



STATISTICS • VISUALIZATION • DATA MANIPULATION • REPORTING

NEW FEATURES AND COMMANDS

Tables	PyStata
Bayesian econometrics	Jupyter Notebook with Stata
Interval-censored Cox model	Faster Stata
Difference in differences (DID)	Bayesian multilevel modeling
Bayesian VAR	New functions for dates and times
Multivariate meta-analysis	Leave-one-out meta-analysis
Treatment-effects lasso	Galbraith plots
Panel-data multinomial logit	Bayesian panel-data models
Zero-inflated ordered logit	Nonparametric tests for trend
Bayesian IRF and FEVD analysis	Lasso with clustered data
Bayesian dynamic forecasting	BIC for lasso penalty selection
Do-file editor enhancements	Bayesian linear and nonlinear DSGEs
Intel Mat Kernel Library (MKL)	H2O integration
Stata on Apple silicon	Java integration
JDBC	And more

STATA 17 | EDITIONS

Stata/MP | The fastest edition of Stata (for quad-core, dual-core, and multicore/multiprocessor computers) that can analyse the largest datasets

Stata/SE | Standard edition; for larger datasets

Stata/BE | Basic edition; for mid-sized datasets

OPERATING SYSTEMS

Windows* 8.1, 10 e 11, Windows Server* 2022, 2019, 2016 e 2012R2

* 64-bit Windows x86-64 Intel® and AMD processors (Core i3 equivalent or better)

macOS 11.0 (Big Sur) or more recent for **Mac** with Apple Silicon e macOS 10.12 (Sierra) or more recent for Macs with Intel 64-bit processors

Any 64-bit (Core i3 equivalent or better) for **Linux**

Minimum requires include the GNU C library (glibc) version 2.17 or more recent

Check the output of **ldd -v** within a terminal



CONTACT US:

TStat S.r.l. | TStat Training
Stata Distributor for Albania, Bosnia and Herzegovina, Croatia, Greece, Italy, Kosovo, North Macedonia, Malta, Montenegro, Serbia, Slovenia and the Slovak Republic
Stata Official Trainer for Europe

Tel. +39 0864 210101 | TSTAT@TSTAT.IT | WWW.TSTAT.IT | WWW.TSTATTRAINING.EU

TABLES

Users have been asking for better tables. In Stata 17 users can now easily create tables that compare regression results or summary statistics, you can create styles and apply them to any table you build, and you can export your tables to MS Word®, PDF, HTML, LaTeX, MS Excel®, or Markdown and include them in reports. The **table** command is revamped. The new **collect** prefix collects as many results from as many commands as you want, builds tables, exports them to many formats, and more. You can also point-and-click to create tables using the new **Tables Builder**.

BAYESIAN ECONOMETRICS

Stata 17 now allows users to implement Bayesian econometrics. Do you need to use probabilistic statements to answer economic questions, for example, are those who participate in a job-training program more likely to stay employed for the next five years? Would you like to incorporate prior knowledge of an economic process? Stata's new Bayesian econometrics features can help. Fit many Bayesian models such as cross-sectional, panel-data, multilevel, and time-series models. Compare models using Bayes factors. Obtain predictions and forecasts. And more!

One of the appeals for using Bayesian methods in econometric modeling is to incorporate the external information about model parameters often available in practice. This information may come from historical data, or it may come naturally from the knowledge of an economic process. Either way, a Bayesian approach allows us to combine that external information with what we observe in the current data to form a more realistic view of the economic process of interest.

Stata 17 offers several new features in the area of Bayesian econometrics:

Bayesian VAR models

Bayesian IRF and FEVD analysis

Bayesian dynamic forecasting

Bayesian longitudinal/panel-data models

Bayesian linear and nonlinear DSGE models

FASTER DATA

Stata values accuracy and it values speed. There is often a tradeoff between the two, but Stata strives to give users the best of both worlds. In Stata 17, algorithms behind **sort** and **collapse** have been updated to make these commands faster. Speed improvements have also been attained for some estimation commands such as **mixed**, which fits multilevel **mixed-effects** models.

