + \sum_{j=1}^M \log \left| \frac{\det(\Delta)}{\det(\mathbf{R}_{11j})} \right| \end{aligned}$$

with the analogous expression for REML.

The variance–covariance matrix of $\hat{\beta}$ is estimated as

$$\widehat{\text{Var}}(\hat{\beta}) = \hat{\sigma}_\epsilon^2 \mathbf{R}_{00}^{-1} (\mathbf{R}_{00}^{-1})'$$

but this does not mean that $\widehat{\text{Var}}(\hat{\beta})$ is identical under both ML and REML because $\mathbf{R}_{00}$ depends on Δ . Because $\hat{\beta}$ is asymptotically uncorrelated with $\{\hat{\alpha}, \log(\hat{\sigma}_\epsilon)\}$, the covariance of $\hat{\beta}$ with the other estimated parameters is treated as 0.

Parameter estimates are stored in `e(b)` as $\{\hat{\beta}, \hat{\alpha}, \log(\hat{\sigma}_\epsilon)\}$, with the corresponding (block-diagonal) variance–covariance matrix stored in `e(V)`. Parameter estimates can be displayed in this metric by specifying the `estmetric` option. However, in `mixed` output, variance components are most often displayed either as variances and covariances or as standard deviations and correlations.

EM iterations are derived by considering the $\mathbf{u}_j$ in (2) as missing data. Here we describe the procedure for maximizing the log likelihood via EM; the procedure for maximizing the restricted log likelihood is similar. The log likelihood for the full data $(\mathbf{y}, \mathbf{u})$ is

$$L_F(\beta, \Sigma, \sigma_\epsilon^2) = \sum_{j=1}^M \{ \log f_1(\mathbf{y}_j | \mathbf{u}_j, \beta, \sigma_\epsilon^2) + \log f_2(\mathbf{u}_j | \Sigma) \}$$

where $f_1(\cdot)$ is the density function for multivariate normal with mean $\mathbf{X}_j\beta + \mathbf{Z}_j\mathbf{u}_j$ and variance $\sigma_\epsilon^2\mathbf{I}_{n_j}$, and $f_2(\cdot)$ is the density for multivariate normal with mean $\mathbf{0}$ and $q \times q$ covariance matrix Σ . As before, we can profile β and σ_ϵ^2 out of the optimization, yielding the following EM iterative procedure:

1. For the current iterated value of $\Sigma^{(t)}$, fix $\hat{\beta} = \hat{\beta}(\Sigma^{(t)})$ and $\hat{\sigma}_\epsilon^2 = \hat{\sigma}_\epsilon^2(\Sigma^{(t)})$ according to (12).
2. Expectation step: Calculate

$$\begin{aligned} D(\Sigma) &\equiv E \left\{ L_F(\hat{\beta}, \Sigma, \hat{\sigma}_\epsilon^2) | \mathbf{y} \right\} \\ &= C - \frac{M}{2} \log \det(\Sigma) - \frac{1}{2} \sum_{j=1}^M E(\mathbf{u}'_j \Sigma^{-1} \mathbf{u}_j | \mathbf{y}) \end{aligned}$$

where C is a constant that does not depend on Σ , and the expected value of the quadratic form $\mathbf{u}'_j \Sigma^{-1} \mathbf{u}_j$ is taken with respect to the conditional density $f(\mathbf{u}_j | \mathbf{y}, \hat{\beta}, \Sigma^{(t)}, \hat{\sigma}_\epsilon^2)$.

3. Maximization step: Maximize $D(\Sigma)$ to produce $\Sigma^{(t+1)}$.

For general, symmetric Σ , the maximizer of $D(\Sigma)$ can be derived explicitly, making EM iterations quite fast.

For general, residual-error structures,

$$\text{Var}(\epsilon_j) = \sigma_\epsilon^2 \Lambda_j$$

where the subscript j merely represents that ϵ_j and Λ_j vary in dimension in unbalanced data, the data are first transformed according to

$$\mathbf{y}_j^* = \hat{\Lambda}_j^{-1/2} \mathbf{y}_j; \quad \mathbf{X}_j^* = \hat{\Lambda}_j^{-1/2} \mathbf{X}_j; \quad \mathbf{Z}_j^* = \hat{\Lambda}_j^{-1/2} \mathbf{Z}_j;$$

and the likelihood-evaluation techniques described above are applied to $\mathbf{y}_j^*$, $\mathbf{X}_j^*$, and $\mathbf{Z}_j^*$ instead. The unique elements of Λ , ρ , are estimated along with the fixed effects and variance components. Because σ_ϵ^2 is always estimated and multiplies the entire Λ_j matrix, $\hat{\rho}$ is parameterized to take this into account.

In the presence of sampling weights, following [Rabe-Hesketh and Skrondal \(2006\)](#), the weighted log pseudolikelihood for a two-level model is given as

$$L(\beta, \Sigma, \sigma_\epsilon^2) = \sum_{j=1}^M w_j \log \left[\int \exp \left\{ \sum_{i=1}^{n_j} w_{i|j} \log f_1(y_{ij} | \mathbf{u}_j, \beta, \sigma_\epsilon^2) \right\} f_2(\mathbf{u}_j | \Sigma) d\mathbf{u}_j \right] \quad (14)$$

where w_j is the inverse of the probability of selection for the j th cluster, $w_{i|j}$ is the inverse of the conditional probability of selection of individual i given the selection of cluster j , and $f_1(\cdot)$ and $f_2(\cdot)$ are the multivariate normal densities previously defined.

Weighted estimation is achieved through incorporating w_j and $w_{i|j}$ into the matrix decomposition methods detailed above to reflect replicated clusters for w_j and replicated observations within clusters for $w_{i|j}$. Because this estimation is based on replicated clusters and observations, frequency weights are handled similarly.

Rescaling of sampling weights can take one of three available forms:

Under `pwscale(size)`,

$$w_{i|j}^* = n_j w_{i|j} \left\{ \sum_{i=1}^{n_j} w_{i|j} \right\}^{-1}$$

Under `pwscale(effective)`,

$$w_{i|j}^* = w_{i|j} \left\{ \sum_{i=1}^{n_j} w_{i|j} \right\} \left\{ \sum_{i=1}^{n_j} w_{i|j}^2 \right\}^{-1}$$

Under both the above, w_j remains unchanged. For method `pwscale(gk)`, however, both weights are modified:

$$w_j^* = n_j^{-1} \sum_{i=1}^{n_j} w_{i|j} w_j; \quad w_{i|j}^* = 1$$

Under ML estimation, robust standard errors are obtained in the usual way (see [P] [_robust](#)) with the one distinction being that in multilevel models, robust variances are, at a minimum, clustered at the highest level. This is because given the form of the log likelihood, scores aggregate at the top-level clusters. For a two-level model, scores are obtained as the partial derivatives of $L_j(\beta, \Sigma, \sigma_\epsilon^2)$ with respect to $\{\beta, \alpha, \log(\sigma_\epsilon)\}$, where L_j is the log likelihood for cluster j and $L = \sum_{j=1}^M L_j$. Robust variances are not supported under REML estimation because the form of the log restricted likelihood does not lend itself to separation by highest-level clusters.

EM iterations always assume equal weighting and an independent, homoskedastic error structure. As such, with weighted data or when error structures are more complex, EM is used only to obtain starting values.

For extensions to models with three or more levels, see [Bates and Pinheiro \(1998\)](#) and [Rabe-Hesketh and Skrondal \(2006\)](#).

Denominator degrees of freedom

When the `dfmethod()` option is specified, `mixed` uses a t distribution with ν_{ddf} degrees of freedom to perform single-hypothesis tests for fixed effects $H_0: \beta_i = 0$ for $i = 1, 2, \dots, p$ or an F distribution with model numerator degrees of freedom and ν_{ddf_m} DDF for a model (joint) test of all coefficients (except the constant) being equal to zero. Denominator degrees of freedom ν_{ddf} and ν_{ddf_m} are computed according to the specified DDF method.

Residual DDF

This method uses the residual degrees of freedom as the DDF, $\nu_{\text{ddf}} = n - p$, where n is the total number of observations, and p is the rank of the design matrix $\mathbf{X}$.

Repeated DDF

This method partitions the residual degrees of freedom into the between-subject degrees of freedom and the within-subject degrees of freedom. This partitioning of the degrees of freedom arises from balanced repeated-measures ANOVA analysis. If levels of a fixed effect change within a subject, then the within-subject degrees of freedom is assigned to the fixed effect of interest; otherwise, the between-subject degrees of freedom is assigned to that fixed effect. See [Schluchter and Elashoff \(1990\)](#) for more computational details and, specifically, for the expressions of between-subject and within-subject degrees of freedom.

ANOVA DDF

This method determines the DDF for a fixed effect depending on whether the corresponding covariate is contained in any of the random-effects equations. If the covariate is contained in a random-effects equation, the DDF ν_{ddf} for the fixed effect is computed as the number of levels of the level variable from that equation minus one. If the covariate is specified in more than one random-effects equation, the DDF ν_{ddf} for the fixed effect is computed as the smallest number of levels of the level variables from those equations minus one and is a conservative estimate of the true DDF. If the covariate is specified only in the fixed-effects equation, the DDF is computed as $\nu_{\text{ddf}} = n - \text{rank}(\mathbf{X}, \mathbf{Z})$.

For example, suppose we have the following mixed model,

```
mixed y A B C || D: A || E: A B
```

where A, B, and C are fixed effects, and D and E are nested random effects. For the fixed effect A, ν_{ddf} is the smaller number of levels of variables D and E minus one because A is included in random-effects equations at both levels D and E. For the fixed effect B, ν_{ddf} is the number of levels of level variable E minus one because B is included in the random-effects equation at the level E. For the fixed effect C, $\nu_{\text{ddf}} = n - \text{rank}(\mathbf{X}, \mathbf{Z})$ because C is not included in any of the random-effects equations.

