sts15	Tests for stationarity of a time series

Christopher F. Baum, Boston College, baum@bc.edu

Abstract: Implements the Elliott–Rothenberg–Stock (1996) DF-GLS test and the Kwiatkowski–Phillips–Schmidt–Shin (1992) KPSS tests for stationarity of a time series. The DF-GLS test is an improved version of the augmented Dickey–Fuller test. The KPSS test has a null hypothesis of stationarity and may be employed in conjunction with the DF-GLS test to detect long memory (fractional integration).

Keywords: stationarity, unit root, time series.

Syntax

`dfgls` *varname* [`if` *exp*] [`in` *range*] [`, maxlag(#) notrend ers`]

`kpss` *varname* [`if` *exp*] [`in` *range*] [`, maxlag(#) notrend`]

Both tests are for use with time series data; you must `tsset` your data before using these tests; see [R] **tsset**. *varname* may contain time series operators; see [U] **14.4.3 Time series varlists**.

Options

`maxlag(#)` specifies the maximum lag order to be considered. The test statistics will be calculated for each lag up to the maximum lag order (which may be zero). If not specified, the maximum lag order for the test is by default calculated from the sample size using a rule provided by Schwert (1989) using $c = 12$ and $d = 4$ in his terminology. Whether the maximum lag is explicitly specified or computed by default, the sample size is held constant over lags at the maximum available sample.

`notrend` specifies that no trend term should be included in the model. The critical values reported differ in the absence of a trend term.

`ERS` (`dfgls` only) specifies that the ERS (and Dickey–Fuller) values are to be used for all levels of significance (eschewing the response surface estimates).

Description

`dfgls` performs the Elliott–Rothenberg–Stock (ERS, 1996) efficient test for an autoregressive unit root. This test is similar to an (augmented) Dickey–Fuller t test, as performed by `dfuller`, but has the best overall performance in terms of small sample size and power, dominating the ordinary Dickey–Fuller test. The `dfgls` test "has substantially improved power when an unknown mean or trend is present" (ERS, 813).

`dfgls` applies a generalized least squares (GLS) detrending (demeaning) step to the *varname*

$$y_t^d = y_t - \widehat{\beta}' z_t$$

For detrending, $z_t = (1, t)'$ and $\widehat{\beta}_0$, $\widehat{\beta}_1$ are calculated by regressing

$$[y_1, (1 - \bar{\alpha}L)\, y_2, \ldots, (1 - \bar{\alpha}L)\, y_T]$$

onto

$$[z_1, (1 - \bar{\alpha}L)\, z_2, ..., (1 - \bar{\alpha}L)\, z_T]$$

where $\bar{\alpha} = 1 + \bar{c}/T$ with $\bar{c} = -13.5$, and L is the lag operator. For demeaning, $z_t = (1)'$ and the same regression is run with $\bar{c} = -7.0$. The values of $\bar{c}$ are chosen so that "the test achieves the power envelope against stationary alternatives (is asymptotically MPI (most powerful invariant)) at 50 percent power" (Stock 1994, 2769; emphasis added). The augmented Dickey–Fuller regression is then computed using the y_t^d series

$$\Delta y_t^d = \alpha + \gamma t + \rho y_{t-1}^d + \sum_{i=1}^{m} \delta_i \Delta y_{t-i}^d + \epsilon_t$$

where m=`maxlag`. The `notrend` option suppresses the time trend in this regression.

Approximate 5% and 10% critical values, by default, are calculated from the response surface estimates of Table 1, Cheung and Lai (1995, 413), which take both the sample size and the lag specification into account. Approximate 1% critical values for

the GLS detrended test are interpolated from Table 1 of ERS (page 825). Approximate 1% critical values for the GLS demeaned test are identical to those applicable to the no-constant, no-trend Dickey–Fuller test and are computed using the `dfuller` code. The `ERS` option specifies that the ERS (and Dickey–Fuller) values are to be used for all levels of significance (eschewing the response surface estimates).

If the maximum lag order exceeds one, the optimal lag order is calculated by the Ng and Perron (1995) sequential t test on the highest order lag coefficient, stopping when that coefficient's p-value is less than 0.10. The lag minimizing the Schwarz criterion (SC, or BIC) is printed with its minimized value.

`kpss` performs the Kwiatkowski–Phillips–Schmidt–Shin test introduced in Kwiatkowski et al. (1992) for stationarity of a time series. This test differs from those in common use (such as `dfuller` and `pperron`) by having a null hypothesis of stationarity. The test may be conducted under the null hypothesis of either trend stationarity (the default) or level stationarity. Inference from this test is complementary to that derived from those based on the Dickey–Fuller distribution (such as `dfgls`, `dfuller` and `pperron`). The KPSS test is often used in conjunction with those tests to investigate the possibility that a series is fractionally integrated; that is, neither $I(1)$ nor $I(0)$; see Lee and Schmidt (1996).

The series is detrended (demeaned) by regressing y on $z_t = (1, t)^{'}$ $\left(z_t = (1)^{'}\right)$, yielding residuals e_t. Let the partial sum series of e_t be s_t. Then the zero-order KPSS statistic $k_0 = T^{-2}\sum_{t=1}^{T} s_t^2 / T^{-1}\sum_{t=1}^{T} e_t^2$. For `maxlag` > 0, the denominator is computed as the Newey–West estimate of the long run variance of the series; see [R] **newey**.

Approximate critical values for the KPSS test are taken from Kwiatkowski et al. (1992).