DIFFERENCE-IN-DIFFERENCES (DID) AND DDD MODELS

New estimation commands **didregress** and **xtdidregress** fit difference-in-differences (DID) and difference-in-difference-in-differences or triple-differences (DDD) models with repeated-measures data. **didregress** works with repeated-cross-sectional data, and **xtdidregress** works with longitudinal/panel data.

DID and DDD models are used to estimate the average treatment effect on the treated (ATET) with repeated-measures data. A treatment effect can be an effect of a drug regimen on blood pressure or an effect of a training program on employment. Unlike with the standard cross-sectional analysis, available with the existing **teffects** command, DID analysis controls for group and time effects when estimating the ATET, where groups identify repeated measures. DDD analysis controls for additional group effects and their interactions with time—users can specify up to three group variables or two group variables and a time variable.

BAYESIAN VAR MODELS

The **bayes** prefix now supports the **var** command to fit Bayesian vector autoregressive (VAR) models.

VAR models study relationships between multiple time series by including lags of outcome variables as model predictors. These models are known to have many parameters: with K outcome variables and p lags, there are at least $p(K^2 + \ln 1)$ parameters. Reliable estimation of the model parameters can be challenging, especially with small datasets.

Bayesian VAR models overcome these challenges by incorporating prior information about model parameters to stabilize parameter estimation.

INTERVAL-CENSORED COX MODEL

A semiparametric Cox proportional hazards regression model is commonly used to analyze uncensored and right-censored event-time data. The new estimation command **stintcox** fits the Cox model to interval-censored event-time data.

Interval-censoring occurs when the time to an event of interest, such as recurrence of cancer, is not directly observed but is known to lie within an interval. For example, the recurrence of cancer can be detected between periodic examinations, but the exact time of recurrence cannot be observed. We know only that cancer recurred sometime between the previous and current examinations. Ignoring interval-censoring may lead to incorrect (biased) results.

Semiparametric estimation, when the baseline hazard function is left completely unspecified, of interval-censored event-time data is challenging because none of the event times are observed exactly. As such, "semiparametric" modeling of these data often resorted to using spline methods or piecewise-exponential models for the baseline hazard function. Genuine semiparametric modeling of interval-censored event-time data was not available until recent methodological advances, which are implemented in the **stintcox** command.

GALBRAITH PLOTS

The new command **meta galbraithplot** produces Galbraith plots for a meta-analysis. These plots are useful for assessing heterogeneity of the studies and for detecting potential outliers. They are also used as an alternative to forest plots for summarizing meta-analysis results when there are many studies.

TREATMENT-EFFECTS LASSO ESTIMATION

If you use **teffects** to estimate treatment effects. You may also use **lasso** to control for many (hundreds, thousands, or more!) covariates. In Stata 17, you can now use **telasso** to estimate treatment effects and control for many covariates.

MULTIVARIATE META-ANALYSIS

Users may want to analyze results from multiple studies. The studies report multiple effect sizes, which are likely to be correlated within a study. Separate meta-analyses, such as those using the existing **meta** command, will ignore the correlation. The new meta **mvregress** command enables users to perform multivariate meta-analysis, which will account for the correlation.

BAYESIAN MULTILEVEL MODELS: NONLINEAR, JOINT, SEM-LIKE, AND MORE

Users can fit breadth of Bayesian multilevel models with the new elegant random-effects syntax of the **bayesmh** command. Univariate linear and nonlinear multilevel models can be fitted more easily. Users can now also fit multivariate linear and nonlinear multilevel models! Think of growth linear and nonlinear multilevel models, joint longitudinal and survival-time models, SEM-type models, and more.

LEAVE-ONE-OUT META-ANALYSIS

Users can now perform leave-one-out meta-analysis by using the new option **leaveoneout** with **meta summarize** and **meta forestplot**.

The leave-one-out meta-analysis performs multiple meta-analyses by excluding one study at each analysis. It is common for studies to produce exaggerated effect sizes, which may distort the overall results. The leave-one-out meta-analysis is useful to investigate the influence of each study on the overall effect-size estimate and to identify influential studies. See **[META] meta summarize** and **[META] meta forestplot**.