For the three methods above, the DDF for a model test of $H_0: \beta = \mathbf{0}$ is computed as follows. If all corresponding single-hypothesis tests $H_0: \beta_i = 0$ have the same DDF ν_{ddf} , then model DDF $\nu_{\text{ddf}_m} = \nu_{\text{ddf}}$. If the single-hypothesis DDF differs, then ν_{ddf_m} is not defined, and the large-sample χ^2 test is reported instead of the F test.

To provide formulas for the Satterthwaite and Kenward–Roger methods, consider a general linear-hypotheses test of fixed effects $H_0: \mathbf{C}'\beta = \mathbf{b}$ with a $p \times l$ matrix of linear hypotheses $\mathbf{C}$ of rank l .

The variance–covariance matrix of $\mathbf{y}$ is $\text{Var}(\mathbf{y}) = \mathbf{V} = \mathbf{ZGZ}' + \mathbf{R} = \mathbf{V}(\sigma)$ and can be viewed as a function of variance components σ ($r \times 1$). Suppose that the first two partial derivatives of $\mathbf{V}(\sigma)$ with respect to σ exist.

Let $\hat{\sigma}$ be the REML estimator of σ . Then, the REML estimator of the fixed effects β is the generalized least-squares estimator

$$\hat{\beta} = \{\mathbf{X}'\mathbf{V}^{-1}(\hat{\sigma})\mathbf{X}\}^{-1} \mathbf{X}'\mathbf{V}^{-1}(\hat{\sigma})\mathbf{Y}$$

where $\widehat{\text{Var}}(\hat{\beta}) = \hat{\Phi} = \Phi(\hat{\sigma}) = \{\mathbf{X}'\mathbf{V}^{-1}(\hat{\sigma})\mathbf{X}\}^{-1}$ is the conventional estimator of the variance–covariance matrix of the fixed effects $\hat{\beta}$, and $\mathbf{V}(\hat{\sigma})$ is the estimator of the covariance matrix of $\mathbf{y}$.

Under the null $H_0: \mathbf{C}'\beta = \mathbf{b}$, the F test statistic is

$$F = \frac{1}{l} (\mathbf{C}'\hat{\beta} - \mathbf{b})' (\mathbf{C}'\hat{\Phi}\mathbf{C})^{-1} (\mathbf{C}'\hat{\beta} - \mathbf{b})$$

and it has an F distribution with l numerator and ν_{ddf_C} DDF.

Satterthwaite DDF

This method is derived from the DDF formula of the original approximation attributable to [Satterthwaite \(1946\)](#):

$$\text{ddf} = \frac{2(\mathbf{C}'\hat{\Phi}\mathbf{C})^2}{\text{Var}(\mathbf{C}'\hat{\Phi}\mathbf{C})}$$

For a single-hypothesis test of $H_0: \mathbf{c}'\beta = \mathbf{b}$, where $\mathbf{c}$ and $\mathbf{b}$ are vectors of known constants, [Giesbrecht and Burns \(1985\)](#) proposed using

$$\nu_{\text{ddf}} = \frac{2(\mathbf{c}'\hat{\Phi}\mathbf{c})^2}{\text{Var}(\mathbf{c}'\hat{\Phi}\mathbf{c})} = \frac{2(\mathbf{c}'\hat{\Phi}\mathbf{c})^2}{\mathbf{d}'\mathbf{W}\mathbf{d}} \quad (15)$$

where $\mathbf{d}$ is a vector of partial derivatives of $\mathbf{c}'\Phi(\sigma)\mathbf{c}$ with respect to σ evaluated at $\hat{\sigma}$, and $\widehat{\text{Var}}(\hat{\sigma}) = \mathbf{W}$ is the estimator of the variance–covariance matrix of $\hat{\sigma}$ computed based on the expected information matrix $\mathbf{I_E}$ in (17) or on the observed information matrix if suboption `oim` of `dfmethod()` is specified.

For a multiple-hypotheses test (when the rank of $\mathbf{C}$ is greater than 1), [Fai and Cornelius \(1996\)](#) proposed an extension of the Giesbrecht–Burns single-degree-of-freedom method. Their method involves the spectral decomposition $\mathbf{C}'\hat{\Phi}\mathbf{C} = \mathbf{P}'\mathbf{D}\mathbf{P}$, where $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_l)$ is an orthogonal matrix of eigenvectors, and $\mathbf{D} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_l)$ is a diagonal matrix of the corresponding eigenvalues. Using this decomposition, we can write the F -test statistic as a sum of l independent approximate t random variates, $F = Q/l$ with

$$Q = \sum_{k=1}^l \frac{\{\mathbf{p}'_k(\mathbf{C}'\hat{\beta} - \mathbf{b})\}^2}{\lambda_k} = \sum_{k=1}^l t_{v_k}^2$$

where v_k is computed using (15). Because t_{v_k} s are independent and have approximate t distributions with v_k degrees of freedom,

$$E(Q) = \sum_{k=1}^l \frac{v_k}{v_k - 2} I(v_k > 2)$$

Then, the DDF for a multiple-hypotheses test can be approximately written as

$$\nu_{\text{ddf}_C} = \frac{2E(Q)}{E(Q) - l}$$

For more computational details of the Satterthwaite method, see [Fai and Cornelius \(1996\)](#).

Kenward–Roger DDF

This method was developed by [Kenward and Roger \(1997\)](#). It is based on adjusting the conventional variance–covariance estimator of fixed effects $\hat{\Phi}$ for small-sample bias and introducing a scaled F test that improves the small-sample performance of the conventional F test of fixed effects.

Kenward and Roger (1997) propose the adjusted estimator,

$$\hat{\Phi}_A = \hat{\Phi} + 2\hat{\Phi} \left\{ \sum_{i=1}^r \sum_{j=1}^r W_{ij} (\mathbf{Q}_{ij} - \mathbf{P}_i \hat{\Phi} \mathbf{P}_j - \frac{1}{4} \mathbf{R}_{ij}) \right\} \hat{\Phi} \quad (16)$$

where $\mathbf{P}_i = \mathbf{X}' \{ \partial \mathbf{V}^{-1}(\sigma) / \partial \sigma_i \} \mathbf{X}$, $\mathbf{Q}_{ij} = \mathbf{X}' \{ \partial \mathbf{V}^{-1}(\sigma) / \partial \sigma_i \} \mathbf{V}(\sigma) \{ \partial \mathbf{V}^{-1}(\sigma) / \partial \sigma_j \} \mathbf{X}$, and $\mathbf{R}_{ij} = \mathbf{X}' \mathbf{V}^{-1}(\sigma) \{ \partial^2 \mathbf{V}(\sigma) / \partial \sigma_i \partial \sigma_j \} \mathbf{V}^{-1}(\sigma) \mathbf{X}$ evaluated at $\hat{\sigma}$ and W_{ij} is the (i, j) th element of $\mathbf{W}$, the estimator of the variance-covariance matrix of $\hat{\sigma}$ computed from the inverse of the expected information matrix $\mathbf{I}_E$, where the element I_E^{ij} of $\mathbf{I}_E$ is defined as

$$2I_E^{ij} = \text{tr} \left(\frac{\partial \mathbf{V}^{-1}}{\partial \sigma_i} \mathbf{V} \frac{\partial \mathbf{V}^{-1}}{\partial \sigma_j} \mathbf{V} \right) - \text{tr}(2\Phi \mathbf{Q}_{ij} - \Phi \mathbf{P}_i \Phi \mathbf{P}_j) \quad (17)$$

Alternatively, you can use the observed information matrix as $\mathbf{W}$ by specifying suboption `oim` in `dfmethod()`.

All terms in (16), except those involving $\mathbf{R}_{ij}$, are invariant under reparameterization of the covariance structures. Also, the second derivative requires more computational resources and may not be numerically stable. For these reasons, the $\mathbf{R}_{ij}$ terms are ignored in the computation of $\hat{\Phi}_A$ in (16).

For multiple-hypotheses testing, Kenward and Roger (1997) propose the scaled F -test statistic, which under the null hypothesis can be written as

$$F_{\text{KR}} = \frac{\lambda}{l} (\mathbf{C}' \hat{\beta} - \mathbf{b})' (\mathbf{C}' \hat{\Phi}_A \mathbf{C})^{-1} (\mathbf{C}' \hat{\beta} - \mathbf{b})$$

and has an F distribution with l numerator and ν_{ddf_C} DDF. The scale factor $\lambda = \nu_{\text{ddf}_C} / (l - 1 + \nu_{\text{ddf}_C})$.

The DDF ν_{ddf_C} and λ are approximated as

$$\nu_{\text{ddf}_C} = 4 + \frac{l + 2}{l \times \rho - 1} \quad \text{and} \quad \lambda = \frac{\nu_{\text{ddf}_C}}{E^* (\nu_{\text{ddf}_C} - 2)}$$

where $\rho = V^* / 2(E^*)^2$ and E^* and V^* are the respective approximate mean and variance of the F_{KR} statistic; see Kenward and Roger (1997, 987) for expressions for E^* and V^* .

Acknowledgments

We thank Badi Baltagi of the Department of Economics at Syracuse University and Ray Carroll of the Department of Statistics at Texas A&M University for each providing us with a dataset used in this entry.

We also thank Mike Kenward of the Medical Statistics Unit at the London School of Hygiene and Tropical Medicine and James Roger (retired) of the Research Statistics Unit at GlaxoSmithKline for answering our questions about their methods.

Charles Roy Henderson (1911–1989) was born in Iowa and grew up on the family farm. His education in animal husbandry, animal nutrition, and statistics at Iowa State was interspersed with jobs in the Iowa Extension Service, Ohio University, and the U.S. Army. After completing his PhD, Henderson joined the Animal Science faculty at Cornell. He developed and applied statistical methods in the improvement of farm livestock productivity through genetic selection, with particular focus on dairy cattle. His methods are general and have been used worldwide in livestock breeding and beyond agriculture. Henderson's work on variance components and best linear unbiased predictions has proved to be one of the main roots of current mixed-model methods.

