The RATS letter

Estima

Volume 32, No. 1 October, 2023

RATS 10.1

We are pleased to announce the release of RATS Version 10.1 for all three platforms (Windows, Macintosh and UNIX). This brings all the versions up to date, using the most modern version of the Scintilla editor that is used for handling the input/output text windows.

Among other changes are substantially improved help files with more detailed descriptions of examples and paper replications. See "Improved Appearance on Help" on page 2.

Graphics Formats

All platforms now support the following formats for export of graphs, both through file operations and the GSAVE instruction: EPS (both landscape and portrait), PDF (same), PNG and JPEG. Windows also includes WMF (Windows meta file) which is a common copy/paste format for Windows applications.

Report file format

10.1 adds a new format for saving and reloading the contents of Report Windows. The `RRFZ` format (RATS Report Format Zipped) is a zipped JSON file which saves all the information that RATS uses to display the report such as the full precision for real numbers (which can be lost in a text file representation of the output). You can reload this to reformat and copy to a spreadsheet or word processor at a later date, or you can reload just to view specific old results for comparisons with new ones.

Specifics about 10.1 on the three platforms are provided in separate sections. 10.1 is free for anyone with a full version 10 license. If you would like to update an older one, go to the update link.

Windows Versions

The Pro version of WinRATS comes with 32 and 64 bit executables. (The 32 bit is included because a few data formats haven't been extended to handle 64 bit architecture.) In the past, the installer has put both of the executables, and their support files, into a single directory. However, as we have added more 3rd party add-ons, we ran into some which were hard-coded to use exactly the same name for the DLLs (dynamic link library) for both architectures—the expectation was that the calling software would be in separate directories for each architecture. As a result, we had to split the directories so the 32 bit executables are in one and the 64 in another.

This will only be an issue if you have your own shortcuts or batch files which look for the 64 bit application in the base "program files" directory. Those will have to be changed to point to the same file, but inside a `\x64` directory. (Any shortcuts created by the installer itself will point to the correct file).

One of those DLLs is for GhostScript, which is used to generate PDF files for graphics. You no longer need to install GhostScript separately.

Macintosh Versions

10.1 is compiled to create a "fat" executable which includes support for the M1 and M2 chips as well as the older Intel chips and has been tested through (at this point) MacOS 14.0 (Sonoma). We don't expect there will be any issues with any updates to the MacOS in the near future. However, the newer compilers come at the cost that the minimal OS they support is 10.14, which is only five years old. Some older Macs can be upgraded to 10.14 (Mojave) or later; most Macs built from 2012 on are compatible.

Beginning with Mac OS 10.15 (Catalina) and continuing through the present, the rules for where applications are to install various components have changed. In particular, the Applications folder is only supposed to have the applications themselves. RATS now installs the examples, procedures and the like into directories within

```
/Users/Shared/RATS
```

You should not try to "migrate" MacRATS software to a new computer as it is likely to put things into the wrong locations. Instead, you need to install off the link from our web site that you were given in the installation email. (If you don't have the link any longer, let us know).

UNIX Versions

The UNIX software has been updated to use GTK3 (which is what the Scintilla editor uses) and uses Cairo for rendering graphs and other windows. It has both a GUI executable (functionally identical to the GUI interfaces on Windows and Macintosh) and a batch executable. It includes both executables and source code, where the latter can be customized—one user added specialized code for accessing their FAME database. If you don't have GTK3 installed (or would prefer not to install it), a "sandboxed" executable which includes it only within the directory for the RATS software is available, so it doesn't interact with any other software.

Sparkle Updater

Version 10.1 now includes an automatic updater for the software using Sparkle™ which is used extensively in the Macintosh software world, and is also available for Windows. This will automatically check our web site for newer builds and give you the option to install now, install later or skip a particular build. If you skip, you can always go back and update later, as there is a **Check for Updates...** operation on the **Help** menu under Windows, and **RATS** menu on the Mac.

This is a big improvement over the separate "WinRATS Updater" application, which required the user to check manually for updates.

The earliest release that includes Sparkle is 10.00f. If you have a version 10 earlier than that, you will first need to install the newer version from our web site using the original installation instructions. After that, the automatic updater will be functional.

Note that this does *not* apply to the network installer. If you are running a network version (10.00f or later), it will let you know that there is a newer version available, but won't permit it to be installed using the updater—we assume that anyone using a network version will need to have someone in the IT department take care of the installation.

Improved Appearance on Help

The most recent help files for RATS are always kept at

https://estima.com/ratshelp

The link within the software (on the **Help** menu) is whatever was current at the time that your particular version of the software was created. However, note that much of what is on the help now are descriptions of examples and procedures, which generally work across different versions of the software.