Examples

Data from Terence Mills' *Econometric Analysis of Financial Time Series* on the UK FTA All Share Index of stock prices (`ftap`) and stock returns (`ftaret`) are analyzed.

```
. use http://fmwww.bc.edu/ec-p/data/Mills2d/fta.dta
. tsset
        time variable:  month, 1965m1 to 1995m12
. dfgls ftap
Number of obs =   355
Maxlag = 16 chosen by Schwert criterion

                     Test        1% Critical       5% Critical      10% Critical
                Statistic            Value             Value             Value
------------------------------------------------------------------------------
DF-GLS(tau)[16]    -0.068           -3.480            -2.818            -2.536
DF-GLS(tau)[15]    -0.155           -3.480            -2.824            -2.542
DF-GLS(tau)[14]    -0.046           -3.480            -2.829            -2.547
DF-GLS(tau)[13]    -0.234           -3.480            -2.835            -2.552
DF-GLS(tau)[12]    -0.131           -3.480            -2.840            -2.557
DF-GLS(tau)[11]    -0.196           -3.480            -2.846            -2.562
DF-GLS(tau)[10]    -0.251           -3.480            -2.851            -2.566
DF-GLS(tau)[9]     -0.173           -3.480            -2.856            -2.571
DF-GLS(tau)[8]     -0.107           -3.480            -2.861            -2.575
DF-GLS(tau)[7]     -0.361           -3.480            -2.865            -2.580
DF-GLS(tau)[6]     -0.391           -3.480            -2.870            -2.584
DF-GLS(tau)[5]     -0.476           -3.480            -2.874            -2.588
DF-GLS(tau)[4]     -0.524           -3.480            -2.879            -2.592
DF-GLS(tau)[3]     -0.484           -3.480            -2.883            -2.595
DF-GLS(tau)[2]     -0.507           -3.480            -2.887            -2.599
DF-GLS(tau)[1]     -0.789           -3.480            -2.891            -2.602
Opt Lag (Ng-Perron sequential t) = 15 with RMSE  35.59803
Min SC =  7.275482 at lag  2 with RMSE    37.0745
. kpss ftap
KPSS test for ftap
Maxlag = 16 chosen by Schwert criterion
Critical values for H0: ftap is trend stationary
10%: 0.119  5% : 0.146  2.5%: 0.176  1% : 0.216
Lag order    Test statistic
    0        7.90141
    1        4.18402
    2        2.86036
    3        2.18027
    4        1.76579
```

```
    5        1.48676
    6         1.2861
    7        1.13475
    8        1.01642
    9        .921225
   10         .84288
   11        .777242
   12        .721428
   13        .673349
   14        .631492
   15        .594708
   16        .562121
. dfgls ftaret
Number of obs =   355
Maxlag = 16 chosen by Schwert criterion
                     Test        1% Critical       5% Critical      10% Critical
                  Statistic           Value             Value             Value
------------------------------------------------------------------------------
DF-GLS(tau)[16]     -4.161           -3.480            -2.818            -2.536
DF-GLS(tau)[15]     -4.119           -3.480            -2.824            -2.542
DF-GLS(tau)[14]     -4.413           -3.480            -2.829            -2.547
DF-GLS(tau)[13]     -4.733           -3.480            -2.835            -2.552
DF-GLS(tau)[12]     -4.663           -3.480            -2.840            -2.557
DF-GLS(tau)[11]     -4.392           -3.480            -2.846            -2.562
DF-GLS(tau)[10]     -4.653           -3.480            -2.851            -2.566
DF-GLS(tau)[9]      -4.795           -3.480            -2.856            -2.571
DF-GLS(tau)[8]      -4.931           -3.480            -2.861            -2.575
DF-GLS(tau)[7]      -6.006           -3.480            -2.865            -2.580
DF-GLS(tau)[6]      -6.203           -3.480            -2.870            -2.584
DF-GLS(tau)[5]      -6.911           -3.480            -2.874            -2.588
DF-GLS(tau)[4]      -7.614           -3.480            -2.879            -2.592
DF-GLS(tau)[3]      -7.769           -3.480            -2.883            -2.595
DF-GLS(tau)[2]      -9.176           -3.480            -2.887            -2.599
DF-GLS(tau)[1]     -13.075           -3.480            -2.891            -2.602
Opt Lag (Ng-Perron sequential t) =  8 with RMSE  .0593867
Min SC =  -5.566828 at lag  2 with RMSE  .0603119
. dfgls ftaret,notrend
Number of obs =   355
Maxlag = 16 chosen by Schwert criterion
                     Test        1% Critical       5% Critical      10% Critical
                  Statistic           Value             Value             Value
------------------------------------------------------------------------------
DF-GLS(mu)[16]      -3.165           -2.580            -1.952            -1.637
DF-GLS(mu)[15]      -3.161           -2.580            -1.955            -1.640
DF-GLS(mu)[14]      -3.430           -2.580            -1.958            -1.643
DF-GLS(mu)[13]      -3.725           -2.580            -1.962            -1.646
DF-GLS(mu)[12]      -3.711           -2.580            -1.965            -1.649
DF-GLS(mu)[11]      -3.528           -2.580            -1.968            -1.652
DF-GLS(mu)[10]      -3.776           -2.580            -1.971            -1.655
DF-GLS(mu)[9]       -3.933           -2.580            -1.974            -1.658
DF-GLS(mu)[8]       -4.087           -2.580            -1.977            -1.660
DF-GLS(mu)[7]       -5.039           -2.580            -1.980            -1.663
DF-GLS(mu)[6]       -5.278           -2.580            -1.982            -1.665
DF-GLS(mu)[5]       -5.966           -2.580            -1.985            -1.668
DF-GLS(mu)[4]       -6.679           -2.580            -1.988            -1.670
DF-GLS(mu)[3]       -6.928           -2.580            -1.990            -1.672
DF-GLS(mu)[2]       -8.312           -2.580            -1.993            -1.675
DF-GLS(mu)[1]      -12.060           -2.580            -1.995            -1.677
Opt Lag (Ng-Perron sequential t) =  8 with RMSE  .0600067
Min SC = -5.53158 at lag  2 with RMSE  .0613843
```

Both tests indicate that `ftap` appears to be nonstationary. `ftaret` appears to be both trend and level stationary.