References

- Andrews, M. J., T. Schank, and R. Upward. 2006. [Practical fixed-effects estimation methods for the three-way error-components model](#). *Stata Journal* 6: 461–481.
- Baltagi, B. H., S. H. Song, and B. C. Jung. 2001. The unbalanced nested error component regression model. *Journal of Econometrics* 101: 357–381.
- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies.
<http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Brown, H., and R. Prescott. 2015. *Applied Mixed Models in Medicine*. 3rd ed. Chichester, UK: Wiley.
- Cameron, A. C., and P. K. Trivedi. 2010. *Microeconometrics Using Stata*. Rev. ed. College Station, TX: Stata Press.
- Canette, I. 2011. Including covariates in crossed-effects models. *The Stata Blog: Not Elsewhere Classified*.
<http://blog.stata.com/2010/12/22/including-covariates-in-crossed-effects-models/>.
- . 2014. Using gsem to combine estimation results. *The Stata Blog: Not Elsewhere Classified*.
<http://blog.stata.com/2014/08/18/using-gsem-to-combine-estimation-results/>.
- Carle, A. C. 2009. Fitting multilevel models in complex survey data with design weights: Recommendations. *BMC Medical Research Methodology* 9: 49.
- Chen, X., and L. Wei. 2003. A comparison of recent methods for the analysis of small-sample cross-over studies. *Statistics in Medicine* 22: 2821–2833.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39: 1–38.
- Diggle, P. J., P. J. Heagerty, K.-Y. Liang, and S. L. Zeger. 2002. *Analysis of Longitudinal Data*. 2nd ed. Oxford: Oxford University Press.
- Elston, D. A. 1998. Estimation of denominator degrees of freedom of F -distributions for assessing Wald statistics for fixed-effect factors in unbalanced mixed models. *Biometrics* 54: 1085–1096.
- Fai, A. H.-T., and P. L. Cornelius. 1996. Approximate F -tests of multiple degree of freedom hypotheses in generalized least squares analyses of unbalanced split-plot experiments. *Journal of Statistical Computation and Simulation* 54: 363–378.
- Fitzmaurice, G. M., N. M. Laird, and J. H. Ware. 2011. *Applied Longitudinal Analysis*. 2nd ed. Hoboken, NJ: Wiley.
- Giesbrecht, F. G., and J. C. Burns. 1985. Two-stage analysis based on a mixed model: Large-sample asymptotic theory and small-sample simulation results. *Biometrics* 41: 477–486.
- Goldstein, H. 1986. Efficient statistical modelling of longitudinal data. *Annals of Human Biology* 13: 129–141.
- Graubard, B. I., and E. L. Korn. 1996. Modelling the sampling design in the analysis of health surveys. *Statistical Methods in Medical Research* 5: 263–281.
- Harville, D. A. 1977. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association* 72: 320–338.
- Henderson, C. R. 1953. Estimation of variance and covariance components. *Biometrics* 9: 226–252.
- Hocking, R. R. 1985. *The Analysis of Linear Models*. Monterey, CA: Brooks/Cole.

- Horton, N. J. 2011. [Stata tip 95: Estimation of error covariances in a linear model](#). *Stata Journal* 11: 145–148.
- Huber, C. 2013a. Multilevel linear models in Stata, part 1: Components of variance. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2013/02/04/multilevel-linear-models-in-stata-part-1-components-of-variance/>.
- . 2013b. Multilevel linear models in Stata, part 2: Longitudinal data. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2013/02/18/multilevel-linear-models-in-stata-part-2-longitudinal-data/>.
- . 2014. How to simulate multilevel/longitudinal data. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2014/07/18/how-to-simulate-multilevel/longitudinal-data/>.
- Kacker, R. N., and D. A. Harville. 1984. Approximations for standard errors of estimators of fixed and random effects in mixed linear models. *Journal of the American Statistical Association* 79: 853–862.
- Kenward, M. G., and J. H. Roger. 1997. Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics* 53: 983–997.
- . 2009. An improved approximation to the precision of fixed effects from restricted maximum likelihood. *Computational Statistics & Data Analysis* 53: 2583–2595.
- Khuri, A. I., T. Mathew, and B. K. Sinha. 1998. *Statistical Tests for Mixed Linear Models*. New York: Wiley.
- Laird, N. M., and J. H. Ware. 1982. Random-effects models for longitudinal data. *Biometrics* 38: 963–974.
- LaMotte, L. R. 1973. Quadratic estimation of variance components. *Biometrics* 29: 311–330.
- Marchenko, Y. V. 2006. [Estimating variance components in Stata](#). *Stata Journal* 6: 1–21.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus. 2008. *Generalized, Linear, and Mixed Models*. 2nd ed. Hoboken, NJ: Wiley.
- Munnell, A. H. 1990. Why has productivity growth declined? Productivity and public investment. *New England Economic Review* Jan./Feb.: 3–22.
- Nichols, A. 2007. [Causal inference with observational data](#). *Stata Journal* 7: 507–541.
- Palmer, T. M., C. M. Macdonald-Wallis, D. A. Lawlor, and K. Tilling. 2014. [Estimating adjusted associations between random effects from multilevel models: The reffadjust package](#). *Stata Journal* 14: 119–140.
- Pantazis, N., and G. Touloumi. 2010. [Analyzing longitudinal data in the presence of informative drop-out: The jmrel command](#). *Stata Journal* 10: 226–251.
- Pfeffermann, D., C. J. Skinner, D. J. Holmes, H. Goldstein, and J. Rasbash. 1998. Weighting for unequal selection probabilities in multilevel models. *Journal of the Royal Statistical Society, Series B* 60: 23–40.
- Pierson, R. A., and O. J. Ginther. 1987. Follicular population dynamics during the estrous cycle of the mare. *Animal Reproduction Science* 14: 219–231.
- Pinheiro, J. C., and D. M. Bates. 2000. *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- Prosser, R., J. Rasbash, and H. Goldstein. 1991. *ML3 Software for 3-Level Analysis: User's Guide for V. 2*. London: Institute of Education, University of London.
- Rabe-Hesketh, S., and A. Skrondal. 2006. Multilevel modelling of complex survey data. *Journal of the Royal Statistical Society, Series A* 169: 805–827.
- . 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.
- Rao, C. R. 1973. *Linear Statistical Inference and Its Applications*. 2nd ed. New York: Wiley.
- Raudenbush, S. W., and A. S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.
- Robson, K., and D. Pevalin. 2016. *Multilevel Modeling in Plain Language*. London: Sage.
- Ruppert, D., M. P. Wand, and R. J. Carroll. 2003. *Semiparametric Regression*. Cambridge: Cambridge University Press.
- Satterthwaite, F. E. 1946. An approximate distribution of estimates of variance components. *Biometrics Bulletin* 2: 110–114.
- Schaalje, G. B., J. B. McBride, and G. W. Fellingham. 2002. Adequacy of approximations to distributions of test statistics in complex mixed linear models. *Journal of Agricultural, Biological, and Environmental Statistics* 7: 512–524.
- Schluchter, M. D., and J. D. Elashoff. 1990. Small-sample adjustments to tests with unbalanced repeated measures assuming several covariance structures. *Journal of Statistical Computation and Simulation* 37: 69–87.

- Schunck, R. 2013. Within and between estimates in random-effects models: Advantages and drawbacks of correlated random effects and hybrid models. *Stata Journal* 13: 65–76.
- Searle, S. R. 1989. Obituary: Charles Roy Henderson 1911–1989. *Biometrics* 45: 1333–1335.
- Searle, S. R., G. Casella, and C. E. McCulloch. 1992. *Variance Components*. New York: Wiley.
- Thompson, W. A., Jr. 1962. The problem of negative estimates of variance components. *Annals of Mathematical Statistics* 33: 273–289.
- Vallejo, G., P. Fernández, F. J. Herrero, and N. M. Conejo. 2004. Alternative procedures for testing fixed effects in repeated measures designs when assumptions are violated. *Psicothema* 16: 498–508.
- Verbeke, G., and G. Molenberghs. 2000. *Linear Mixed Models for Longitudinal Data*. New York: Springer.
- West, B. T., K. B. Welch, and A. T. Galecki. 2015. *Linear Mixed Models: A Practical Guide Using Statistical Software*. 2nd ed. Boca Raton, FL: Chapman & Hall/CRC.
- Wiggins, V. L. 2011. Multilevel random effects in xtmixed and sem—the long and wide of it. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2011/09/28/multilevel-random-effects-in-xtmixed-and-sem-the-long-and-wide-of-it/>.
- Winer, B. J., D. R. Brown, and K. M. Michels. 1991. *Statistical Principles in Experimental Design*. 3rd ed. New York: McGraw-Hill.

Also see

- [ME] **mixed postestimation** — Postestimation tools for mixed
- [ME] **meglm** — Multilevel mixed-effects generalized linear model
- [ME] **menl** — Nonlinear mixed-effects regression
- [ME] **me** — Introduction to multilevel mixed-effects models
- [MI] **estimation** — Estimation commands for use with mi estimate
- [BAYES] **bayes: mixed** — Bayesian multilevel linear regression
- [SEM] **intro 5** — Tour of models (*Multilevel mixed-effects models*)
- [XT] **xtrc** — Random-coefficients model
- [XT] **xtreg** — Fixed-, between-, and random-effects and population-averaged linear models
- [U] **20 Estimation and postestimation commands**

Postestimation commands	predict	margins
test and testparm	lincom	contrast
pwcompare	Remarks and examples	Stored results
Methods and formulas	References	Also see

Postestimation commands

The following postestimation commands are of special interest after `mixed`:

Command	Description
estat df	calculate and display degrees of freedom for fixed effects
estat group	summarize the composition of the nested groups
estat icc	estimate intraclass correlations
estat recovariance	display the estimated random-effects covariance matrices
estat sd	display variance components as standard deviations and correlations
estat wcorrelation	display model-implied within-cluster correlations and standard deviations

The following standard postestimation commands are also available:

Command	Description
contrast	contrasts and ANOVA-style joint tests of estimates
estat ic	Akaike's and Schwarz's Bayesian information criteria (AIC and BIC)
estat summarize	summary statistics for the estimation sample
estat vce	variance–covariance matrix of the estimators (VCE)
estimates	cataloging estimation results
hausman	Hausman's specification test
lincom	point estimates, standard errors, testing, and inference for linear combinations of coefficients
lrtest	likelihood-ratio test
margins	marginal means, predictive margins, marginal effects, and average marginal effects
marginsplot	graph the results from margins (profile plots, interaction plots, etc.)
nlcom	point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients
predict	predictions, residuals, influence statistics, and other diagnostic measures
predictnl	point estimates, standard errors, testing, and inference for generalized predictions
pwcompare	pairwise comparisons of estimates
test	Wald tests of simple and composite linear hypotheses
testnl	Wald tests of nonlinear hypotheses

predict

Description for predict

`predict` creates a new variable containing predictions such as linear predictions, standard errors, fitted values, residuals, and standardized residuals.