BAYESIAN LONGITUDINAL/PANEL-DATA MODELS

Users fit random-effects panel-data or longitudinal models by using **xtreg** for continuous outcomes, **xtlogit** or **xtprobit** for binary outcomes, **xtologit** or **xtoprobit** for ordinal outcomes, and more. In Stata 17, one can fit Bayesian versions of these models by simply prefixing them with **bayes**. See the full list of supported commands in **[BAYES] Bayesian estimation** under *Panel-data models*.

PANEL-DATA MULTINOMIAL LOGIT MODEL

Stata's new estimation command **xtmlogit** fits panel-data multinomial logit (MNL) models to categorical outcomes observed over time. Suppose that we have data on choices of restaurants from individuals collected over several weeks. Restaurant choices are categorical outcomes that have no natural ordering, so we could use the existing **mlogit** command (with cluster-robust standard errors). But **xtmlogit** models individual characteristics directly and thus may produce more efficient results. And it can properly account for characteristics that might be correlated with covariates.

BAYESIAN DYNAMIC FORECASTING

Dynamic forecasting is a common prediction tool after fitting multivariate time-series models, such as vector autoregressive (VAR) models. If you use **fcast** to compute dynamic forecasts after fitting a classical **var** model. You can now use **bayesfcast** to compute Bayesian dynamic forecasts after fitting a Bayesian VAR model using **bayes: var**.

Bayesian dynamic forecasts produce an entire sample of predicted values instead of a single prediction as in classical analysis. This sample can be used to answer various modeling questions, such as how well the model predicts future observations without making the asymptotic normality assumption when estimating forecast uncertainty. This is particularly appealing for small datasets for which the asymptotic normality assumption may be suspect.

ZERO-INFLATED ORDERED LOGIT MODEL

The new estimation command **ziologit** fits zero-inflated ordered logistic regression models. This model is used when data exhibit a higher fraction of observations in the lowest category than would be expected from a standard ordered logistic model. We refer to observations in the lowest category as zeros because they typically correspond to the absence of a behavior or trait.

Zero inflation is accounted for by assuming that the zeros come from both a logistic model and an ordered logistic model. Each model can have different covariates, and the results can be displayed as odds ratios instead of the default coefficients. See **[R] ziologit**.

NONPARAMETRIC TESTS FOR TREND

The **np trend** command now supports four tests for trend across ordered groups. You can choose between the Cochran-Armitage test, the Jonckheere-Terpstra test, the linear-by-linear trend test, and the Cuzick test using ranks. The first three tests are new, and the fourth test was performed by **np trend** previously.

BAYESIAN IRF AND FEVD ANALYSIS

Impulse-response functions (IRFs), dynamic-multiplier functions, and forecast-error variance decompositions (FEVDs) are commonly used to describe the results from multivariate time-series models such as VAR models. VAR models have many parameters, which may be difficult to interpret. IRFs and other functions combine the effect of multiple parameters into one summary (per time period). For instance, IRFs measure the effect of a shock (change) in one variable on a given outcome variable. Bayesian IRFs (and other functions) produce the results using the "exact" posterior distribution of IRFs, which does not rely on the assumption of asymptotic normality. They may also provide more stable estimates for small datasets because they incorporate prior information about the model parameters.

BIC FOR LASSO PENALTY SELECTION

Selection of a penalty parameter is fundamental to lasso analysis. Using a small penalty may include too many variables. Using a large penalty may omit potentially important variables.

Lasso estimation already provides several penalty-selection methods, including cross-validation, adaptive, and plugin. Users can now use the Bayesian information criterion (BIC) to select the penalty parameter after lasso for prediction and lasso for inference by specifying the **selection(bic)** option. Also, the new postestimation command **bicplot** plots the BIC values as a function of a penalty parameter after fitting a lasso model. This provides a convenient graphical representation for the value of the penalty parameter that minimizes the BIC function.

LASSO FOR CLUSTERED DATA

Users can now account for clustered data in their lasso analysis. Ignoring clustering may lead to incorrect results in the presence of correlation between observations within the same cluster. With lasso commands for prediction such as **lasso** and **elasticnet**, one can specify the new **cluster({\it clustvar})** option. With lasso commands for inference such as **poregress**, users can specify the new **vce(cluster {\it clustvar})** option.

