

Unit Roots and Cointegration E-Course

Registration is now open for our newest e-course, which will cover *Unit Roots and Cointegration*. The course will run for roughly eight weeks beginning May 2 and will examine the practical and theoretical issues regarding “unit root” behavior of data and its importance for modern macroeconometric analysis, and its effects on other types of models, such as GARCH and related models. We will look at the effects on testing procedures of allowance for structural breaks. We will also cover the related (and generally more technically demanding) long-memory methods of fractional integration and differencing.

Installments of the course will be posted on a private part of the RATS forum roughly once per week. Participants can work through them at their own pace, make suggestions and ask questions. You will also receive the finished workbook and all future editions of it. Registration is \$100. To order on-line, go to

<https://estima.com/shopcart/courses.shtml>

The new course is the first product on that page. The others are the existing course materials (developed from similar forum-based courses done in the past).

RATS Version 10

If you haven't yet updated to Version 10, you can check <https://estima.com/ratslatest.shtml>

to see what you're missing. Any single-user license version 9.2 or earlier can be updated to version 10 for just \$150. (If you have a multiple-user license, contact us, or your reseller, for pricing). You don't need 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 before placing your order. You can get the update either with or without printed manuals—the latter avoids shipping charges. In either case, the software ships with PDF versions of all of the manuals.

If you would like to switch from Windows to Macintosh (or vice versa), there's no additional charge when you are doing a major update like this.

The direct link for ordering a single user update is

https://estima.com/shopcart/rats_update.shtml

Conditional Forecasting with Restricted Shocks

In a current working paper, Antolin-Diaz, Petrella and Rubio-Ramirez apply conditional forecasting to a Vector Autoregression, but need to restrict which shocks can be used to accomplish the post-sample restrictions. In a standard implementation of conditional forecasting, the most likely set of shocks may not be very useful from a policy standpoint—for instance, to hit optimistic goals for GDP into the future, in most cases, the most likely path is to have large positive shocks to GDP in the first few periods.

There is nothing particularly complicated about this—the stock `@CONDITION` procedure can be used. All you need to do is feed it a “factor” matrix which isn't full rank and only includes columns for the shocks that you want to include. For instance, in the `CONDITION.RPF` example, to use only money and interest rate shocks, this first does a Cholesky factor with M (#3) and R (#5) ordered first

```
compute mrfirst=%psdfactor(%sigma,$
    |13,5,1,2,4,6|)
```

then creates a matrix which zeros out all but those two columns:

```
dec rect msectoronly(%nvar,%nvar)
ewise msectoronly(i,j)=$
    %if(j==3.or.j==5,mrfirst(i,j),0.0)
```

Adding `FACTOR=MSECTORONLY` to the `@CONDITION` procedure call will solve the conditional forecast with the desired restrictions. Note that (orthogonalized) shocks to M and R will (as this is written) have contemporaneous effects on the other variables because the Cholesky factor is ordered with those before the other four variables—as the authors point out, this type of exercise depends upon an identification of the shocks, while a standard conditional forecast (what they call “conditional-on-observables” forecasting) doesn't.

You need to be careful in using this that you don't have too many conditions for the number of shocks that you permit. If, for instance, you are only allowing one shock, you can have no more than one condition per horizon (or more precisely, the count of restrictions through every horizon can't exceed the number of steps—you could, for instance, impose two restrictions at two steps out as long as you have no restrictions at one step).

Antolin-Diaz, J & Petrella, I. & Rubio-Ramírez, J (2018), “Structural Scenario Analysis with SVARs,” CEPR Discussion Papers 12579, C.E.P.R. Discussion Papers.

New and Updated Programs

The following two papers have been available before, but they have been substantially revised and simplified, and now have detailed descriptions as part of the on-line help.

Sadorsky(2012) “Correlations and volatility spillovers between oil prices and the stock prices of clean energy and technology companies”, *Energy Economics*, vol 34, pp 248-255 analyzes a three variable VAR-GARCH model with (returns on) oil, “clean” energy stocks and general technology stocks. Using a DCC-VARMA model for the variance, it computes time-varying hedge ratios and portfolio weights. Note that the portfolio calculations are (assuming mean zero returns) computed the same way regardless of the multivariate GARCH model that was used to estimate the conditional covariance matrices. These are presented both graphically, and as summary statistics across the sample (again, these are time-varying with the period-to-period estimates of the covariance matrices out of the GARCH model).

Chan and Maheu(2002), “Conditional Jump Dynamics in Stock Market Returns”, *Journal of Business and Economic Statistics*, vol 20, no. 3, 377-389, estimates GARCH models with added Poisson jump processes with either fixed or “ARJI” Poisson probabilities. The ARJI-GARCH model was introduced in this paper to allow the jump probabilities to be time-varying using an ARMA-like model. The program has been available since 2012, but this is the first time that we’ve added any technical description other than the comments.

The paper uses an irregular daily data set from 1928 to 1984; we’ve rewritten the example files to use the mapped date scheme added with RATS version 10 (see story on page 4.)