Menu for predict

Statistics > Postestimation

Syntax for predict

Syntax for obtaining BLUPs of random effects and the BLUPs' standard errors

```
predict [type] newvarsspec [if] [in], reffects [reses(newvarsspec)
relevel(levelvar) ]
```

Syntax for obtaining scores after ML estimation

```
predict [type] newvarsspec [if] [in], scores
```

Syntax for obtaining other predictions

```
predict [type] newvar [if] [in] [ , statistic relevel(levelvar) ]
```

newvarsspec is *stub** or *newvarlist*.

statistic	Description
Main	
<code>xb</code>	linear prediction for the fixed portion of the model only; the default
<code>stdp</code>	standard error of the fixed-portion linear prediction
<code><u>f</u>itted</code>	fitted values, fixed-portion linear prediction plus contributions based on predicted random effects
<code><u>r</u>esiduals</code>	residuals, response minus fitted values
* <code><u>r</u>standard</code>	standardized residuals

Unstarred statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample. Starred statistics are calculated only for the estimation sample, even when `if e(sample)` is not specified.

Options for predict

Main

xb, the default, calculates the linear prediction $\mathbf{x}\beta$ based on the estimated fixed effects (coefficients) in the model. This is equivalent to fixing all random effects in the model to their theoretical mean value of 0.

stdp calculates the standard error of the linear predictor $\mathbf{x}\beta$.

reffects calculates best linear unbiased predictions (BLUPs) of the random effects. By default, BLUPs for all random effects in the model are calculated. However, if the **relevel(levelvar)** option is specified, then BLUPs for only level *levelvar* in the model are calculated. For example, if *classes* are nested within *schools*, then typing

```
. predict b*, reffects relevel(school)
```

would produce BLUPs at the school level. You must specify *q* new variables, where *q* is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you.

Rabe-Hesketh and Skrondal (2012, sec. 2.11.2) discuss the link between the empirical Bayes predictions and BLUPs and how these predictions are unbiased. They are unbiased when the groups associated with the random effects are expected to vary in repeated samples. If you expect the groups to be fixed in repeated samples, then these predictions are no longer unbiased.

reses(newvarsspec) calculates the standard errors of the BLUPs of the random effects. By default, standard errors for all BLUPs in the model are calculated. However, if the **relevel(levelvar)** option is specified, then standard errors for only level *levelvar* in the model are calculated; see the **reffects** option.

You must specify *q* new variables, where *q* is the number of random-effects terms in the model (or level). However, it is much easier to just specify *stub** and let Stata name the variables *stub1*, *stub2*, ..., *stubq* for you. The new variables will have the same storage type as the corresponding random-effects variables.

The **reffects** and **reses()** options often generate multiple new variables at once. When this occurs, the random effects (or standard errors) contained in the generated variables correspond to the order in which the variance components are listed in the output of **mixed**. Still, examining the variable labels of the generated variables (with the **describe** command, for instance) can be useful in deciphering which variables correspond to which terms in the model.

fitted calculates fitted values, which are equal to the fixed-portion linear predictor *plus* contributions based on predicted random effects, or in mixed-model notation, $\mathbf{x}\beta + \mathbf{Zu}$. By default, the fitted values take into account random effects from all levels in the model; however, if the **relevel(levelvar)** option is specified, then the fitted values are fit beginning with the topmost level down to and including level *levelvar*. For example, if *classes* are nested within *schools*, then typing

```
. predict yhat_school, fitted relevel(school)
```

would produce school-level predictions. That is, the predictions would incorporate school-specific random effects but not those for each class nested within each school.

residuals calculates residuals, equal to the responses minus fitted values. By default, the fitted values take into account random effects from all levels in the model; however, if the **relevel(levelvar)** option is specified, then the fitted values are fit beginning at the topmost level down to and including level *levelvar*.

rstandard calculates standardized residuals, equal to the residuals multiplied by the inverse square root of the estimated error covariance matrix.

`scores` calculates the parameter-level scores, one for each parameter in the model including regression coefficients and variance components. The score for a parameter is the first derivative of the log likelihood (or log pseudolikelihood) with respect to that parameter. One score per highest-level group is calculated, and it is placed on the last record within that group. Scores are calculated in the estimation metric as stored in `e(b)`.

`scores` is not available after restricted maximum-likelihood (REML) estimation.

`relevel(levelvar)` specifies the level in the model at which predictions involving random effects are to be obtained; see the options above for the specifics. *levelvar* is the name of the model level and is either the name of the variable describing the grouping at that level or is `_all`, a special designation for a group comprising all the estimation data.

margins

Description for margins

`margins` estimates margins of response for linear predictions.

Menu for margins

Statistics > Postestimation

Syntax for margins

```
margins [marginlist] [ , options ]
margins [marginlist] , predict(statistic ...) [options]
```

statistic	Description
xb	linear predictor for the fixed portion of the model only; the default
<u>r</u> effects	not allowed with margins
<u>s</u> cores	not allowed with margins
stdp	not allowed with margins
<u>f</u> itted	not allowed with margins
<u>r</u> esiduals	not allowed with margins
<u>s</u> tandard	not allowed with margins

Statistics not allowed with `margins` are functions of stochastic quantities other than `e(b)`.
For the full syntax, see [R] [margins](#).

test and testparm

Description for test and testparm

`test` and `testparm`, by default, perform χ^2 tests of simple and composite linear hypotheses about the parameters for the most recently fit `mixed` model. They also support F tests with a small-sample adjustment for fixed effects.

Menu for test and testparm

Statistics > Postestimation

Syntax for test and testparm

```
test (spec) [(spec) ...] [, test_options small]
```

```
testparm varlist [, testparm_options small]
```

Options for test and testparm

Options

test_options; see [R] **test** options. Options `df()`, `common`, and `nosvyadjust` may not be specified together with `small`.

testparm_options; see options of `testparm` in [R] **test**. Options `df()` and `nosvyadjust` may not be specified together with `small`.

`small` specifies that F tests for fixed effects be carried out with the denominator degrees of freedom (DDF) obtained by the same method used in the most recently fit `mixed` model. If option `dfmethod()` is not specified in the previous `mixed` command, option `small` is not allowed. For certain methods, the DDF for some tests may not be available. See *Small-sample inference for fixed effects* in [ME] **mixed** for more details.

lincom

Description for lincom

`lincom`, by default, computes point estimates, standard errors, z statistics, p -values, and confidence intervals for linear combinations of coefficients after `mixed`. `lincom` also provides t statistics for linear combinations of the fixed effects, with the degrees of freedom calculated by the `DF` method specified in option `dfmethod()` of `mixed`.

Menu for lincom

Statistics > Postestimation

Syntax for lincom

```
lincom exp [, lincom_options small]
```

Options for lincom

lincom_options; see [R] [lincom](#) options. Option `df()` may not be specified together with `small`.

`small` specifies that t statistics for linear combinations of fixed effects be displayed with the degrees of freedom obtained by the same method used in the most recently fit `mixed` model. If option `dfmethod()` is not specified in the previous `mixed` command, option `small` is not allowed. For certain methods, the degrees of freedom for some linear combinations may not be available. See [Small-sample inference for fixed effects](#) in [ME] [mixed](#) for more details.

contrast

Description for contrast

`contrast`, by default, performs χ^2 tests of linear hypotheses and forms contrasts involving factor variables and their interactions for the most recently fit `mixed` model. `contrast` also supports tests with small-sample adjustments after `mixed`, `dfmethod()`.

Menu for contrast

Statistics > Postestimation

Syntax for contrast

```
contrast termlist [ , contrast_options small ]
```

Options for contrast

contrast_options; see [R] **contrast** options. Options `df()` and `nosvyadjust` may not be specified together with `small`.

`small` specifies that tests for contrasts be carried out with the DDF obtained by the same method used in the most recently fit `mixed` model. If option `dfmethod()` is not specified in the previous `mixed` command, option `small` is not allowed. For certain methods, the DDF for some contrasts may not be available. See *Small-sample inference for fixed effects* in [ME] **mixed** for more details.

pwcompare

Description for pwcompare

`pwcompare` performs pairwise comparisons across the levels of factor variables from the most recently fit mixed model. `pwcompare`, by default, reports the comparisons as contrasts (differences) of margins along with z tests or confidence intervals for the pairwise comparisons. `pwcompare` also supports t tests with small-sample adjustments after `mixed`, `dfmethod()`.

Menu for pwcompare

Statistics > Postestimation

Syntax for pwcompare

```
pwcompare marginlist [ , pwcompare_options small ]
```

Options for pwcompare

pwcompare_options; see [R] [pwcompare](#) options. Option `df()` may not be specified together with `small`.

`small` specifies that t tests for pairwise comparisons be carried out with the degrees of freedom obtained by the same method used in the most recently fit mixed model with the `dfmethod()` option. If option `dfmethod()` is not specified in the previous `mixed` command, option `small` is not allowed. For certain methods, the degrees of freedom for some pairwise comparisons may not be available. See [Small-sample inference for fixed effects](#) in [ME] [mixed](#) for more details.