Saved Results

dfgls saves the following scalars in r():

r(N)	number of observations
r(optlag)	optimal lag order
r(sc*n*)	Schwarz criterion at lag *n*
r(rmse*n*)	root mean square error at lag *n*
r(dft*n*)	DF-GLS statistic at lag *n*

kpss saves the following scalars in r():

r(N)	number of observations
r(dft*n*)	KPSS statistic at lag *n*

Acknowledgments

I acknowledge useful conversations with Serena Ng, James Stock, and Vince Wiggins. The KPSS code was adapted from John Barkoulas' RATS code for that test. Thanks also to Richard Sperling for tracking down a discrepancy between published work and the dfgls output and alerting me to the Cheung and Lai estimates. Any remaining errors are my own.

References

Cheung, Y. W. and K.-S. Lai. 1995. Lag order and critical values of a modified Dickey–Fuller test. *Oxford Bulletin of Economics and Statistics* 57: 411–419.

Elliott, G., T. J. Rothenberg, and J. H. Stock. 1996. Efficient tests for an autoregressive unit root. *Econometrica* 64: 813–836.

Kwiatkowski, D., P. C. Phillips, P. Schmidt, and Y. Shin. 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54: 159–178.

Lee, D. and P. Schmidt. 1996. On the power of the KPSS test of stationarity against fractionally-integrated alternatives. *Journal of Econometrics* 73: 285–302.

Ng, S. and P. Perron. 1995. Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag. *Journal of the American Statistical Association* 90: 268–281.

Schwert, G. W. 1989. Tests for unit roots: A Monte Carlo investigation. *Journal of Business and Economic Statistics* 7: 147–160.

Stock, J. H. 1994. Unit roots, structural breaks and trends. In *Handbook of Econometrics IV*, ed. R. F. Engle and D. L. McFadden. Amsterdam: Elsevier.

sts16 | Tests for long memory in a time series

Christopher F. Baum, Boston College, baum@bc.edu
Vince Wiggins, Stata Corporation, vwiggins@stata.com

Abstract: Implements the Geweke/Porter-Hudak log periodogram estimator (1983), the Phillips modified log periodogram estimator (1999b) and the Robinson log periodogram estimator (1995) for the diagnosis of long memory, or fractional integration, in a time series. The Robinson estimator may be applied to a set of time series.

Keywords: fractional integration, long memory, stationarity, time series.

Syntax

gphudak *varname* [if *exp*] [in *range*] [, powers(*numlist*)]

modlpr *varname* [if *exp*] [in *range*] [, powers(*numlist*) notrend]

roblpr *varlist* [if *exp*] [in *range*] [, powers(*numlist*) l(#) j(#) constraints(*numlist*)]

These tests are for use with time series data; you must tsset your data before using these tests; see [R] **tsset**. *varname* or *varlist* may contain time series operators; see [U] **14.4.3 Time-series varlists**.

Options

powers(*numlist*) indirectly specifies the number of ordinates to be included in the regression. A number of ordinates equal to the integer part of T raised to the powers(*numlist*) will be used. Powers ranging from 0.50 to 0.75 are commonly employed for gphudak and modlpr. These routines use the default power of 0.5. roblpr uses the default power of 0.9. For roblpr, multiple powers may only be specified if a single variable appears in *varlist*.

notrend specifies that detrending is not to be applied by modlpr. By default, a linear trend will be removed from the series.

`l(#)` specifies the number of initial ordinates to be removed from the regression for `roblpr`. Some researchers have found that such exclusion improves the properties of tests based on log-periodogram regressions. The default value of `l` is zero.

`j(#)` specifies that the log periodogram employed in `roblpr` is to be computed as an average of adjacent ordinates. The default value of `j` is 1, so that no averaging is performed. If `j` is 2, the number of ordinates is halved; with a `j` of 3, divided by three, and so on. When `j` is greater than 1, the value of `powers` should be set large enough so that the averaged ordinates are sufficient in number.

`constraints(`*numlist*`)` specifies the constraint numbers of the linear constraints to be applied during estimation in `roblpr`. The default is to perform unconstrained estimation. This option allows the imposition of linear constraints prior to estimation of the pooled coefficient vector. For instance, if *varlist* contains prices, dividends, and returns, and your prior (or previous findings) states that prices' and dividends' order of integration is indistinguishable, one might impose that constraint to improve the power of the F test provided by `roblpr`. You would specify the constraints prior to the `roblpr` command and then provide the list of constraints in the constraints option to `roblpr`.

