The RATSletter

RATS Version 8.3 Now Available

With RATS version 8.3, we've added several important new tools to make you more productive. We are now using a public domain "programmer's editor" called Scintilla. We've used a proprietary text editor since WinRATS was introduced, as the text editor supplied with Windows at the time was limited to 32K bytes, which could be completely filled with one set of impulse responses on a reasonably large model. Unfortunately, that editor was orphaned several years ago, and wasn't portable to the Macintosh or UNIX.

Scintilla is portable to all three platforms and is specifically designed as a programmer's editor, which is what is needed for RATS. One very useful feature is a "bracket matcher", which shows the (apparent) matches for ( ) and { } pairs (showing the mates in green), to help fix parenthesis nesting errors. If one of these characters has no obvious partner, it is shown in red as you cursor over it. We've also added the ability to "mark" a location in a text file so you can easily return to it.

A big productivity aid is the new *Find in Files* operation. Among standard, textbook and paper replication examples, we have well over 1000 running examples. With *Find in Files*, you can search across these for particular text, like `MV=DCC` to find examples which do DCC GARCH models. Both *Find in Files*, and the standard *Edit-Find* operations also now support "regular expressions", a standard set of codings for text patterns. For instance, a search pattern of `sur(.*model=` will locate all instances of a **SUR** instruction with the `MODEL` option, with `.*` matching any number of characters. A simple `sur(model=` will only find examples where `MODEL` is the first option listed.

RATS 8.3 now uses multiple "threads" with the editor interface in one and the statistical engine in another. This makes it possible for the editor part to still be responsive while calculations are going on. It also slightly improves performance, particularly on computers with multiple cores, since it can give over one entire core to the numerical calculations while running the interface (and anything else you're doing) on another.

The 64-bit executable (on WinRATS Pro) now runs about 15-20% faster than the 32-bit on the same computer due to improvements in compiler technology. We still provide both executables, since some of the added data formats aren't available for

ARCH/GARCH Course Materials

The workbook, example programs, and procedures from our recent e-course on *ARCH/GARCH and Stochasic Volatility Models* are now available for sale. The price for the course materials is $50. The package includes the 291-page PDF workbook, thirty example programs plus programs replicating examples from several textbooks.

This examines the practical and theoretical issues with estimating and using ARCH, GARCH, and stochastic volatility models. It covers many variants of univariate and multivariate GARCH estimation, including specification of the variance and mean models, overcoming numerical issues, and diagnostic checking.

Among the papers covered in detail are

Bauwens and Laurent(2005), "A New Class of Multivariate Skew Densities, With Application to Generalized Autoregressive Conditional Heteroscedasticity Models," *Journal of Business and Economic Statistics*, vol 23, pp 346-354

Elder and Serletis(2010), "Oil Price Uncertainty," *Journal of Money, Credit and Banking*, vol 42, no. 6, 1137–1159

Gospodinov(2005), "Testing For Threshold Nonlinearity in Short-Term Interest Rates," *Journal of Financial Econometrics*, vol 3, no 3, 344–371.

Hafner and Herwartz(2006), "Volatility impulse responses for multivariate GARCH models: An exchange rate illustration", *Journal of International Money and Finance*, vol 25, no 5, 719-740.

West and Cho(1995), "The predictive ability of several models of exchange rate volatility," *Journal of Econometrics*, vol 69, no 2, 367–391.

The course also explains the workings of the Gibbs sampling and bootstrap examples that are provided with RATS, both for inference on the model parameters and for inference on out-of-sample behavior (such as calculation of Value at Risk).