BAYESIAN LINEAR AND NONLINEAR DSGE MODELS

Bayesian linear and nonlinear dynamic stochastic general equilibrium (DSGE) models can now be fitted in Stata 17 by prefixing **dsge** and **dsgenl** with **bayes:**. Incorporate information about the ranges of model parameters by choosing from over 30 different prior distributions. Perform Bayesian IRF analysis, perform interval hypothesis testing, use Bayes factors to compare models, and more.

NEW FUNCTIONS FOR DATES AND TIMES

Stata 17 has added new convenience functions for handling dates and times in both Stata and Mata. The new functions can be grouped into three categories:

1. Datetime durations: functions that are designed to get durations, such as ages (see **[D] Datetime durations**).
2. Relative dates: functions that return dates based on other dates, such as the next birthday relative to a given date (see **[D] Datetime relative dates**).
3. Datetime components: functions that extract different components from datetime values (see **[D] Datetime**).

The new functions factor in leap years, leap days, and leap seconds, wherever applicable. Leap seconds are one-second adjustments that are occasionally applied to Coordinated Universal Time (UTC).

DO-FILE EDITOR IMPROVEMENTS: NAVIGATION CONTROL, ENHANCED BOOKMARKS ...

The Do-file Editor in Stata 17 has the following improvements:

1. Bookmarks: now saved with do-files.
2. New Navigation Control: navigate do-files easier.
3. Syntax highlighting support now includes Java and XML.
4. Autocompletion of quotes, parentheses, and brackets around a selection. For example, select the text **mymacro** and then type a left quote, `; the Do-file Editor will then bind the text selection with single quotes, changing the selection to **`mymacro'**.

Bookmarks. One of the most requested features for the Do-file Editor is the ability to save bookmarks in a do-file. Bookmarks are used to mark lines of interest to navigate to them later more easily. Bookmarks are particularly useful in navigating long do-files. Stata 17 users can now add bookmarks to sections of their do-files that perform data management, display summary statistics, and perform statistical analysis. Users can then use the menus, the toolbar, or the new Navigation Control to quickly move back and forth between those sections without having to scroll through several lines of code to find the section that they are looking for.

Navigation Control. Stata 17 makes navigating do-files easier with the new Navigation Control, which displays a list of bookmarks and their labels. Selecting a bookmark from the Navigation Control will move the Do-file Editor to the bookmark's line. In addition to bookmarks, the Navigation Control will also display a list of programs within a do-file. Selecting a program from the Navigation Control will move the Do-file Editor to the program's definition. No extra effort is needed to add a program to the Navigation Control. The Do-file Editor will automatically add a program's definition to the Navigation Control.

INTEL MATH KERNEL LIBRARY (MKL)

Stata 17 introduces usage of the Intel Math Kernel Library (MKL) on compatible hardware (all Intel- and AMD-based 64-bit computers), providing deeply optimized LAPACK routines.

LAPACK stands for Linear Algebra PACKage and is a freely available set of routines for solving systems of simultaneous equations, eigenvalue problems, and singular-value problems, among others. Mata operators and functions (such as **qrd()**, **lud()**, and **cholesky()**) leverage LAPACK where possible for many numerical operations.

LAPACK backed by the Intel MKL provides the latest LAPACK routines heavily optimized for the 64-bit Intel x86-64 instruction set used by both modern Intel and modern AMD processors. Mata functions and operators using MKL benefit greatly in terms of performance.

And most importantly, users do not need to do anything to take advantage of the speed gains. Stata commands using these Mata functions and operators, and the Mata functions and operators themselves, will automatically use the Intel MKL on compatible hardware.

JDBC

Connecting Stata with databases has become even easier. Stata 17 adds support for JDBC (Java Database Connectivity). JDBC is a cross-platform standard for exchanging data between programs and databases. The new **jdbc** command supports the JDBC standard for exchanging data with relational databases or nonrelational database-management systems that have rectangular data. Users can import data from some of the most popular database vendors, such as Oracle, MySQL, Amazon Redshift, Snowflake, Microsoft SQL Server, and many more. What is great about **jdbc** is that it is a cross-platform solution, so Stata 17's JDBC setup works the same way for Windows, Mac, and Unix systems.