Remarks and examples

Various predictions, statistics, and diagnostic measures are available after fitting a mixed model using `mixed`. For the most part, calculation centers around obtaining BLUPs of the random effects. Random effects are not estimated when the model is fit but instead need to be predicted after estimation. Calculation of intraclass correlations, estimating the dependence between responses for different levels of nesting, may also be of interest.

► Example 1: Obtaining predictions of random effects and checking model fit

In [example 3](#) of [ME] [mixed](#), we modeled the weights of 48 pigs measured on nine successive weeks as

$$\text{weight}_{ij} = \beta_0 + \beta_1 \text{week}_{ij} + u_{0j} + u_{1j} \text{week}_{ij} + \epsilon_{ij} \quad (1)$$

for $i = 1, \dots, 9$, $j = 1, \dots, 48$, $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$, and u_{0j} and u_{1j} normally distributed with mean 0 and variance–covariance matrix

$$\Sigma = \text{Var} \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} = \begin{bmatrix} \sigma_{u0}^2 & \sigma_{01} \\ \sigma_{01} & \sigma_{u1}^2 \end{bmatrix}$$

```
. use http://www.stata-press.com/data/r15/pig
(Longitudinal analysis of pig weights)
. mixed weight week || id: week, covariance(unstructured)

Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -868.96185
Iteration 1:   log likelihood = -868.96185
Computing standard errors:
Mixed-effects ML regression              Number of obs   =       432
Group variable: id                      Number of groups =        48
                                         Obs per group:
                                         min =          9
                                         avg =         9.0
                                         max =          9
                                         Wald chi2(1)    =    4649.17
                                         Prob > chi2      =     0.0000

Log likelihood = -868.96185
```

weight	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
week	6.209896	.0910745	68.18	0.000	6.031393	6.388399
_cons	19.35561	.3996387	48.43	0.000	18.57234	20.13889

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: Unstructured				
var(week)	.3715251	.0812958	.2419532	.570486
var(_cons)	6.823363	1.566194	4.351297	10.69986
cov(week,_cons)	-.0984378	.2545767	-.5973991	.4005234
var(Residual)	1.596829	.123198	1.372735	1.857505

LR test vs. linear model: chi2(3) = 764.58 Prob > chi2 = 0.0000
Note: LR test is conservative and provided only for reference.

Rather than see the estimated variance components listed as variance and covariances as above, we can instead see them as correlations and standard deviations using `estat sd`; see [ME] [estat sd](#).

```
. estat sd
```

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
id: Unstructured				
sd(week)	.6095286	.0666874	.4918874	.7553052
sd(_cons)	2.612157	.2997895	2.085976	3.271064
corr(week,_cons)	-.0618257	.1575911	-.3557072	.243182
sd(Residual)	1.263657	.0487466	1.171638	1.362903

We can use `estat recovariance` to display the estimated variance components $\hat{\Sigma}$ as a correlation matrix; see [ME] [estat recovariance](#).

```
. estat recovariance, correlation
Random-effects correlation matrix for level id
```

	week	_cons
week	1	
_cons	-.0618257	1

Finally, we can use `estat wcorrelation` to display the within-cluster marginal standard deviations and correlations for one of the clusters; see [ME] [estat wcorrelation](#).

```
. estat wcorrelation, format(%4.2g)
```

Standard deviations and correlations for id = 1:

Standard deviations:

obs	1	2	3	4	5	6	7	8	9
sd	2.9	3.1	3.3	3.7	4.1	4.5	5	5.5	6.1

Correlations:

obs	1	2	3	4	5	6	7	8	9
1	1								
2	.8	1							
3	.77	.83	1						
4	.72	.81	.86	1					
5	.67	.78	.85	.89	1				
6	.63	.75	.83	.88	.91	1			
7	.59	.72	.81	.87	.91	.93	1		
8	.55	.69	.79	.86	.9	.93	.94	1	
9	.52	.66	.77	.85	.89	.92	.94	.95	1

Because within-cluster correlations can vary between clusters, `estat wcorrelation` by default displays the results for the first cluster. In this example, each cluster (pig) has the same number of observations, and the timings of measurements (`week`) are the same between clusters. Thus the within-cluster correlations are the same for all the clusters. In [example 1](#) of [ME] [estat wcorrelation](#), we fit a model where different clusters have different within-cluster correlations and show how to display these correlations.

We can also obtain BLUPs of the pig-level random effects (u_{0j} and u_{1j}). We need to specify the variables to be created in the order `u1 u0` because that is the order in which the corresponding variance components are listed in the output (`week _cons`). We obtain the predictions and list them for the first 10 pigs.

```
. predict u1 u0, reffects
. by id, sort: generate tolist = (_n==1)
. list id u0 u1 if id <=10 & tolist
```

	id	u0	u1
1.	1	.2369444	-.3957636
10.	2	-1.584127	.510038
19.	3	-3.526551	.3200372
28.	4	1.964378	-.7719702
37.	5	1.299236	-.9241479
46.	6	-1.147302	-.5448151
55.	7	-2.590529	.0394454
64.	8	-1.137067	-.1696566
73.	9	-3.189545	-.7365507
82.	10	1.160324	.0030772

If you forget how to order your variables in `predict`, or if you use `predict stub*`, remember that `predict` labels the generated variables for you to avoid confusion.

```
. describe u0 u1
```

variable name	storage type	display format	value label	variable label
u0	float	%9.0g		BLUP r.e. for id: _cons
u1	float	%9.0g		BLUP r.e. for id: week


Examining (1), we see that within each pig, the successive weight measurements are modeled as simple linear regression with intercept $\beta_0 + u_{j0}$ and slope $\beta_1 + u_{j1}$. We can generate estimates of the pig-level intercepts and slopes with

```
. generate intercept = _b[_cons] + u0
. generate slope = _b[week] + u1
. list id intercept slope if id<=10 & tolist
```

	id	interc~t	slope
1.	1	19.59256	5.814132
10.	2	17.77149	6.719934
19.	3	15.82906	6.529933
28.	4	21.31999	5.437926
37.	5	20.65485	5.285748
46.	6	18.20831	5.665081
55.	7	16.76509	6.249341
64.	8	18.21855	6.040239
73.	9	16.16607	5.473345
82.	10	20.51594	6.212973

Thus we can plot estimated regression lines for each of the pigs. Equivalently, we can just plot the fitted values because they are based on both the fixed and the random effects:

```
. predict fitweight, fitted
. twoway connected fitweight week if id<=10, connect(L)
```




We can also generate standardized residuals and see whether they follow a standard normal distribution, as they should in any good-fitting model:

```
. predict rs, rstandard
```

```
. summarize rs
```

Variable	Obs	Mean	Std. Dev.	Min	Max
rs	432	1.01e-09	.8929356	-3.621446	3.000929

```
. qnorm rs
```



◀

► Example 2: Estimating the intraclass correlation

Following [Rabe-Hesketh and Skrondal \(2012, chap. 2\)](#), we fit a two-level random-effects model for human peak-expiratory-flow rate. The subjects were each measured twice with the Mini-Wright peak-flow meter. It is of interest to determine how reliable the meter is as a measurement device. The intraclass correlation provides a measure of reliability. Formally, in a two-level random-effects model, the intraclass correlation corresponds to the correlation of measurements within the same individual and also to the proportion of variance explained by the individual random effect.

First, we fit the two-level model with mixed:

```
. use http://www.stata-press.com/data/r15/pefrate, clear
(Peak-expiratory-flow rate)
. mixed wm || id:
Performing EM optimization:
Performing gradient-based optimization:
Iteration 0:   log likelihood = -184.57839
Iteration 1:   log likelihood = -184.57839
Computing standard errors:
Mixed-effects ML regression              Number of obs   =       34
Group variable: id                      Number of groups =       17
Obs per group:
                                         min =         2
                                         avg =        2.0
                                         max =         2
                                         Wald chi2(0)    =         .
                                         Prob > chi2     =         .

Log likelihood = -184.57839
```

wm	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_cons	453.9118	26.18617	17.33	0.000	402.5878	505.2357

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
id: Identity	var(_cons)	11458.94	3998.952	5782.176	22708.98
	var(Residual)	396.441	135.9781	202.4039	776.4942

LR test vs. linear model: chibar2(01) = 46.27 Prob >= chibar2 = 0.0000

Now we use `estat icc` to estimate the intraclass correlation:

```
. estat icc
Intraclass correlation
```

Level	ICC	Std. Err.	[95% Conf. Interval]	
id	.9665602	.0159495	.9165853	.9870185

This correlation is close to 1, indicating that the Mini-Wright peak-flow meter is reliable. But as noted by [Rabe-Hesketh and Skrondal \(2012\)](#), the reliability is not only a characteristic of the instrument but also of the between-subject variance. Here we see that the between-subject standard deviation, `sd(_cons)`, is much larger than the within-subject standard deviation, `sd(Residual)`.

In the presence of fixed-effects covariates, `estat icc` reports the residual intraclass correlation, the correlation between measurements conditional on the fixed-effects covariates. This is equivalent to the correlation of the model residuals.

In the presence of random-effects covariates, the intraclass correlation is no longer constant and depends on the values of the random-effects covariates. In this case, `estat icc` reports conditional intraclass correlations assuming 0 values for all random-effects covariates. For example, in a two-level model, this conditional correlation represents the correlation of the residuals for two measurements on the same subject, which both have random-effects covariates equal to 0. Similarly to the interpretation of intercept variances in random-coefficient models ([Rabe-Hesketh and Skrondal 2012](#), chap. 4),

interpretation of this conditional intraclass correlation relies on the usefulness of the 0 baseline values of random-effects covariates. For example, mean centering of the covariates is often used to make a 0 value a useful reference.

See [ME] [estat icc](#) for more information.