This is the sixth in our series of e-courses, with the other five being on *Bayesian Econometrics*, *Vector Autoregressions*, *State-Space and DSGE Models*, *Structural Breaks and Switching Models* and *Panel Data*. For more information on any of the courses, see:

www.estima.com/courses_completed.shtml

New and Updated Program and Procedures

The files below contain example programs and (in most cases) data files replicating results from significant econometrics papers or demonstrating other useful techniques. These are available on our website, and will be included with version 8.3 of RATS. Several of these are described in considerable detail in the *ARCH/GARCH course* (page one)

bauwens_laurent_jbes2005.zip
Replication file for Bauwens and Laurent(2005), "A New Class of Multivariate Skew Densities, With Application to Generalized Autoregressive Conditional Heteroscedasticity Models," *Journal of Business and Economic Statistics*, vol 23, pp 346-354. This estimates a DCC GARCH model with errors having the multivariate skew-*t* density proposed in the paper.

camacho_jel2011.zip
Replication file for Camacho(2011), "Markov-switching models and the unit root hypothesis in real U.S. GDP", *Economics Letters*, vol. 112, 161-164. This tests for a unit root allowing for a Markov-switching intercept. It demonstrates estimation of a Markov-switching univariate regression with both fixed and switching coefficients and the simulation of a Markov-switching autoregression.

cushman_zha_jme1997.zip
Replication file for Cushman and Zha(1997), "Identifying monetary policy in a small open economy under flexible exchange rates," *Journal of Monetary Economics*, vol 39, 433-448. This analyzes a "near-SVAR" using a specialized Gibbs sampler designed for large VAR's with an exogenous block.

fry_pagan_jel2010.zip
Example of calculation of Fry-Pagan (2011) median target estimator. "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, vol. 49, no. 4, 938-60.

koutmos_jfba1996.zip
Replication file for the multivariate EGARCH model with mean and asymmetric volatility spillovers from Koutmos(1996) "Modeling the Dynamic Interdependence of Major European Stock Markets", *Journal of Business Finance and Accounting*, vol 23, 975-988.

l1trend.zip
Example from Kim, Koh, Boyd, Gorinevsky(2009), "L1 Filtering", *SIAM Review*, vol 51, pp 339-360, which computes a "robust" estimate of an H-P filter.

These examples have had significant improvements. Many of these are explained in detail as part of the *ARCH/GARCH course* (see page 1).

ckls_jof1992.zip
Chan, Karolyi, Longstaff and Sanders(1992), "An Empirical Comparison of Models of the Short-Term Interest Rate", Journal of Finance, vol 47, no 3, 1209-1227. This estimates models of the interest rate using GMM.

ehrmann_ellison_valla_el2003.zip
Ehrmann, Ellison, Valla (2003), "Regime-dependent impulse response functions in a Markov-switching vector autoregression model", *Economics Letters*, Vol. 78, pp. 295-299.

elderserletis_jmcb2010.zip
Elder and Serletis(2010), "Oil Price Uncertainty", *Journal of Money, Credit and Banking*, Vol. 42, 1137-1159.

hamilton_susmel_joe1994.zip
Hamilton and Susmel(1994), "Autoregressive Conditional Heteroskedasticity and Changes in Regime," *Journal of Econometrics*, vol 64, pp 307-333.

garchgibbs.rpf
garchimport.rpf
These are two examples of MCMC methods applied to a univariate GARCH model. Compared with earlier versions, both have been optimized and we've added quite a few comments. The Elder-Serletis replication (above) is an example of MCMC applied to a *multivariate* GARCH model.

montesur.rpf
This estimates a SUR model using Gibbs sampling techniques. This can be used for SUR models of relatively modest size (perhaps 200-300 parameters total in the regressions). The SUR model in the Cushman-Zha paper (column to the left) is more like 1100 parameters, and so needs a more computationally efficient technique to be get the time required down to a manageable level.

Some important procedures have been improved with 8.3. In particular, the Markov-switching support procedures now include calculation of standardized residuals, as described in the story on page 3. The three main "post-processors" for analyzing Monte Carlo draws in Vector Autoregressions (**@MCFEVDTABLE**, **@MCGRAPHIRF** and **@MCPROCESSIRF)** now have additional options for formatting and handling output.