An example of some of the improvements can be seen at

https://estima.com/ratshelp/argarchsimrpf.html

Many of the pages of examples (and procedures as well) now include detailed descriptions of the various steps taken as part of the analysis. Both the full code and code snippets can easily be copied and pasted out of the help into a program. (Similar code in the PDF documentation can be corrupted by the page layout program making subtle changes to improve appearance).

New/Updated Examples

The following examples are either new, or substantially updated either by improved coding, improved documentation or both. There is a separate story of the substantial changes to the "Diebold-Yilmaz Papers" on page 4.

BOXCOX.RPF
Example of estimation of a model with a Box-Cox transformation on the dependent variable.

CAGAN.RPF
This is an example of estimation of the parameters underlying a small DSGE model (Cagan's model of hyperinflation).

CASSKOOPMANS.RPF
This is an example of a simple dynamic linear model (Cass-Koopmans with production linear in capital) solved using the `DSGE` instruction.

DSGEHISTORY.RPF
This shows how to compute the historical decomposition for a DSGE model that's been fit to data. As in the VAR literature, the historical decomposition computes the part of a series which uses only a particular set of shocks.

EGARCHBOOTSTRAP.RPF
EGARCHSIMULATE.RPF
These are two (very) similar programs which simulate an EGARCH model, either using parametric bootstrap (`EGARCHBOOTSTRAP.RPF`) or random number simulation (`EGARCHSIMULATE.RPF`) for forecasting variance out-of-sample (EGARCH has no closed-form forecasting formula) and/or simulating data.

GARCHBOOT.RPF
This is an example of bootstrapping a GARCH model. It does a Value at Risk (VaR) calculation for a ten-period return for the dollar/yen exchange rate.

GARCHFLUX.RPF
This is an example of the Nyblom fluctuations test (`@FLUX` procedure) for a GARCH model.

GARCHMVSIMULATE.RPF
This is an example of the simulation of a multivariate (DVECH) GARCH process.

GARCHUVFLEX.RPF
This is an example of estimation of a univariate GARCH model with standard variance recursions, but non-standard densities. This uses specialized versions of procedures for the skew-t and skew-GED distributions.

(continued on page 3)

HASBROUCK.RPF
HASBROUCK JOF 1995

Hasbrouck(1995) proposes the calculation of the long-run decomposition of variance for a set of prices for a single security in several markets. The concept of the long-run decomposition of variance can be applied to any VECM, not just one with a single trend. JOF 1995 example is most similar to the published work. The RPF example is an alternative which includes an estimated (rather than assumed) cointegrating vector.

HPFILTER.RPF

This examines the state-space model underlying the Hodrick-Prescott filter using the `DLM` instruction, relaxing the variance ratio used in the HP filter.

MONTESVAR_MH.RPF

This is an example of Monte Carlo integration of a structural VAR (SVAR) using Random Walk Metropolis. `MONTESVAR.RPF` is similar, but uses importance sampling. Importance sampling is generally superior if it can be made to work, however, that typically requires a structural model where the parameters are well-estimated by **`CVMODEL`**, which usually means the model is a restriction of a Cholesky factor, or something close to it. Random Walk Metropolis likely requires more experimentation to come up with a good method for generating candidates, but can much more easily handle likelihood surfaces with odd shapes.

MONTEVECM.RPF

This is an example of Monte Carlo integration for the impulse responses for a VECM (Vector Error Correction Model).

MSVARIANCES.RPF

This does a Markov switching model of the variance of a series of returns. This is a (possible) alternative to a GARCH model, as high volatility periods will persist under the control of the switching model, thus you will see the same type of volatility clustering that you would with a GARCH model.

PANELCAUSE.RPF

This is an example of a test for (Granger) causality in a panel data set allowing for heterogeneity in the coefficients and variances.

TARMODELS.RPF

This is an example of estimation and analysis of a STAR (Smooth Transition Autoregression) model. It analyzes population data on lynx (a small wild cat) which has natural fluctuations which are asymmetrical because the declines are usually sharper than the gains.

VARGARCHSIMULATE.RPF

This shows how to simulate a VAR with GARCH (DVECH) errors. It first simulates a mean zero GARCH error process to form the shocks, then simulates the VAR taking the GARCH shocks as input.

VARMAGARCH.RPF

This is an example of estimation of a multivariate GARCH model with VAR and VARMA mean models.

VECMCAUSE.RPF

This is an example of testing for causality in a VECM. It has separate tests for causality and long-run causality.

VECMGARCH.RPF

This is an example of estimation of a VECM model with a GARCH error process.