Using %EQNxxx Functions in Non-linear models

Quite a few fairly complicated non-linear models (specialty GARCH models, state-space models, threshold models) are based upon relatively simple linear mean models. However, hard-coding a specific mean model can make the task of adapting to a different data set with a different mean model more complicated. For instance, the residual formula for an AR(1) for variable R can be written

```
nonlin(parmset=meanparms) a0 a1
frml residf = r-(a0+a1*r{1})
```

or as

```
frml meanf = a0+a1*r{1}
```

where r-meanf is used later to compute the residual. However, an AR(1) might not be the appropriate mean model for a different series and you shouldn’t simply adopt the mean model that was used in a published paper without at least checking whether it makes sense in your application.

A more flexible (and actually faster) way to handle this is to use a linear equation (which can be set up with either a **LINREG** with **DEFINE**, or with an **EQUATION**) and then use either %EQNVALUE or %EQNRVALUE to evaluate the value of the mean model or the value of the residual for a given entry and setting for the parameter VECTOR. For instance, the following will estimate a linear equation (here a 2 lag autoregression) by least squares (which is often the best way to give guess values), defines MEANEQ from that, and then defines the VECTOR MU for the coefficients and creates a PARMSET whose only element is that VECTOR.

```
linreg(define=meaneq) dm
# constant dm{1 2}
frml(equation=meaneq,vector=mu,$
      parmset=meanparms) meanf
```

This has the additional benefit of creating labels for the estimation output which include the regressor information, such as MU (2) =DM{1 }

For a log likelihood that needs a calculation based upon the residual from that equation using parameters in MU, you can use %EQNRVALUE as is shown in this:

```
frml logl = u=%eqnrvalue(meaneq,t,mu), $
           (uu(t)=u^2), (h(t)=varf(t)), $
           %logdensity(h,u)
```

%EQNRVALUE takes three parameters: the EQUATION, the entry at which it is to be evaluated (almost always T, when used inside a FRML) and the coefficient VECTOR at which the equation is to be evaluated. The EQUATION keeps track of the dependent variable (y) and the explanatory variables (X), so given the time period t and coefficient VECTOR θ , %EQNRVALUE computes the residual $y_t - X_t\theta$. To apply this to a different mean model, all you need to do is to change the supplementary card on the **LINREG**.

The companion function is %EQNVALUE, which evaluates just the right-hand side of the EQUATION. That is particularly useful in state-space models, where it would most commonly be used in the Z or MU option on **DLM**.

Spillover/Causality-in-Variance Tests

A common topic for people writing a master's thesis or undergraduate paper is to study the "spillover" between markets A and B using GARCH models. While seemingly a relatively straightforward case of "estimate and test some coefficients for zero", it is, in fact, much less clear whether the results from this are of real value. There are several reasons for this:

- a. Volatility isn't observable.
- b. Unlike causality in the mean, where (Granger) "causality" means that a series is forecast *better* with the extra information, causality in variance means that the variance forecasts are merely *different*. (Because of (a), you can't really tell whether they are better.)
- c. GARCH models are all approximations, and models with similar fits can produce very different results.

Cheung and Ng(1996) tried to avoid some of these problems by proposing a test based upon the results of *univariate* GARCH models, testing cross-correlations of squared univariate standardized residuals. That fails as a useful strategy for several reasons. First, the squared residuals are known to be a poor proxy for volatility. Second, this is basically an application of the Pierce-Haugh(1977) test for Granger causality which was very quickly discarded as it produced unreliable results.

The typical parametric method for testing causality-in-variance is to use a multivariate GARCH model in which some parameters create cross-variable effects on future volatility estimates, and test whether those are zero. The most common choice for this is the BEKK model, where the off-diagonal elements would have that type of effect. (Given the standard parameterization, A(1,2) would create a cross-effect from series 1 to series 2.) However, a BEKK often produces a similar fit to a DVECH model, and the DVECH model (by construction) permits no "spillover" as the volatility of each series depends only upon its own past.

To look at the effect this can have, we'll look at the model used in Hafner and Herwartz(2006). (Note that H&H do not do formal tests for variance-in-causality, but look at that through the volatility

impulse responses). The BEKK model used in the paper has 15 free parameters (4 in the mean model and 11 in the GARCH part). The log likelihood is 28606.8. A Wald test for volatility spillover from series 2 (British pound) to series 1 (Deutsche Mark) is 11.5, which, for a χ^2 with 2 degrees of freedom, has a p -value of .0032. The test in the opposite direction is 69.4, which has a p -value which is 0 to as many decimals as you might want to show. (These tests can be set up fairly easily using the *Statistics—Regression Tests* wizard, as it's just *Exclusion Restrictions* on the proper elements of the BEKK parameters, such as jointly on $\mathbf{A}(1, 2)$ and $\mathbf{B}(1, 2)$ for the 1 to 2 direction). From this, one would conclude that there is apparently fairly strong evidence of spillover, particularly from Germany to the UK. However, the DVECH model (which again, allows no possibility of spillover), has 13 free parameters (here 9 in the GARCH part) with a log likelihood of 28610.8. If we were picking between the two models using AIC, they would be tied to five decimals. With the more stringent SBC, the smaller DVECH model is strongly preferred (estimates use 3718 observations, so the SBC penalty for the larger BEKK model is 16.4 with a log likelihood difference of only 4). The difference in inference regarding causality between the two models is due to a combination of (b) and (c)—the BEKK estimates are different, but apparently not really better. Clearly one has to be careful in picking a multivariate GARCH model based solely upon whether it admits a test for "spillover".