► Example 3: Estimating residual intraclass correlations

In [example 4](#) of [ME] [mixed](#), we estimated a Cobb–Douglas production function with random intercepts at the region level and at the state-within-region level:

$$y_{jk} = \mathbf{X}_{jk}\beta + u_k^{(3)} + u_{jk}^{(2)} + \epsilon_{jk}$$

```
. use http://www.stata-press.com/data/r15/productivity
(Public Capital Productivity)
. mixed gsp private emp hwy water other unemp || region: || state:
(output omitted)
```

We can use [estat group](#) to see how the data are broken down by state and region:

```
. estat group
```

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
region	9	51	90.7	136
state	48	17	17.0	17

We are reminded that we have balanced productivity data for 17 years for each state.

We can use [predict, fitted](#) to get the fitted values

$$\hat{y}_{jk} = \mathbf{X}_{jk}\hat{\beta} + \hat{u}_k^{(3)} + \hat{u}_{jk}^{(2)}$$

but if we instead want fitted values at the region level, that is,

$$\hat{y}_{jk} = \mathbf{X}_{jk}\hat{\beta} + \hat{u}_k^{(3)}$$

we need to use the [relevel\(\)](#) option:


```
. predict gsp_region, fitted relevel(region)
. list gsp gsp_region in 1/10
```

	gsp	gsp_re~n
1.	10.25478	10.40529
2.	10.2879	10.42336
3.	10.35147	10.47343
4.	10.41721	10.52648
5.	10.42671	10.54947
6.	10.4224	10.53537
7.	10.4847	10.60781
8.	10.53111	10.64727
9.	10.59573	10.70503
10.	10.62082	10.72794

❏ Technical note

Out-of-sample predictions are permitted after `mixed`, but if these predictions involve BLUPs of random effects, the integrity of the estimation data must be preserved. If the estimation data have changed since the mixed model was fit, `predict` will be unable to obtain predicted random effects that are appropriate for the fitted model and will give an error. Thus to obtain out-of-sample predictions that contain random-effects terms, be sure that the data for these predictions are in observations that augment the estimation data.



We can use `estat icc` to estimate residual intraclass correlations between productivity years in the same region and in the same state and region.

```
. estat icc
Residual intraclass correlation
```

Level	ICC	Std. Err.	[95% Conf. Interval]	
region	.159893	.127627	.0287143	.5506202
state region	.8516265	.0301733	.7823466	.9016272

`estat icc` reports two intraclass correlations for this three-level nested model. The first is the level-3 intraclass correlation at the region level, the correlation between productivity years in the same region. The second is the level-2 intraclass correlation at the state-within-region level, the correlation between productivity years in the same state and region.

Conditional on the fixed-effects covariates, we find that annual productivity is only slightly correlated within the same region, but it is highly correlated within the same state and region. We estimate that state and region random effects compose approximately 85% of the total residual variance.



➤ Example 4: Small-sample adjusted tests for fixed effects

To illustrate the use of `test` and `testparm` with the `small` option for small-sample adjusted tests for fixed effects, we refit the dental veneer data from [example 14](#) of [\[ME\] mixed](#) using the Satterthwaite method (option `dfmethod(satterthwaite)`) to compute the DF for fixed effects.

```
. use http://www.stata-press.com/data/r15/veneer, clear
(Dental veneer data)

. mixed gcf followup base_gcf cda age
> || patient: followup, covariance(unstructured)
> || tooth:, reml nolog dfmethod(satterthwaite)
```

Mixed-effects REML regression Number of obs = 110

Group Variable	No. of Groups	Observations per Group		
		Minimum	Average	Maximum
patient	12	2	9.2	12
tooth	55	2	2.0	2

DF method: Satterthwaite	DF:	min =	10.41
		avg =	28.96
		max =	50.71
	F(4, 16.49)	=	1.87
Log restricted-likelihood = -420.92761	Prob > F	=	0.1638

gcf	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
followup	.3009815	1.936863	0.16	0.879	-3.963754 4.565717
base_gcf	-.0183127	.1433094	-0.13	0.899	-.3065704 .269945
cda	-.329303	.5292525	-0.62	0.537	-1.39197 .7333636
age	-1.5773932	.2139656	-2.70	0.022	-1.051598 -.1031885
_cons	45.73862	12.55497	3.64	0.001	19.90352 71.57372

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
patient: Unstructured				
var(followup)	41.88772	18.79997	17.38009	100.9535
var(_cons)	524.9851	253.0205	204.1287	1350.175
cov(followup,_cons)	-140.4229	66.57623	-270.9099	-9.935907
tooth: Identity				
var(_cons)	47.45738	16.63034	23.8792	94.3165
var(Residual)	48.86704	10.50523	32.06479	74.47382

LR test vs. linear model: $\chi^2(4) = 91.12$ Prob > $\chi^2 = 0.0000$

Note: LR test is conservative and provided only for reference.

Now we can, for example, test the hypotheses that all fixed effects are zero by typing

```
. testparm *, small
( 1) [gcf]followup = 0
( 2) [gcf]base_gcf = 0
( 3) [gcf]cda = 0
( 4) [gcf]age = 0

F( 4, 16,49) = 1.87
Prob > F = 0.1638
```

The F statistic for the overall test is 1.87, and the DDF is estimated to be 16.49. These results are different from the model test using the Kenward–Roger DDF method reported in the header of the estimation output in [example 1](#) of [ME] **estat df** (the F statistic is 1.47, and the model DDF is 27.96).

The results differ because the Kenward–Roger method uses an adjusted F -test statistic and adjusts the fixed-effects variance–covariance estimator for a small sample. Both methods, however, lead to the same conclusion of no joint significance of the fixed effects.

Without option `small`, the commands `test` and `testparm` report large-sample χ^2 Wald tests. We can compare the small-sample and large-sample tests of the joint hypotheses that the coefficient on `followup` and the coefficient on `age` equal zero.

```
. test followup = age = 0, small
( 1) [gcf]followup - [gcf]age = 0
( 2) [gcf]followup = 0
      F( 2, 10.75) =      3.65
      Prob > F      =      0.0617

. test followup = age = 0
( 1) [gcf]followup - [gcf]age = 0
( 2) [gcf]followup = 0
      chi2( 2) =      7.30
      Prob > chi2 =      0.0260
```

The DDF of the F test, which is computed using the Satterthwaite method from our posted results, is 10.75. The p -values are very different (0.0617 versus 0.0260), and they lead to different conclusions of whether we should reject the null hypotheses at the $\alpha = 0.05$ level.

Similarly, you can use the `small` option with `lincom` to perform small-sample inference for linear combinations of fixed effects.



► Example 5: Small-sample adjusted contrasts

As we did with `test`, after fitting a mixed model with the `dfmethod()` option for small-sample adjustment, we can use the `small` option with `contrast` to adjust for a small sample when estimating contrasts. Suppose we have collected data on a vigilance performance test. This experiment has been designed to test the response latency scores of two modes of signal during a four-hour monitoring period. This is a split-plot factorial design where `signal` is the whole-plot factor, `hour` is the subplot factor, and `subject` is the block factor. The whole-plot factor and the subplot factor are fixed; the block factor is random. Also, suppose that two measurements are missing in this dataset.

```
. use http://www.stata-press.com/data/r15/vptscores, clear
(Vigilance Performance Test Scores With Missing Data)
. tabdisp subject hour, cellvar(score) by(signal) concise missing
```

Signal and Subject ID	Monitoring period			
	1	2	3	4
Auditory				
1	3	4	7	7
2	6	5	.	8
3	3	4	7	9
4	3	3	6	8
Visual				
5	1	2	5	10
6	2	3	6	.
7	2	4	5	9
8	2	3	6	11

We start by fitting a mixed model. Because the dataset is small and unbalanced, we apply the Kenward–Roger method for small-sample adjustment:

```

. mixed score signal##hour || subject:, reml dfmethod(kroger) nolog nogroup
Mixed-effects REML regression          Number of obs   =          30
DF method: Kenward-Roger              DF:           min =         16.02
                                      avg =         16.76
                                      max =         18.29
                                      F(7,    16.08)    =         43.84
                                      Prob > F        =         0.0000

Log restricted-likelihood =    -32.9724

```

score	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
signal						
Visual	-2	.6288677	-3.18	0.005	-3.319693	-.6803071
hour						
2	.25	.5359916	0.47	0.647	-.8861371	1.386137
3	3.108222	.5911044	5.26	0.000	1.859163	4.357281
4	4.25	.5359916	7.93	0.000	3.113863	5.386137
signal#hour						
Visual#2	1	.7580066	1.32	0.206	-.6067405	2.606741
Visual#3	.6417778	.7979294	0.80	0.433	-1.046666	2.330221
Visual#4	4.044205	.7979294	5.07	0.000	2.355762	5.732649
_cons	3.75	.4446766	8.43	0.000	2.816836	4.683164

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
subject: Identity				
var(_cons)	.2163751	.2345718	.0258477	1.811312
var(Residual)	.574574	.2062107	.2843515	1.161011

```

LR test vs. linear model: chibar2(01) = 1.55          Prob >= chibar2 = 0.1069

```

We can test the main effects and the interaction effects by typing the `contrast` command. With the `small` option, `contrast` reports small-sample adjusted F tests. Without the `small` option, `contrast` performs large-sample χ^2 Wald tests. Below is the comparison of the small-sample and the large-sample contrasts:

```
. contrast signal##hour, small
Contrasts of marginal linear predictions
Margins      : asbalanced
```

	df	ddf	F	P>F
score				
signal	1	5.95	1.78	0.2307
hour	3	16.35	100.62	0.0000
signal#hour	3	16.35	9.66	0.0007

```
. contrast signal##hour
Contrasts of marginal linear predictions
Margins      : asbalanced
```

	df	chi2	P>chi2
score			
signal	1	1.79	0.1810
hour	3	304.95	0.0000
signal#hour	3	29.35	0.0000

From these results, we can see that the *p*-values for the main effect of `signal` and the interaction effect vary between small-sample and large-sample tests. However, both tests indicate that the `hour` effect and the interaction effects are significant. We can decompose the interaction effect into separate interaction contrasts for further investigation.

```
. contrast r.signal#ar.hour, small
Contrasts of marginal linear predictions
Margins      : asbalanced
```

	df	ddf	F	P>F
score				
signal#hour				
(Visual vs Auditory) (2 vs 1)	1	16.02	1.74	0.2056
(Visual vs Auditory) (3 vs 2)	1	16.37	0.20	0.6594
(Visual vs Auditory) (4 vs 3)	1	16.66	16.57	0.0008
Joint	3	16.35	9.66	0.0007

	Contrast	Std. Err.	DF	[95% Conf. Interval]	
score					
signal#hour					
(Visual vs Auditory)					
(2 vs 1)	1	.7580066	16.0	-.6067405	2.606741
(Visual vs Auditory)					
(3 vs 2)	-.3582222	.7979294	16.4	-2.046666	1.330221
(Visual vs Auditory)					
(4 vs 3)	3.402427	.8359478	16.7	1.635991	5.168863

From previous analysis, we already knew the overall interaction was significant. From the decomposition, we can easily see that the overall significance is driven by differences in the third and fourth hours; the change in response latency from hour three to hour four is greater for visual signals than for auditory signals.