If your database vendor provides a JDBC driver, you can download and install that driver and then read from, write to, and execute SQL on your database via **jdbc**. You can load an entire database table into Stata or use **SQL SELECT** to just load specific columns from a table into Stata. You can also insert all your variables into a database table or just insert a subset of your dataset.

STATA ON APPLE SILICON

Stata 17 for Mac is available as a universal application that will run natively on both Macs with Apple Silicon and Macs with Intel processors. Macs with Apple Silicon include the new MacBook Air, MacBook Pro, and Mac mini, all with M1 processors. The M1 chip promises greater performance and greater power efficiency. This is noteworthy to Stata-for-Mac users, many of whom use Mac laptops.

Although this first set of M1 Macs is considered to be an entry-level set, StataCorp has found that M1 Macs natively running Stata outperform Intel Macs by 30-35%. They even greatly outperform or perform as well as Intel Macs that cost more than twice as much! And users who insist on only Apple-Silicon-native software on their Apple Silicon Macs, will be happy to know that no part of Stata 17, from the installer to the application itself, requires the use of Rosetta 2.

Stata functions the same way whether you are running Stata natively on an M1 Mac or on an Intel Mac, and no special license is required for the M1 Mac. Intel Mac users should note that we will continue to support and release new versions of Stata for Macs with Intel processors for years to come.

PYSTATA

Stata 17 introduces a concept called PyStata. PyStata is a term that encompasses all the ways Stata and Python can interact.

Stata 16 featured the ability to call Python code from Stata. Stata 17 greatly expands on this by allowing users to invoke Stata from a stand-alone Python environment via a new **pystata** Python package.

Users can access Stata and Mata conveniently in an IPython kernel-based environment (for example, Jupyter Notebook and console and Jupyter Lab and console); in other environments that support the IPython kernel (for example, Spyder IDE and PyCharm IDE); or when accessing Python from a command line (for example, Windows Command Prompt, macOS terminal, Unix terminal, and Python's IDLE).

JAVA INTEGRATION

In Stata 17, users can now embed and execute Java code directly in Stata. Users could already create and use Java plugins in previous versions of Stata, but this required them to compile their code and bundle it in a Jar file. Executing Java in a do-file would give users the freedom to execute Java code tied directly to their Stata code. In Stata 17, users can now write their Java code in do-files or ado-files, or even invoke Java interactively (like JShell) from within Stata.

One of Java's strengths is in its extensive APIs that are packaged with the Java virtual machine. There are also many useful third-party libraries available. Depending on what users need to do, one may even be able to write parallelized code to take advantage of multiple cores.

The Java code users write compiles on the fly—there is no need to use an external compiler! Additionally, the Stata Function Interface (SFI) Java package is included, providing a bidirectional connection between Stata and Java. The SFI package has classes to access Stata's current dataset, frames, macros, scalars, matrices, value labels, characteristics, global Mata matrices, date and time values, and more.

Stata bundles the Java Development Kit (JDK) with its installation, so there is no additional setup involved.

JUPYTER NOTEBOOK WITH STATA

Jupyter Notebook is a powerful and easy-to-use web application that allows users to combine executable code, visualizations, mathematical equations and formulas, narrative text, and other rich media in a single document (a "notebook") for interactive computing and developing. It has been widely used by researchers and scientists to share their ideas and results for collaboration and innovation.

In Stata 17, as part of PyStat, users can invoke Stata and Mata from Jupyter Notebook with the **IPython** (interactive Python) kernel. This means users can combine the capabilities of both Python and Stata in a single environment to make their work easily reproducible and shareable with others.

The invocation of Stata from Jupyter Notebook is driven by the new **pystata** Python package.

H2O INTEGRATION

In Stata 17, developers at Statacorp have been experimenting with connecting to H2O, a scalable and distributed open-source machine-learning and predictive analytics platform. Further information regarding H2O is available at <https://docs.h2o.ai/>.