Technical note on constraints. When constraints are imposed it is difficult to identify the number of numerator degrees of freedom in the test for equality of d coefficients reported at the bottom of `roblpr`'s output. Since constraints can be of any general form and it is possible to specify constraints that are not unique, `roblpr` determines the degrees of freedom from the rank of the matrix used to compute the Wald statistic. Determining that matrix rank from a numerical standpoint can be problematic, in which case `roblpr` may overstate the number of constraints being tested and thereby incorrectly compute the numerator degrees of freedom for the test. This rarely has a meaningful impact on the statistical test, but you may wish to test only the unconstrained coefficients if the computed degrees of freedom are wrong.

For example, after the final example below, we could perform the test by typing `test ftap == ftaret`. In this case, the degrees of freedom were correct, so we needn't have gone to the trouble.

Description

The model of an autoregressive fractionally integrated moving average process of a time series of order $(p,\ d,\ q)$, denoted by ARFIMA$(p,\ d,\ q)$, with mean μ, may be written using operator notation as

$$\Phi(L)(1-L)^d\,(y_t-\mu)=\Theta(L)\epsilon_t,\ \epsilon_t\sim i.i.d.(0,\sigma_\epsilon^2) \tag{1}$$

where L is the backward-shift operator,

$$\Phi(L)=1-\phi_1 L-\cdots-\phi_p L^p$$

$\Theta(L)=1+\vartheta_1 L+\cdots+\vartheta_q L^q$, and $(1-L)^d$ is the fractional differencing operator defined by

$$(1-L)^d=\sum_{k=0}^{\infty}\frac{\Gamma(k-d)L^k}{\Gamma(-d)\Gamma(k+1)} \tag{2}$$

with $\Gamma(\cdot)$ denoting the gamma (generalized factorial) function. The parameter d is allowed to assume any real value. The arbitrary restriction of d to integer values gives rise to the standard autoregressive integrated moving average (ARIMA) model. The stochastic process y_t is both stationary and invertible if all roots of $\Phi(L)$ and $\Theta(L)$ lie outside the unit circle and $|d|<0.5$. The process is nonstationary for $d\geq 0.5$, as it possesses infinite variance; for example, see Granger and Joyeux (1980).

Assuming that $d\in[0,0.5)$, Hosking (1981) showed that the autocorrelation function, $\rho(\cdot)$, of an ARFIMA process is proportional to k^{2d-1} as $k\to\infty$. Consequently, the autocorrelations of the ARFIMA process decay hyperbolically to zero as $k\to\infty$ in contrast to the faster, geometric decay of a stationary ARMA process. For $d\in(0,0.5)$, $\sum_{j=-n}^{n}|\rho(j)|$ diverges as $n\to\infty$, and the ARFIMA process is said to exhibit long memory, or long-range positive dependence. The process is said to exhibit intermediate memory (anti-persistence), or long-range negative dependence, for $d\in(-0.5,0)$. The process exhibits short memory for $d=0$, corresponding to stationary and invertible ARMA modeling. For $d\in[0.5,1)$ the process is mean reverting, even though it is not covariance stationary, as there is no long-run impact of an innovation on future values of the process.

If a series exhibits long memory, it is neither stationary ($I(0)$) nor is it a unit root ($I(1)$) process; it is an $I(d)$ process, with d a real number. A series exhibiting long memory, or persistence, has an autocorrelation function that damps hyperbolically, more slowly than the geometric damping exhibited by "short memory" (ARMA) processes. Thus, it may be predictable at long horizons. Long memory models originated in hydrology and have been widely applied in economics and finance. An excellent survey of long memory models is given by Baillie (1996).

There are two approaches to the estimation of an ARFIMA $(p,\ d,\ q)$ model: exact maximum likelihood estimation, as proposed by Sowell (1992), and semiparametric approaches, as described in this insert. Sowell's approach requires specification of the p and q values, and estimation of the full ARFIMA model conditional on those choices. This involves all the attendant

difficulties of choosing an appropriate ARMA specification, as well as a formidable computational task for each combination of p and q to be evaluated. The methods described here assume that the short memory or ARMA components of the time series are relatively unimportant, so that the long memory parameter d may be estimated without fully specifying the data-generating process. These methods are thus described as semiparametric.

`gphudak` performs the Geweke and Porter-Hudak (GPH 1983) semiparametric log periodogram regression, often described as the "GPH test," for long memory (fractional integration) in a time series. The GPH method uses nonparametric methods—a spectral regression estimator—to evaluate d without explicit specification of the ARMA parameters of the series. The series is usually differenced so that the resulting d estimate will fall in the $[-0.5, 0.5]$ interval.

Geweke and Porter-Hudak (1983) proposed a semiparametric procedure to obtain an estimate of the memory parameter d of a fractionally integrated process X_t in a model of the form

$$(1-L)^d X_t = \epsilon_t, \tag{3}$$

where ϵ_t is stationary with zero mean and continuous spectral density $f_\epsilon(\lambda) > 0$. The estimate $\widehat{d}$ is obtained from the application of ordinary least squares to

$$\log\left(I_x(\lambda_s)\right) = \widehat{c} - \widehat{d}\log\left|1-e^{i\lambda_s}\right|^2 + \text{residual} \tag{4}$$

computed over the fundamental frequencies $\{\lambda_s = 2\pi s/n, s = 1, \ldots, m < n\}$. We define

$$\omega_x(\lambda_s) = \frac{1}{\sqrt{2\pi n}}\sum_{t=1}^{n} X_t e^{it\lambda_s}$$

as the discrete Fourier transform (DFT) of the time series X_t, $I_x(\lambda_s) = \omega_x(\lambda_s)\,\omega_x(\lambda_s)^*$ as the periodogram, and $x_s = \log\left|1-e^{i\lambda_s}\right|$. Ordinary least squares on (4) yields

$$\widehat{d} = \frac{\sum_{s=1}^{m} x_s \log I_x(\lambda_s)}{2\sum_{s=1}^{m} x_s^2} \tag{5}$$