@REGSTRTEST is a new procedure for testing for smooth transition effects in a general regression, as the existing **@STARTEST** does for an autoregression specifically.

Markov-Switching Model Residuals

We've had quite a few users who wanted to compute residuals from their Markov-switching model for use in diagnostics. It's not a simple matter to come up with a definition of "residual" that has desirable properties. The most obvious try (and the most common request) would be the regime-specific residuals; that is, the model residuals across time for a specific regime. The problem with this is that the switching model has a predictive density for y_t under regime i only if the switching process is *in* regime i at t—it says nothing about what happens if the process *isn't* in regime i. To take a straightforward case, suppose that we have a simple model with a Markov switching mean, where regime 1 is "low" and regime 2 is "high". Where the process is actually in regime 1, the regime-specific residuals using the regime 2 predictions will be systematically negative, and where the process is actually in regime 2, the residuals computed from regime 1 will be systematically positive. Thus, if we tried to do any type of diagnostics on the regime-specific residuals across time, we would expect that the mean would be non-zero, they would be serially correlated, and not Normal, so there would be no useful diagnostics.

Since the regime isn't known, the prediction for y_t will be a probability-weighted combination of the predictions under the regimes. In order to get residuals which are serially uncorrelated (if the model is correct), we need to use predictive probabilities; that is, the probabilities of the regimes at t given data through t-1. For concreteness, assume that we have the regime models as:

$$y_t \sim N\left(x_t\beta_i, \sigma_i^2\right)$$

Given probability weights (at t) of p_i, the mean of y_t is

$$\sum_i p_i x_t \beta_i$$

so the expected value of y_t minus this is zero. The variance is a bit more complicated, as, while expected values can be computed using probability weighting, variances can't. Instead, we have compute the variance using $\mathrm{var}(y_t) = E y_t^2 - \left(E y_t\right)^2$. Since

$$E y_t^2 = \sum_i p_i \left(\sigma_i^2 + \left(x_t\beta_i\right)^2\right)$$

the variance is

$$\hat{\sigma}_t^2 = \sum_i p_i \left(\sigma_i^2 + \left(x_t\beta_i\right)^2\right) - \left(\sum_i p_i x_t\beta_i\right)^2$$

Note that this does *not* simplify noticeably if the variances are identical across regimes—if the re-

gime is unclear (none of the p are near one), the major driver of the predictive variance is likely to be the difference in the means across regimes. Thus we get the standardized predictive residual as:

$$\hat{e}_t = (y_t - \hat{y}_t) / \hat{\sigma}_t$$

which will be (if the model is correct) mean zero, variance one and serially uncorrelated, so you can use a Ljung-Box or West-Cho test for serial correlation, or test for constancy of variance. One thing which these *won't* be is Normal—even if the underlying assumption of Normality is correct, these have a mixture distribution and won't be Normal.

In a multivariate model, the basic calculations are as above (with outer products replacing squares). Again, the predictive covariance matrix is time-varying, even if the model assumes the covariance is fixed across regimes. In order to provide useful diagnostics, the predictive errors at t need to be standardized by being pre-multiplied by the inverse of a factor of the predictive covariance matrix at t. Each possible factor scheme will produce a (somewhat) different set of standardized residuals; what's built into the new versions of the **@MSVAR-Setup** and **@MSSysRegression** procedures is the Choleski factor.

The procedures to generate the standardized residuals in the updated versions of the Markov-switching support procedures are

@MSRegStdResiduals *resids start end*

if you're using **@MSRegression** for a single equation,

@MSSysRegStdResids *start end resids*

if you're using **@MSSysRegression** and

@MSVARStdResiduals *start end resids*

if you're using **@MSVARSetup**. For the last two, *resids* is a VECTOR[SERIES] which will be mean zero, serially uncorrelated with an identity covariance matrix under the null that the model is correct.