With the integration of H2O, users can start, connect to, and query an H2O cluster from Stata. In addition, Stata 17 includes a suite of commands to manipulate data (H2O frames) on the cluster. For example, users can create new H2O frames by importing data files or loading Stata's current dataset. They can also split, combine, and query H2O frames from within Stata.

Although this is still in the experimental stage for StataCorp, it has been made available so that Stata users can try it out. Note however, that as it is an experimental feature, syntax and features are subject to change. When using Stata commands that provide access to a given feature of H2O, keep in mind that that is an H2O feature. Though users may be accessing it via a Stata command, what it does is up to H2O and is outside of Stata.



TStat S.r.l. | TStat Training

Stata Distributor for Albania, Bosnia and Herzegovina, Croatia, Greece, Italy, Kosovo, North Macedonia, Malta, Montenegro, Serbia, Slovenia and the Slovak Republic
Stata Official Trainer for Europe

Tel. +39 0864 210101 | TSTAT@TSTAT.IT
WWW.TSTAT.IT | WWW.TSTATTRAINING.EU



Stata is a registered trademark of StataCorp LP, College Station, TX, USA, and the Stata logo is used with the permission of StataCorp.

ADDED SINCE RELEASE DATE

INTERFACE

The Do-file Editor has many new enhancements.

You can now increase the level of indentation for a bookmark's label in the Navigation Control by adding the character # to the bookmark comment. For example, bookmark comment *****# Bookmark 2** will be indented one level more than bookmark comment *****# Bookmark 1**.

Java and Python code blocks are now added to the Navigation Control.

You can now execute only the current line that the cursor is on by clicking **Tools > Execute (do) line**. The cursor will then automatically advance to the next executable line, bypassing empty lines and comments.

The Do-file Editor now supports code folding of Java code blocks.

When you create a bookmark by toggling or clicking in the bookmark margin for a line, the Do-file Editor now automatically adds a default name for this bookmark.

You can now add a bookmark to a Mata or Java code block by clicking in the bookmark margin of a line. The Do-file Editor will add the bookmark comment **//#**, which is a valid comment for both Mata and Java.

You can now select text between bookmarks by clicking **Edit > Select** between bookmarks.

DATA MANAGEMENT

When you are using **jdbc** to load, write, or view data from a database with a Java API and your JDBC driver needs multiple JAR files to work, you can use the new **jarpath()** option to specify the directory where the driver JAR file is stored along with driver dependencies.

When you use **egen** functions **iqr()**, **median()**, and **pctile()**, you can now specify the **autotype** option to automatically create the smallest type of variable (byte, int, long, or double) possible to hold the results.

REPORTING

Publications often require headers and footers that alternate on odd and even pages. **putdocx begin** and **putdocx sectionbreak** now make it easy to create Word documents with content in the headers and footers that varies across pages.

When you need to append multiple documents using **putdocx append** or **putdocx save**, you can now specify which of the appended documents' styles should be used in the final document.

You can now customize a table's titles and notes—specifying the font, color, alignment, and more—when using **putdocx table** or **putpdf table** to add a table to your document.

putdocx table now allows you to specify individual row heights and column widths using the **height()** and **width()** options. Alternatively, you can specify the widths of all columns by providing a matrix of column widths.

BAYESIAN ANALYSIS

The **bayesmh** command has a number of enhancements:

The new prior **mvnscaled()** provides a multivariate normal prior with a scaled covariance matrix. The new distribution can be used to specify a conjugate prior for the regression coefficients of a linear regression model.

Gibbs sampling is now available for the combination of a probit likelihood and a multivariate normal prior for regression coefficients.

Time-series operators are now allowed with independent variables in linear, nonlinear, and multiequation models. This allows you to fit a wide variety of Bayesian time series models such as **Bayesian threshold autoregressive models**.

SURVIVAL ANALYSIS

If you are using Stata/MP and fitting a Cox proportional hazards model for interval-censored survival-time data with 200 or more observations, estimation is now faster.