VOLATILITYESTIMATES.RPF

This shows various methods of estimating volatility from historical data. This uses the methods from Garman and Klass(1980). Note that these are relatively crude estimates which are designed for treating volatility as an "observable" for further analysis, such as in the Diebold-Yilmaz papers.

New/Updated Procedures

lagppolyroots.src

Produces a table of the inverted roots of the input lag polynomial, showing the modulus and (for complex roots) the period.

rangrid.src

Is a function which takes a matched pair of `VECTORS` with a grid of values and corresponding densities and returns a draw from the density approximated by linear interpolation between the grid points. This is a last resort for doing a draw from a non-standard (univariate) distribution when another method such as the rejection method isn't available.

stepprobit.src

Does a backwards stepwise reduction of a probit model, dropping variables (one at a time) as long as the smallest t-statistic is less than a chosen threshold in absolute value.

vratio.src

Performs a variance ratio test using one of several methods for evaluating the significance. The variance ratio test is a test of one of the implications of a random walk hypothesis for financial time series, that the multi-step returns should increase linearly with the time gap.

Diebold-Yilmaz Papers

In a pair of papers, Diebold and Yilmaz(2009, 2012) proposed calculation of "spillover" indexes by examining the decomposition of variance (in 2009) and a "generalized" decomposition of variance (2012) for a VAR on a large set of series and aggregating them into "to" and "from" measures for flow of variance.

Since we first posted replications for these, we have had quite a few questions about extensions, or about applications to the user's own data. So we have recently updated both the programs themselves and the descriptions of them to answer those questions. There are separate pages for the two papers, but we have also included a separate page which describes the common parts of the analysis.

One of the more basic questions is what the relationship is between the "volatility spillover" measures used in these papers, and the same done using multivariate GARCH models. The simple answeer to that is none, other than using the same terminology. A GARCH model typically takes the period to period returns and posits a connection for volatility both across time and among variables. "Spillover" is determined by examining the parameters of that model for cross-effects between variables. By contrast, the two DY papers take volatility as "data", and analyze it using VAR techniques, defining spillover to mean response of one volatility series to shocks in another. The raw data aren't included in the data for the DY papers, but what they describe is estimation from multiple measures of price as is done in the (recently written)

VolatilityEstimates.rpf

example—they use the maximum, minimum, opening and closing prices to estimate volatility within a period.

There are two advantages of this approach:

1. It can handle a much larger set of series than would typically be possible with a GARCH model. DY(2009) uses 19 global equity markets, while a GARCH model which admits some type of "spillover" such as a BEKK typically is limited to about five variables—models which can handle more than that are generally highly restricted.
2. It uses more information than the typical closing prices (only) used in generating returns.

The big *disadvantage* is that the period by period estimates of volatility are computed independently of each other and are subject to rather considerable error. If you do not have the extra pricing detail (particularly maximum and minimum), the efficiency of using just the close to close returns for estimating volatility is quite low—in fact, if you have a data point where there is no change from the previous day, even if there was some intraday movement, the volatility datum is zero, which will break the analysis in log form.

It's possible to use a multi-period volatility estimate (such as overlapping ten period variances), though it would be recommended that such estimates be done assuming a mean of zero rather than centering around the sample mean. However, that has, by construction, a very high level of persistence (since last period's value will use nine of the ten data points which are used in this period). The VAR analysis can go a long way to unwind that, but it's much more likely that the VAR on that constructed data will have unstable roots.

One other common feature to both DY papers is the use of rolling sample analysis of the spillover measures This is subject to the overinterpretation problem that is shown in the

RollingCausality.rpf

example—where sampling error makes some windows seem to show "causality" where there is none. Aside from that, a common error is to use rolling windows which are too small for the model of interest.

In both papers, DY chose to exclude from the graphs of rolling window statistics any range which had unstable roots. One of the changes we made was to make that optional. However, even a modestly unstable root will, over ten periods, tend to dominate the results, causing outlier points in the graphs.

One other change we made was to the graphs done in the DY2012 paper. That worked with a much smaller set of series (4 rather than 19), which allowed for greater detail in examining pairs of series. (There are only 6 pairs vs 171.) Extending that from 4 to 5 or 6 is possible, but the original graphs were hard-coded for a specific number of series. The newer program figures out the optimal way to handle a different number of series:

```
compute rows=fix(sqrt(%nvar))
compute cols=(%nvar-1)/rows+1
spgraph(vfields=rows,hfields=cols,$...
```

This favors adding columns before rows, which generally produces a better appearance with time series graphs. So a matrix of 5 graphs will be 2 x 3 with the bottom right cell empty.

2717 Harrison St
Evanston, IL 60201 USA
sales@estima.com
1-847-864-8772