We can also calculate the pairwise differences of the hourly marginal means by typing the `pwcompare` command. With the `small` option, `pwcompare` reports small-sample adjusted pairwise comparisons along with the degrees of freedom for each pairwise comparison.

```
. pwcompare hour, small
Pairwise comparisons of marginal linear predictions
Margins      : asbalanced
```

	Contrast	Std. Err.	DF	Unadjusted [95% Conf. Interval]	
score					
hour					
2 vs 1	.75	.3790033	16.0	-.0533703	1.55337
3 vs 1	3.429111	.3989647	16.4	2.584889	4.273333
4 vs 1	6.272103	.3989647	16.4	5.427881	7.116324
3 vs 2	2.679111	.3989647	16.4	1.834889	3.523333
4 vs 2	5.522103	.3989647	16.4	4.677881	6.366324
4 vs 3	2.842991	.4179739	16.7	1.959774	3.726209

When we compare these results with the large-sample results below, we can see that the confidence interval of hour 2 versus hour 1 changes to include 0. Therefore, after adjusting for small-sample

size, we would not reject the hypothesis that the means for hour 1 and hour 2 are equivalent at the 5% significance level.

```
. pwcompare hour
Pairwise comparisons of marginal linear predictions
Margins      : asbalanced
```

		Contrast	Std. Err.	Unadjusted [95% Conf. Interval]	
score	hour				
	2 vs 1	.75	.3790033	.0071672	1.492833
	3 vs 1	3.429111	.3971529	2.650706	4.207516
	4 vs 1	6.272103	.3971529	5.493697	7.050508
	3 vs 2	2.679111	.3971529	1.900706	3.457516
	4 vs 2	5.522103	.3971529	4.743697	6.300508
	4 vs 3	2.842991	.4145085	2.03057	3.655413



Stored results

pwcompare with option small stores the following in r():

```
Matrices
r(L_df)      degrees of freedom for each margin difference
r(M_df)      degrees of freedom for each margin estimate
```

pwcompare with options post and small stores the following in e():

```
Matrices
e(L_df)      degrees of freedom for each margin difference
e(M_df)      degrees of freedom for each margin estimate
```

Methods and formulas

Methods and formulas are presented under the following headings:

- Prediction
- Small-sample inference

Prediction

Following the notation defined throughout [ME] **mixed**, BLUPs of random effects **u** are obtained as

$$\tilde{\mathbf{u}} = \tilde{\mathbf{G}}\mathbf{Z}'\tilde{\mathbf{V}}^{-1}\left(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}\right)$$

where $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{V}}$ are $\mathbf{G}$ and $\mathbf{V} = \mathbf{Z}\mathbf{G}\mathbf{Z}' + \sigma_e^2\mathbf{R}$ with maximum likelihood (ML) or REML estimates of the variance components plugged in. Standard errors for BLUPs are calculated based on the iterative technique of Bates and Pinheiro (1998, sec. 3.3) for estimating the BLUPs themselves. If estimation is done by REML, these standard errors account for uncertainty in the estimate of $\boldsymbol{\beta}$, while for ML the standard errors treat $\boldsymbol{\beta}$ as known. As such, standard errors of REML-based BLUPs will usually be larger.

Fitted values are given by $\mathbf{X}\hat{\boldsymbol{\beta}} + \mathbf{Z}\tilde{\mathbf{u}}$, residuals as $\hat{\boldsymbol{\epsilon}} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{Z}\tilde{\mathbf{u}}$, and standardized residuals as

$$\hat{\boldsymbol{\epsilon}}_* = \hat{\sigma}_{\epsilon}^{-1} \hat{\mathbf{R}}^{-1/2} \hat{\boldsymbol{\epsilon}}$$

If the `relevel(levelvar)` option is specified, fitted values, residuals, and standardized residuals consider only those random-effects terms up to and including level *levelvar* in the model.

For details concerning the calculation of scores, see [Methods and formulas](#) in [ME] **mixed**.

Small-sample inference

For small-sample computations performed when the `small` option is used with `test`, `testparm`, `lincom`, `contrast`, or `pwcompare`, see [Denominator degrees of freedom](#) in [Methods and formulas](#) of [ME] **mixed**.

References

- Bates, D. M., and J. C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical Memorandum BL0112140-980226-01TM*. Murray Hill, NJ: Bell Labs, Lucent Technologies.
<http://ect.bell-labs.com/sl/project/nlme/CompMulti.pdf>.
- Mitchell, M. N. 2015. *Stata for the Behavioral Sciences*. College Station, TX: Stata Press.
- Rabe-Hesketh, S., and A. Skrondal. 2012. *Multilevel and Longitudinal Modeling Using Stata*. 3rd ed. College Station, TX: Stata Press.

Also see

[ME] **mixed** — Multilevel mixed-effects linear regression

[U] **20 Estimation and postestimation commands**

Glossary

ANOVA denominator degrees of freedom (DDF) method. This method uses the traditional ANOVA for computing DDF. According to this method, the DDF for a test of a fixed effect of a given variable depends on whether that variable is also included in any of the random-effects equations. For traditional ANOVA models with balanced designs, this method provides exact sampling distributions of the test statistics. For more complex mixed-effects models or with unbalanced data, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

approximation denominator degrees of freedom (DDF) methods. The Kenward–Roger and Satterthwaite DDF methods are referred to as approximation methods because they approximate the sampling distributions of test statistics using t and F distributions with the DDF specific to the method for complicated mixed-effects models and for simple mixed models with unbalanced data. Also see *exact denominator degrees of freedom (DDF) methods*.

between–within denominator degrees of freedom (DDF) method. See *repeated denominator degrees of freedom (DDF) method*.

BLUPs. BLUPs are best linear unbiased predictions of either random effects or linear combinations of random effects. In linear models containing random effects, these effects are not estimated directly but instead are integrated out of the estimation. Once the fixed effects and variance components have been estimated, you can use these estimates to predict group-specific random effects. These predictions are called BLUPs because they are unbiased and have minimal mean squared errors among all linear functions of the response.

canonical link. Corresponding to each family of distributions in a generalized linear model (GLM) is a canonical link function for which there is a sufficient statistic with the same dimension as the number of parameters in the linear predictor. The use of canonical link functions provides the GLM with desirable statistical properties, especially when the sample size is small.

conditional hazard function. In the context of mixed-effects survival models, the conditional hazard function is the hazard function computed conditionally on the random effects. Even within the same covariate pattern, the conditional hazard function varies among individuals who belong to different random-effects clusters.

conditional hazard ratio. In the context of mixed-effects survival models, the conditional hazard ratio is the ratio of two conditional hazard functions evaluated at different values of the covariates. Unless stated differently, the denominator corresponds to the conditional hazard function at baseline, that is, with all the covariates set to zero.

conditional overdispersion. In a negative binomial mixed-effects model, conditional overdispersion is overdispersion conditional on random effects. Also see *overdispersion*.

containment denominator degrees of freedom (DDF) method. See *ANOVA denominator degrees of freedom (DDF) method*.

continuous-time autoregressive structure. A generalization of the autoregressive structure that allows for unequally spaced and noninteger time values.

covariance structure. In a mixed-effects model, covariance structure refers to the variance–covariance structure of the random effects.

crossed-effects model. A crossed-effects model is a mixed-effects model in which the levels of random effects are not nested. A simple crossed-effects model for cross-sectional time-series data would contain a random effect to control for panel-specific variation and a second random effect

to control for time-specific random variation. Rather than being nested within panel, in this model a random effect due to a given time is the same for all panels.

crossed-random effects. See *crossed-effects model*.

EB. See *empirical Bayes*.

empirical Bayes. In generalized linear mixed-effects models, empirical Bayes refers to the method of prediction of the random effects after the model parameters have been estimated. The empirical Bayes method uses Bayesian principles to obtain the posterior distribution of the random effects, but instead of assuming a prior distribution for the model parameters, the parameters are treated as given.

empirical Bayes mean. See *posterior mean*.

empirical Bayes mode. See *posterior mode*.

error covariance, error covariance structure. Variance–covariance structure of the errors within the *lowest-level group*. For example, if you are modeling random effects for classes nested within schools, then error covariance refers to the variance–covariance structure of the observations within classes, the lowest-level groups. With a slight abuse of the terminology, error covariance is sometimes also referred to as residual covariance or residual error covariance in the literature.

exact denominator degrees of freedom (DDF) methods. Residual, repeated, and ANOVA DDF methods are referred to as exact methods because they provide exact t and F sampling distributions of test statistics for special classes of mixed-effects models—linear regression, repeated-measures designs, and traditional ANOVA models—with balanced data. Also see *approximation denominator degrees of freedom (DDF) methods*.

fixed effects. In the context of multilevel mixed-effects models, fixed effects represent effects that are constant for all groups at any level of nesting. In the ANOVA literature, fixed effects represent the levels of a factor for which the inference is restricted to only the specific levels observed in the study. See also *fixed-effects model* in [XT] *Glossary*.

free parameter. Free parameters are parameters that are not defined by a *linear form*. Free parameters are displayed with a forward slash in front of their names or their equation names.