Various authors have proposed methods for the choice of m, the number of Fourier frequencies included in the regression. The regression slope estimate is an estimate of the slope of the series' power spectrum in the vicinity of the zero frequency; if too few ordinates are included, the slope is calculated from a small sample. If too many are included, medium and high-frequency components of the spectrum will contaminate the estimate. A choice of $\sqrt{T}$ or 0.5 for `power` is often employed. To evaluate the robustness of the GPH estimate, a range of power values (from 0.40 to 0.75) is commonly calculated as well. Two estimates of the d coefficient's standard error are commonly employed: the regression standard error, giving rise to a standard t test, and an asymptotic standard error, based upon the theoretical variance of the log periodogram of $\pi^2/6$. The statistic based upon that standard error has a standard normal distribution under the null.

`modlpr` computes a modified form of the GPH estimate of the long memory parameter, d, of a time series, proposed by Phillips (1999a, 1999b). Phillips (1999a) points out that the prior literature on this semiparametric approach does not address the case of $d = 1$, or a unit root, in (3), despite the broad interest in determining whether a series exhibits unit-root behavior or long memory behavior, and his work showing that the $\widehat{d}$ estimate of (5) is inconsistent when $d > 1$, with $\widehat{d}$ exhibiting asymptotic bias toward unity. This weakness of the GPH estimator is solved by Phillips' modified log periodogram regression estimator, in which the dependent variable is modified to reflect the distribution of d under the null hypothesis that $d = 1$. The estimator gives rise to a test statistic for $d = 1$ which is a standard normal variate under the null. Phillips suggests that deterministic trends should be removed from the series before application of the estimator. Accordingly, the routine will automatically remove a linear trend from the series. This may be suppressed with the `notrend` option. The comments above regarding `power` apply equally to `modlpr`.

The Phillips (1999b) modification of the GPH estimator is based on an exact representation of the DFT in the unit root case. The modification expresses

$$\omega_x(\lambda_s) = \frac{\omega_u(\lambda_s)}{1-e^{i\lambda_s}} - \frac{e^{i\lambda_s}}{1-e^{i\lambda_s}}\frac{X_n}{\sqrt{2\pi n}}$$

and the modified DFT as

$$v_x(\lambda_s) = \omega_x(\lambda_s) + \frac{e^{i\lambda_s}}{1-e^{i\lambda_s}}\frac{X_n}{\sqrt{2\pi n}}$$

with associated periodogram ordinates $I_v(\lambda_s) = v_x(\lambda_s)\, v_x(\lambda_s)^*$ (Phillips 1999b, 9). He notes that both $v_x(\lambda_s)$ and, thus, $I_v(\lambda_s)$ are observable functions of the data. The log-periodogram regression is now the regression of $\log I_v(\lambda_s)$ on $a_s = \log\left|1 - e^{i\lambda_s}\right|$. Defining $\bar{a} = m^{-1}\sum_{s=1}^{m} a_s$ and $x_s = a_s - \bar{a}$, the modified estimate of the long-memory parameter becomes

$$\tilde{d} = \frac{\sum_{s=1}^{m} x_s \log I_\nu(\lambda_s)}{2\sum_{s=1}^{m} x_s^2} \tag{6}$$

Phillips proves that, with appropriate assumptions on the distribution of ϵ_t, the distribution of $\tilde{d}$ follows

$$\sqrt{m}\left(\tilde{d} - d\right) \rightarrow N\left(0, \frac{\pi^2}{24}\right) \tag{7}$$

in distribution, so $\tilde{d}$ has the same limiting distribution at $d = 1$ as does the GPH estimator in the stationary case so $\tilde{d}$ is consistent for values of d around unity. A semiparametric test statistic for a unit root against a fractional alternative is then based upon the statistic (Phillips 1999a, 10)

$$z_d = \frac{\sqrt{m}\left(\tilde{d} - 1\right)}{\pi/24} \tag{8}$$

with critical values from the standard normal distribution. This test is consistent against both $d < 1$ and $d > 1$ fractional alternatives.

`roblpr` computes the Robinson (1995) multivariate semiparametric estimate of the long memory (fractional integration) parameters, $d(g)$, of a set of G time series, $y(g)$, $g = 1, G$ with $G \geq 1$. When applied to a set of time series, the $d(g)$ parameter for each series is estimated from a single log-periodogram regression which allows the intercept and slope to differ for each series. One of the innovations of Robinson's estimator is that it is not restricted to using a small fraction of the ordinates of the empirical periodogram of the series, that is, the reasonable values of `power` need not exclude a sizable fraction of the original sample size. The estimator also allows for the removal of one or more initial ordinates and for the averaging of the periodogram over adjacent frequencies. The rationale for using non-default values of either of these options is presented in Robinson (1995).

Robinson (1995) proposes an alternative log-periodogram regression estimator which he claims provides "modestly superior asymptotic efficiency to $\bar{d}(0)$", ($\bar{d}(0)$ being the Geweke and Porter-Hudak estimator) Robinson (1995, 1052). Robinson's formulation of the log-periodogram regression also allows for the formulation of a multivariate model, providing justification for tests that different time series share a common differencing parameter. Normality of the underlying time series is assumed, but Robinson claims that other conditions underlying his derivation are milder than those conjectured by GPH.