If you want the regime-specific residuals (for some purpose other than direct diagnostics), we've altered the "residuals" procedures within each of these to let you choose to evaluate for a specifi regime. These were designed to be used in MCMC estimators, so (by default) they use the current set of sampled regimes. For the **@MSVARSetup** procedures, the calculation can be done with

@MSVARResids(specific=reg) *v start end*

with *reg* as the specific regime (1 to the number of regimes) for which you want residuals. The procedures to compute these are similar for the other procedure sets.

RATS Version 8.3, contd.

the 64-bit at this point. The install process for Win-RATS Pro has been simplified, with the download-able version now including both 32 and 64 bit.

The following changes have been made to existing instructions:

BOXJENK has a MEANEQ option to define the equation for the mean (only) in a RegARIMA model

SSTATS adds a START option similar to those on **SET** and **MAXIMIZE**. It also will now ignore entries which evaluate to NA so you don't have to use a SMPL option to skip over them.

THEIL now has a FORECASTS option which allows you to input a set of forecasts generated by some method other than the standard techniques available, while still being able to let **THEIL** do the forecast error calculations.

The STACKED bar graph styles on **GRAPH** will now allow a mix of positive and negative values at a given time period, stacking away from zero in each direction.

The VTICKS and HTICKS options on the graphics instructions will now accept a VECTOR of values at which you want tick marks, allowing you to override the automatic choices made by RATS.

8.3 adds several new functions:

%MODELLARGESTROOT(model) returns the absolute value of the largest root of the companion matrix of a model. This is helpful in doing Monte Carlo simulations on a VAR, as you can easily reject draws that are unstable.

%STRESCAPE(s,chars,escape) returns a string with any character from chars "escaped" by prefixing with it the character or string escape.

%STRFIND(s,f) looks for the string f in s, returning the position if found and zero otherwise

%STRTRIM(s) returns a copy of a string with leading and trailing quotes and spaces removed

In addition, quite a few existing functions have been altered to allow them to apply to each element of a vector or matrix rather than just to single real values: COS, SIN, TAN, %PLUS and %MINUS. %VALID, when applied to a matrix will return 1 only if no element of the matrix is missing.

Prices and Ordering

Users with current update subscriptions will automatically receive version 8.3 on CD or by download. Users with single-user licenses for versions 8.0 to 8.2 who don't have update subscriptions can update to 8.3 for just $25. CD and download updates include updated documentation in PDF form.

If you have version 7.3 or older, you can update a single-user license for just $150. You don't have to have the older software still installed in order to get an update. You *do* need to have the serial number. If you cannot find your serial number, contact us at *sales@estima.com* or 800-822-8038 before placing your order. You can get the update either with or without printed manuals—the latter avoids shipping charges. The software ships with PDF versions of all of the manuals.

If you have the standard version of RATS, you can upgrade to the Professional version for an additional $150. The Pro version gives you the X12 seasonal adjustment capability, the ability to access the online FRED database, as well as support for reading ODBC/SQL, FAME, and CRSP databases.

The prices for updating a single-user license for Windows or Macintosh are shown below. Please contact us if you have questions, or need to update multi-user or UNIX/Linux licenses.

What you have:	What you want: 8.3 Std	8.3 Pro
RATS Pro v. 8.0-8.2	—	$25
RATS Standard v. 8.0-8.2	$25	$150
RATS Pro (7.3 or earlier)	—	$150
RATS Std. (7.3 or earlier)	$150	$300

Update Subscriptions

Update subscriptions provide a convenient way to budget your software expenditures, and are the easiest way to ensure that you receive all updates to RATS as soon as they are released. We offer two forms of update subscriptions: Version and Annual.

Most (current) multi-user licenses are on annual subscriptions. These provide all updates for one year. The expiration date is shown on the CD envelope label and license agreement. If you (think you) need to renew yours, please contact us for pricing.

Single-user licenses are generally handled using Version subscriptions, which provide all updates up to and including the next major release.

If you have version 8 or are updating to 8.3, a version update subscription is just $150 (plus shipping for customers outside the U.S.) and would give you all updates up to and including version 9.0.