CUSTOMIZABLE TABLE

The new **etable** command creates a table of estimation results from one or more estimation commands. This table can be customized in many ways—statistics to be reported, numeric and string formats, stars for significance, notes, title, labels, and more. The table can be exported directly to Microsoft Word, HTML, PDF, Microsoft Excel, LaTeX, Markdown, SMCL, or plain text. In most cases, **etable** is used alone to create, customize, and export a table in one step. However, because **etable** creates a collection, the table can be further customized using the **collect** suite of commands.

The **table** command is now more powerful and convenient.

Variables defining the rows, columns, and separate tables can now be string variables. Previously, only numeric variables were allowed in this context.

table now displays results using the formats of numeric variables that used to define the rows, columns, and separate tables and using the formats of factor variables specified in the **statistic()** option.

With the new **zerocounts** option, **table** will now report a 0 instead of leaving a cell empty when a cell count of 0 is encountered.

The **collect** suite of commands has enhancements for customizing and exporting tables.

With the new **collect composite** command, you can now create a single result composed of multiple other results. This means that you can create a result such as a mean and a standard deviation and place it in a single cell of your table rather than two separate cells.

You can now specify that, when a result exceeds a specified minimum or maximum value, some alternative text will be shown in the table. This is particularly useful for reporting p-values as "<0.0001" when they are less than this value.

You can now add customized titles and notes to your tables.

You can now append tables when exporting to SMCL, plain text, Markdown, HTML, and LaTeX documents.

When you create a table that includes stars for significance, you can now request that a note be added to the bottom of the table, indicating the significance level represented by each symbol.

collect export is now faster when exporting big tables with many empty cells to PDF, Microsoft Word, or Microsoft Excel.

You can now specify a color of **nil** or **none** to remove any existing color specification.

When exporting a table to an Excel file, you can now use the **open** option to open the Excel file in memory for modification.

You can now specify styles for exporting tables to LaTeX that determine whether the table is centered horizontally and whether the LaTeX table environment is used.

TREATMENT EFFECTS

The new **estat grangerplot** command can now be used after **didregress** and **xtdidregress** to graphically assess whether treatment effects vary over time.

If you are using **didregress** or **xtdidregress** for difference-in-differences estimation and you are using the wild bootstrap to compute confidence intervals and p-values, you will now obtain results much faster. In addition, you can now perform the wild bootstrap in blocks if you need to improve memory use.

GENERAL STATISTICS

The **churdle** command for fitting hurdle models now has improved numerical precision for linear models. In some cases, models that did not converge previously or for which initial values were not feasible can now be fit.

FUNCTIONS

Four new matrix functions—**vech()**, **vecp()**, **invvech()**, and **invvecp()**—are now available for performing transformations between square matrices and vectors.

META-ANALYSIS

Performing subgroup analysis? You can now obtain prediction intervals for each subgroup's overall effect size when you use **meta forestplot** and **meta summarize** to create forest plots and meta-analysis summaries.

When you create a subgroup forest plot, the within-group significance test for each group-specific overall effect size is now displayed.

When you create forest plots, you now have fine control of the text that is displayed for the heterogeneity statistics, the homogeneity test, the significance test of overall effect size, and the subgroup variations of these statistics.

PROGRAMMING

In the PyStata module, you can now include the **splash** argument in the **init()** function. **splash** was added to control whether to display the splash message when Stata is initialized.

If you are using Mata solvers and inverters, you can now use the **mata set matasolvetol** command to set the scalar factor in the computation of tolerances.

Mata's **solve_tol()** function is now faster. Therefore, the many Mata solvers and inverters that rely on **solve_tol()**, such as **qrsolve()**, **lusolve()**, and **cholsolve()**, are also faster. Likewise, Stata estimation commands (for example, **sem**) that use these functions are now faster.

SPEED

If you are using **didregress** or **xtdidregress** for difference-in-differences estimation and you are using the wild bootstrap to compute confidence intervals and p-values, you will now obtain results much faster.

If you are using Stata/MP and fitting a Cox proportional hazards model for interval-censored event-time data using **stintcox** with 200 or more observations, estimation is now faster.

TO DOWNLOAD THE LATEST FEATURES

If you already have Stata 17, type

```
update
```

```
update all
```

then type

```
help whatsnew
```

to discover the new features added since Stata 17's official release.