We present here Robinson's multivariate formulation, which applies to a single time series as well. Let X_t represent a G-dimensional vector with g^{th} element $X_{gt}, g = 1, \ldots, G$. Assume that X_t has a spectral density matrix $\int_{-\pi}^{\pi} e^{ij\lambda} f(\lambda)\, d\lambda$, with (g, h) element denoted as $f_{gh}(\lambda)$. The gth diagonal element, $f_{gg}(\lambda)$, is the power spectral density of X_{gt}. For $0 < C_g < \infty$ and $-1/2 < d_g < 1/2$, assume that $f_{gg}(\lambda) \sim C_g \lambda^{-2d_g}$ as $\lambda \rightarrow 0+$ for $g = 1, \ldots, G$. The periodogram of X_{gt} is then denoted as

$$I_g(\lambda) = (2\pi n)^{-1}\left|\sum_{t=1}^{n} X_{gt} e^{it\lambda}\right|^2, g = 1, \ldots, G \tag{9}$$

Without averaging the periodogram over adjacent frequencies nor omission of l initial frequencies from the regression, we may define $Y_{gk} = \log I_g(\lambda_k)$. The least squares estimates of $c = (c_1, \ldots, c_G)'$ and $d = (d_1, \ldots, d_G)'$ are given by

$$\begin{bmatrix} \tilde{c} \\ \tilde{d} \end{bmatrix} = \text{vec}\left\{Y'Z(Z'Z)^{-1}\right\} \tag{10}$$

where $Z = (Z_1, \ldots, Z_m)'$, $Z_k = (1, -2\log\lambda_k)'$, $Y = (Y_1, \ldots, Y_G)$, and $Y_g = (Y_{g,1}, \ldots, Y_{g,m})'$ for m periodogram ordinates. Standard errors for $\tilde{d}_g$ and for a test of the restriction that two or more of the d_g are equal may be derived from the estimated covariance matrix of the least squares coefficients. The standard errors for the estimated parameters are derived from a pooled estimate of the variance in the multivariate case, so that their interval estimates differ from those of their univariate counterparts. Modifications to this derivation when the frequency-averaging (`j`) or omission of initial frequencies (`l`) options are selected may be found in Robinson (1995).

Examples

Data from Terence Mills' *Econometric Analysis of Financial Time Series* on UK FTA All Share stock returns (`ftaret`) and dividends (`ftadiv`) are analyzed.

```
. use http://fmwww.bc.edu/ec-p/data/Mills2d/fta.dta

. tsset
        time variable:  month, 1965m1 to 1995m12

. gphudak ftaret,power(0.5 0.6 0.7)

GPH estimate of fractional differencing parameter
------------------------------------------------------------------------------
                                                        Asy.
Power   Ords     Est d   StdErr  t(H0: d=0)  P>|t|    StdErr  z(H0: d=0)  P>|z|
------------------------------------------------------------------------------
  .50     20  -.00204  .160313    -0.0127   0.990   .187454    -0.0109   0.991
  .60     35  .228244  .145891     1.5645   0.128   .130206     1.7529   0.080
  .70     64  .141861  .089922     1.5776   0.120   .091267     1.5544   0.120
------------------------------------------------------------------------------


. modlpr ftaret, power(0.5 0.55:0.8)

Modified LPR estimate of fractional differencing parameter
------------------------------------------------------------------------------
Power   Ords      Est d    Std Err  t(H0: d=0)    P>|t|    z(H0: d=1)    P>|z|
------------------------------------------------------------------------------
  .50     19   .0231191    .139872     0.1653     0.870      -6.6401     0.000
  .55     25   .2519889   .1629533     1.5464     0.135      -5.8322     0.000
  .60     34   .2450011   .1359888     1.8016     0.080      -6.8650     0.000
  .65     46   .1024504   .1071614     0.9560     0.344      -9.4928     0.000
  .70     63   .1601207   .0854082     1.8748     0.065     -10.3954     0.000
  .75     84   .1749659     .08113     2.1566     0.034     -11.7915     0.000
  .80    113   .0969439   .0676039     1.4340     0.154     -14.9696     0.000
------------------------------------------------------------------------------

. roblpr ftaret

Robinson estimates of fractional differencing parameter
------------------------------------------------------------
Power   Ords      Est d    Std Err  t(H0: d=0)    P>|t|
------------------------------------------------------------
  .90    205   .1253645   .0446745     2.8062     0.005
------------------------------------------------------------

. roblpr ftap ftadiv

Robinson estimates of fractional differencing parameters
Power =  .90                            Ords         =    205
------------------------------------------------------------
Variable         |     Est d       Std Err      t       P>|t|
-----------------+------------------------------------------
ftap             |  .8698092      .0163302  53.2640     0.000
ftadiv           |  .8717427      .0163302  53.3824     0.000
------------------------------------------------------------
Test for equality of d coefficients:   F(1,406) =  .00701   Prob > F = 0.9333

. constraint define 1 ftap=ftadiv

. roblpr ftap ftadiv ftaret, c(1)

Robinson estimates of fractional differencing parameters
Power =  .90                            Ords         =    205
------------------------------------------------------------
Variable         |     Est d       Std Err      t       P>|t|
-----------------+------------------------------------------
ftap             |  .8707759      .0205143  42.4473     0.000
ftadiv           |  .8707759      .0205143  42.4473     0.000
ftaret           |  .1253645      .0290116   4.3212     0.000
------------------------------------------------------------
Test for equality of d coefficients:   F(1,610) =  440.11   Prob > F = 0.0000
```

The GPH test, applied to the stock returns series, generates estimates of the long memory parameter that cannot reject the null at the ten percent level using the t test. Phillips' modified LPR, applied to this series, finds that $d = 1$ can be rejected for all powers tested, while $d = 0$ (stationarity) may be rejected at the ten percent level for powers 0.6, 0.7, and 0.75. Robinson's estimate for the returns series alone is quite precise. Robinson's multivariate test, applied to the price and dividends series, finds that each series has $d > 0$. The test that they share the same d cannot be rejected. Accordingly, the test is applied to all three series subject to the constraint that price and dividends series have a common d, yielding a more precise estimate of the difference in d parameters between those series and the stock returns series.