Gauss–Hermite quadrature. In the context of generalized linear mixed models, Gauss–Hermite quadrature is a method of approximating the integral used in the calculation of the log likelihood. The quadrature locations and weights for individual clusters are fixed during the optimization process.

generalized linear mixed-effects model. A generalized linear mixed-effects model is an extension of a generalized linear model allowing for the inclusion of random deviations (effects).

generalized linear model. The generalized linear model is an estimation framework in which the user specifies a distributional family for the dependent variable and a link function that relates the dependent variable to a linear combination of the regressors. The distribution must be a member of the exponential family of distributions. The generalized linear model encompasses many common models, including linear, probit, and Poisson regression.

GHQ. See *Gauss–Hermite quadrature*.

GLM. See *generalized linear model*.

GLME model. See *generalized linear mixed-effects model*.

GLMM. Generalized linear mixed model. See *generalized linear mixed-effects model*.

hierarchical model. A hierarchical model is one in which successively more narrowly defined groups are nested within larger groups. For example, in a hierarchical model, patients may be nested within doctors who are in turn nested within the hospital at which they practice.

intraclass correlation. In the context of mixed-effects models, intraclass correlation refers to the correlation for pairs of responses at each nested level of the model.

Kenward–Roger denominator degrees of freedom (DDF) method. This method implements the [Kenward and Roger \(1997\)](#) method, which is designed to approximate unknown sampling distributions of test statistics for complex linear mixed-effects models. This method is supported only with restricted maximum-likelihood estimation.

Laplacian approximation. Laplacian approximation is a technique used to approximate definite integrals without resorting to quadrature methods. In the context of mixed-effects models, Laplacian approximation is as a rule faster than quadrature methods at the cost of producing biased parameter estimates of variance components.

Lindstrom–Bates algorithm. An algorithm used by the [linearization method](#).

linear form. A linear combination is what we call a “linear form” as long as you do not refer to its coefficients or any subset of the linear combination anywhere in the expression. Linear forms are beneficial for some nonlinear commands such as [nl](#) because they make derivative computation faster and more accurate. In contrast to [free parameters](#), parameters of a linear form are displayed without forward slashes in the output. Rather, they are displayed as parameters within an equation whose name is the linear combination name. Also see [Linear forms versus linear combinations in \[ME\] menl](#).

linear mixed model. See [linear mixed-effects model](#).

linear mixed-effects model. A linear mixed-effects model is an extension of a linear model allowing for the inclusion of random deviations (effects).

linearization log likelihood. Objective function used by the [linearization method](#) for optimization. This is the log likelihood of the linear mixed-effects model used to approximate the specified nonlinear mixed-effects model.

linearization method, Lindstrom–Bates method. Method developed by [Lindstrom and Bates \(1990\)](#) to approximate for fitting [nonlinear mixed-effects models](#). The linearization method uses a first-order Taylor-series expansion of the specified nonlinear mean function to approximate it with a linear function of fixed and random effects. Thus a nonlinear mixed-effects model is approximated by a [linear mixed-effects model](#), in which the fixed-effects and random-effects design matrices involve derivatives of the nonlinear mean function with respect to fixed effects (coefficients) and random effects, respectively. Also see [Introduction](#) in [\[ME\] menl](#).

link function. In a generalized linear model or a generalized linear mixed-effects model, the link function relates a linear combination of predictors to the expected value of the dependent variable. In a linear regression model, the link function is simply the identity function.

LME model. See [linear mixed-effects model](#).

lowest-level group. The second level of a multilevel model with the observations composing the first level. For example, if you are modeling random effects for classes nested within schools, then classes are the lowest-level groups.

MCAGH. See [mode-curvature adaptive Gauss–Hermite quadrature](#).

mean–variance adaptive Gauss–Hermite quadrature. In the context of generalized linear mixed models, mean–variance adaptive Gauss–Hermite quadrature is a method of approximating the integral used in the calculation of the log likelihood. The quadrature locations and weights for

individual clusters are updated during the optimization process by using the posterior mean and the posterior standard deviation.

mixed model. See *mixed-effects model*.

mixed-effects model. A mixed-effects model contains both fixed and random effects. The fixed effects are estimated directly, whereas the random effects are summarized according to their (co)variances. Mixed-effects models are used primarily to perform estimation and inference on the regression coefficients in the presence of complicated within-subject correlation structures induced by multiple levels of grouping.

mode-curvature adaptive Gauss–Hermite quadrature. In the context of generalized linear mixed models, mode-curvature adaptive Gauss–Hermite quadrature is a method of approximating the integral used in the calculation of the log likelihood. The quadrature locations and weights for individual clusters are updated during the optimization process by using the posterior mode and the standard deviation of the normal density that approximates the log posterior at the mode.

MVAGH. See *mean–variance adaptive Gauss–Hermite quadrature*.

named substitutable expression. A named substitutable expression is a *substitutable expression* defined within `menl`'s `define()` option; see *Substitutable expressions* in [ME] `menl`.

nested random effects. In the context of mixed-effects models, nested random effects refer to the nested grouping factors for the random effects. For example, we may have data on students who are nested in classes that are nested in schools.

NLME model. See *nonlinear mixed-effects model*.

nonlinear mixed-effects model. A model in which the conditional mean function given random effects is a nonlinear function of fixed and random effects. A linear mixed-effects model is a special case of a nonlinear mixed-effects model.

one-level model. A one-level model has no multilevel structure and no random effects. Linear regression is a one-level model.

overdispersion. In count-data models, overdispersion occurs when there is more variation in the data than would be expected if the process were Poisson.

posterior mean. In generalized linear mixed-effects models, posterior mean refers to the predictions of random effects based on the mean of the posterior distribution.

posterior mode. In generalized linear mixed-effects models, posterior mode refers to the predictions of random effects based on the mode of the posterior distribution.

QR decomposition. QR decomposition is an orthogonal-triangular decomposition of an augmented data matrix that speeds up the calculation of the log likelihood; see *Methods and formulas* in [ME] `mixed` for more details.

quadrature. Quadrature is a set of numerical methods to evaluate a definite integral.

random coefficient. In the context of mixed-effects models, a random coefficient is a counterpart to a slope in the fixed-effects equation. You can think of a random coefficient as a randomly varying slope at a specific level of nesting.

random effects. In the context of mixed-effects models, random effects represent effects that may vary from group to group at any level of nesting. In the ANOVA literature, random effects represent the levels of a factor for which the inference can be generalized to the underlying population represented by the levels observed in the study. See also *random-effects model* in [XT] *Glossary*.

random intercept. In the context of mixed-effects models, a random intercept is a counterpart to the intercept in the fixed-effects equation. You can think of a random intercept as a randomly varying intercept at a specific level of nesting.

random-effects substitutable expression. A random-effects substitutable expression is a [substitutable expression](#) containing random-effects terms; see [Random-effects substitutable expressions](#) in [ME] [menl](#).

REML. See [restricted maximum likelihood](#).

repeated denominator degrees of freedom (DDF) method. This method uses the repeated-measures ANOVA for computing DDF. It is used with balanced repeated-measures designs with spherical correlation error structures. It partitions the residual degrees of freedom into the between-subject degrees of freedom and the within-subject degrees of freedom. The repeated method is supported only with two-level models. For more complex mixed-effects models or with unbalanced data, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

residual covariance, residual error covariance. See [error covariance](#).

residual denominator degrees of freedom (DDF) method. This method uses the residual degrees of freedom, $n - \text{rank}(X)$, as the DDF for all tests of fixed effects. For a linear model without random effects with independent and identically distributed errors, the distributions of the test statistics for fixed effects are t or F distributions with the residual DDF. For other mixed-effects models, this method typically leads to poor approximations of the actual sampling distributions of the test statistics.

restricted maximum likelihood. Restricted maximum likelihood is a method of fitting linear mixed-effects models that involves transforming out the fixed effects to focus solely on variance–component estimation.

Satterthwaite denominator degrees of freedom (DDF) method. This method implements a generalization of the [Satterthwaite \(1946\)](#) approximation of the unknown sampling distributions of test statistics for complex linear mixed-effects models. This method is supported only with restricted maximum-likelihood estimation.

substitutable expression. Substitutable expressions are like any other mathematical expressions involving scalars and variables, such as those you would use with Stata's `generate` command, except that the parameters to be estimated are bound in braces. See [Substitutable expressions](#) in [ME] [menl](#).

three-level model. A three-level mixed-effects model has one level of observations and two levels of grouping. Suppose that you have a dataset consisting of patients overseen by doctors at hospitals, and each doctor practices at one hospital. Then a three-level model would contain a set of random effects to control for hospital-specific variation, a second set of random effects to control for doctor-specific random variation within a hospital, and a random-error term to control for patients' random variation.

two-level model. A two-level mixed-effects model has one level of observations and one level of grouping. Suppose that you have a panel dataset consisting of patients at hospitals; a two-level model would contain a set of random effects at the hospital level (the second level) to control for hospital-specific random variation and a random-error term at the observation level (the first level) to control for within-hospital variation.

variance components. In a mixed-effects model, the variance components refer to the variances and covariances of the various random effects.

within-group errors. In a two-level model with observations nested within groups, within-group errors refer to error terms at the observation level. In a higher-level model, they refer to errors within the [lowest-level groups](#).

References

- Kenward, M. G., and J. H. Roger. 1997. Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics* 53: 983–997.
- Lindstrom, M. J., and D. M. Bates. 1990. Nonlinear mixed effects models for repeated measures data. *Biometrics* 46: 673–687.
- Satterthwaite, F. E. 1946. An approximate distribution of estimates of variance components. *Biometrics Bulletin* 2: 110–114.

Subject and author index

See the [combined subject index](#) and the [combined author index](#) in the *Glossary and Index*.