Cheung, Y. W. and L.K. Ng, (1996), "A causality-in-variance test and its application to Financial market prices," *Journal of Econometrics*, vol 72, no 1-2, pp 33-48.

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

Pierce, D.A., & L.D. Haugh (1977), "Causality in temporal systems: Characterization and a survey". *Journal of Econometrics*, vol 5, no 3, pp 265-293.

Davies (1987) Likelihood Ratio Test

Various types of switching models face a technical problem in testing for their very existence (vs a simpler linear model) because the more complicated model has one or more parameters which are unidentified under the null (that the second regime is unnecessary). This violates a key assumption in the theorem that gives rise to an asymptotic chi-squared distribution for the likelihood ratio test.

Davies(1987) proposes an adjustment to the critical value of the likelihood ratio test to correct for this, under certain conditions. Garcia and Perron(1996) employed this for the case of testing for a Markov Switching model vs a linear model and (citing G&P) it has been used by other authors in later papers on the subject. However, Garcia and Perron isn't a proper use of Davies. First of all, Davies allows for only a single parameter that's unidentified under the null, while a two-regime Markov switching has two. (Davies adjustment is based upon approximating an integral over that single parameter). It also only applies to a least squares model where the "regime" is determined (exactly) based upon that one parameter, while the Markov Switching model only provides a probability of each regime.

Hansen(1996) actually looks at the behavior of the Davies adjustment when applied to the simpler case of the threshold autoregression (which *does* meet the two requirements described above) and shows (analytically) that the likelihood ratio statistic in that case fails two other key assumptions required by Davies, and that (on simulation) it fails rather badly to provide a good approximation.

Formal testing procedures for the number of regimes (and, in particular, for one regime vs two) are generally quite complicated, involving simulations. Frühwirth-Schnatter(2006) cites a number of papers which have chosen the relatively simple Schwarz Bayesian Criterion (which requires only estimating both models). While the SBC has similar problems with the lack of identification under the null (in deriving the result that it gives an "asymptotic" estimate of the posterior odds), it seems to work adequately in practice.

Davies, R.(1987), "Hypothesis testing when a nuisance parameter is present only under the alternative", *Biometrika*, vol 74, no 1, pp 33–43.

Frühwirth-Schnatter, S. (2006), *Finite Mixture and Markov Switching Models*. New York:Springer.

Garcia, R. and P. Perron(1996), "An Analysis of the Real Interest Rate under Regime Shifts," *Review of*

Economics and Statistics, vol. 78, no 1, pp 111-125.

Hansen, B.E. (1996), "Inference When a Nuisance Parameter is Not Identified Under the Null Hypothesis", *Econometrica*, vol 64, pp 413-430.

Mapped Dates

The mapped date scheme was added with version 10 to allow for use of date fields, for both input and output, for data files with (typically) daily data with gaps for non-trading days. For example, the data file in Chan and Maheu(2002) is daily from 1928 to 1999 with non-trading days removed. To demonstrate the point that a simple "jump" GARCH model has a problem with stability over that range, they estimate a model over three subranges (1928-1950, 1951-1969 and 1970-1984) so it's necessary to be able to (easily) locate the beginning and end of each year.

The original data set has YEAR, MONTH and DAY columns, which is one of three ways to transmit date information (the others being numeric codings like 19281029 or appropriately coded date columns). The **CALENDAR** instruction to create the mapped date scheme will here be:

```
cal (julian=%julianfromymd (year , month , day ) )
```

To pick out the last entry in 1950, the simplest method is to take one less than the first entry in 1951:

```
compute end1950=1951:1-1
```

The `year:period` calculation works with ordinary daily and weekly data as well—it returns the entry "period" entries into the given year. In addition to this, you might not know that the RATS date calculations will let you "overindex" the month and day fields: while 1950:13:1 seems like it should produce an error, consider how you would figure out what the start is of the next month from an arbitrary entry. If overindexing *isn't* permitted, you would have to figure out if the entry was in December, so you could add one to the year and switch the month to 1, while other months can be handled by adding one to the month. That's all done automatically in RATS, so `%year (entry) : (%month (entry) +1) : 1` is, for any entry, the first entry in the next month. (Note that Excel allows exactly the same calculation, presumably for the same reason).

The RATSletter

© 2019 Estima

1560 Sherman Ave, Suite 1029,

Evanston, IL 60201 USA

sales@estima.com

847-864-8772