Saved Results

`gphudak` saves in `e()`:

`e(N_powers)`	number of powers (scalar)
`e(depvar)`	dependent variable name (macro)
`e(gph)`	matrix of results, 9 by `N_powers`

`modlpr` saves in `e()`:

`e(N_powers)`	number of powers (scalar)
`e(depvar)`	dependent variable name (macro)
`e(modlpr)`	matrix of results, 8 by `N_powers`

`roblpr` saves the following scalars in `r()`:

`r(N)`	number of observations
`r(rob)`	d estimate
`r(se)`	estimated standard error of d
`r(t)`	t statistic
`r(p)`	p-value of t statistic

If more than one power is specified in `roblpr`, the saved results pertain to the last power used.

Acknowledgments

The first author acknowledges John Barkoulas' original exposition of the ARFIMA model, and thanks Peter Phillips for clarifying comments on his working papers. Any remaining errors are the authors' responsibility.

References

Baillie, R. 1996. Long memory processes and fractional integration in econometrics. *Journal of Econometrics* 73: 5–59.

Geweke, J. and S. Porter-Hudak. 1983. The estimation and application of long memory time series models. *Journal of Time Series Analysis* 4: 221–238.

Granger, C. W. J. and R. Joyeux. 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1: 15–39.

Hosking, J. R. M. 1981. Fractional differencing. *Biometrika* 68: 165–176.

Phillips, P. C. B. 1999a. Discrete Fourier transforms of fractional processes. Unpublished working paper No. 1243, Cowles Foundation for Research in Economics, Yale University. `http://cowles.econ.yale.edu/P/cd/d12a/d1243.pdf`

——. 1999b. Unit root log periodogram regression. Unpublished working paper No. 1244, Cowles Foundation for Research in Economics, Yale University. `http://cowles.econ.yale.edu/P/cd/d12a/d1244.pdf`

Robinson, P. M. 1995. Log-periodogram regression of time series with long range dependence. *Annals of Statistics* 23: 1048–1072.

Sowell, F. 1992. Maximum likelihood estimation of stationary univariate fractionally-integrated time-series models, *Journal of Econometrics* 53: 165–188.

sts17 Compacting time series data

Christopher F. Baum, Boston College, baum@bc.edu

Abstract: `tscollap` provides the ability to compact data of monthly, quarterly or half-yearly frequencies to a lower frequency by one or more methods (e.g., average, sum, last value per period, and so on).

Keywords: time series, data frequency, collapse.

Syntax

`tscollap` *clist* `, to(`*freq*`)` [`generate(`*freqvar*`)`]

where *clist* is either

[(*stat*)] *varlist* [[(*stat*)] ...]

or

[(*stat*) *target_var* = *varname* [*target_var* = *varname* ...] [[(*stat*) ...]]

or any combination of the *varlist* or *target_var* forms, and *stat* is one of

```
Y-variable: logwt
Grouped by: foreign (Car type)
Group numbers:
   Car type |      Freq.     Percent        Cum.
------------+-----------------------------------
   Domestic |         52       70.27       70.27
    Foreign |         22       29.73      100.00
------------+-----------------------------------
      Total |         74      100.00
Transformation: Fisher´s z
95% confidence interval(s) for percentile ratio(s)
between values of exp(logwt) in first and second groups:
      Percent   Pctl_Rat    Minimum    Maximum
r1         25  1.1935375  1.0341465  1.3533567
r2         50  1.4806389  1.3101849  1.6280196
r3         75   1.744916  1.6079542  1.8772724
```

We note that, typically, American cars are 148% heavier than foreign cars, with confidence limits ranging from 131% to 163% as heavy. The 25th percentile ratio (103% to 135%) shows that the two car types do not overlap a great deal.

Saved results

`cendif` saves in `r()`:

Scalars

`r(N)`	number of observations	`r(N_clust)`	number of clusters
`r(N_1)`	sample size N_1	`r(N_2)`	sample size N_2
`r(df_r)`	residual degrees of freedom (if `tdist` present)		

Macros

`r(depvar)`	name of dependent variable	`r(by)`	name of `by` variable defining groups
`r(clustvar)`	name of cluster variable	`r(tdist)`	`tdist` if specified
`r(wtype)`	weight type	`r(wexp)`	weight expression
`r(centiles)`	list of percents for percentiles	`r(Dslist)`	list of D^*-values for percentiles
`r(transf)`	transformation specified by `transf`	`r(tranlab)`	transformation label in output
`r(level)`	confidence level	`r(eform)`	`eform` if specified

Matrices

`r(cimat)`	confidence intervals for differences or ratios	`r(Dsmat)`	upper and lower limits for $D^*(\theta)$

Acknowledgments

I would like to thank Nick Cox of Durham University, UK, and Bill Gould of Stata Corporation for some very helpful advice on the coding of infinite confidence limits, such as those occasionally resulting from Equation (9).

References

Campbell, M. J. and M. J. Gardner. 1988. Calculating confidence intervals for some non-parametric analyses. *British Medical Journal* 296: 1454–1456.

Conover, W. J. 1980. *Practical Nonparametric Statistics*. 2d ed. New York: John Wiley & Sons.

Gardner, M. J. and D. G. Altman. 1989. *Statistics with Confidence*. London: British Medical Journal.

Newson, R. 2000. snp15.2: Update to somersd. *Stata Technical Bulletin* 58: 30.

Royston, P. 1998. CID: Stata module to calculate confidence intervals for means or differences. On the Ideas list at http://ideas.uqam.ca/ideas/data/Softwares/bocbocodeS338001.html.

Somers, R. H. 1962. A New asymmetric measure of association for ordinal variables. *American Sociological Review* 27: 799–811.

sts15.1 | Tests for stationarity of a time series: update

Christopher F. Baum, Boston College, baum@bc.edu
Richard Sperling, The Ohio State University, rsperling@boo.net

Abstract: Enhances the Elliott–Rothenberg-Stock DF-GLS test and the Kwiatkowski–Phillips–Schmidt–Shin KPSS tests for stationarity of a time series introduced in Baum (2000) and corrects an error in both routines.

Keywords: stationarity, unit root, time series.

Changes to dfgls

`dfgls` did not handle missing initial values properly. That is, if the time series variable specified had initial values not excluded by `if` or `in` conditions, those values were improperly considered in the construction of the sample size. This would apply as well to the consideration of variables with time series operators, such as `D.gdp`, since those variables will have at least one missing observation at the outset. This has been corrected.

The `dfgls` routine has been enhanced to add a very powerful lag selection criterion, the "modified AIC" (MAIC) criterion proposed by Ng and Perron (2000). They have established that use of this MAIC criterion may provide "huge size improvements" in the `dfgls` test. The criterion, indicating the appropriate lag order, is printed on `dfgls` output and may be used to select the test statistic from which inference is to be drawn.

It should be noted that all of the lag length criteria employed by `dfgls` (the sequential *t* test of Ng and Perron 1995, the SC, and the MAIC) are calculated, for various lags, by holding the sample size fixed at that defined for the longest lag. These criteria cannot be meaningfully compared over lag lengths if the underlying sample is altered to use all available observations. That said, if the optimal lag length (by whatever criterion) is found to be much less than that picked by the Schwert criterion, it would be advisable to rerun the test with the `maxlag` option specifying that optimal lag length, especially when using samples of modest size.

New syntax for kpss

`kpss` *varname* [`if` *exp*] [`in` *range*] [`, maxlag(#) notrend qs auto`]

`kpss` did not make use of all available observations in the computation of the autocovariance function. This has been corrected. The online help file now provides instructions for reproducing the statistics reported in Kwiatkowski et al. (1992) from a dataset available online.

The `kpss` routine has been enhanced to add two options recommended by the work of Hobijn et al. (1998). An automatic bandwidth selection routine has been added, rendering it unnecessary to evaluate a range of test statistics for various lags. An option to weight the empirical autocovariance function by the quadratic spectral kernel, rather than the Bartlett kernel employed by KPSS, has also been introduced. These options may be used separately or in combination. It is in combination that Hobijn et al. found the greatest improvement in the test: "Our Monte Carlo simulations show that the best small sample results of the test in case the process exhibits a high degree of persistence are obtained using both the automatic bandwidth selection procedure and the Quadratic Spectral kernel" (1998, 14).

New options

`qs` specifies that the autocovariance function is to be weighted by the quadratic spectral kernel, rather than the Bartlett kernel. Andrews (1991) and Newey and West (1994) "indicate that it yields more accurate estimates of σ_ϵ^2 than other kernels in finite samples" (Hobijn et al. 1998, 6).

`auto` specifies that the automatic bandwidth selection procedure proposed by Newey and West (1994), as described by Hobijn et al. (1998, 7), is used to determine `maxlag` in two stages. First, the "a priori nonstochastic bandwidth parameter" n_T is chosen as a function of the sample size and the specified kernel. The autocovariance function of the estimated residuals is calculated and used to generate γ as a function of sums of autocorrelations. The `maxlag` to be used in computing the long-run variance, $\widehat{m}_T$, is then calculated as $\min\left[T, \text{int}\left[\widehat{\gamma}T^{\theta}\right]\right]$ where $\theta = 1/3$ for the Bartlett kernel and $\theta = 1/5$ for the quadratic spectral kernel.

Additional saved results

`dfgls` saves the modified AIC at lag n in `r(maic`*n*`)`.

References

Andrews, D. W. K. 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59: 817–858.

Baum, C. F. 2000. sts15: Test for stationarity of a time series. *Stata Technical Bulletin* 57: 36–39.

Hobijn, B., P. H. Franses, and M. Ooms. 1998. Generalizations of the KPSS-test for stationarity. Econometric Institute Report 9802/A, Econometric Institute, Erasmus University Rotterdam. `http://www.eur.nl/few/ei/papers`.

Kwiatkowski, D., P. C. B. Phillips, P. Schmidt, and Y. Shin. 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54: 159–178.

Newey, W. K. and K. D. West. 1994. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies* 61: 631–653.

Ng, S. and P. Perron. 1995. Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag. *Journal of the American Statistical Association* 90: 268–281.

——. 2000. Lag length selection and the construction of unit root tests with good size and power. *Econometrica*, in press.