AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

A seasonal unit-root test with Stata

Domenico Depalo
Bank of Italy
Roma, Italy
domenico.depalo@bancaditalia.it

Abstract. Many economic time series exhibit important systematic fluctuations within the year, i.e., seasonality. In contrast to usual practice, I argue that using original data should always be considered, although the process is more complicated than that of using seasonally adjusted data. Motivations to use unadjusted data come from the information contained in their peaks and troughs and from economic theory. One major complication is the possible unit root at seasonal frequencies. In this article, I tackle the issue of implementing a test to identify the source of seasonality. In particular, I follow Hylleberg et al. (1990, *Journal of Econometrics* 44: 215–238) for quarterly data.

Keywords: st0172, sroot, unit roots, seasonality

1 Introduction

Many economic time series exhibit important systematic fluctuations within the year, i.e., seasonality. Although applied econometricians have long used seasonally adjusted data, there exists increasing consensus that this practice is suboptimal for at least two reasons. First, peaks and troughs convey information that is lost during the adjustment; second, seasonally adjusted data often conflicts with the economic theory. Consider the rational expectation hypothesis or the permanent income hypothesis. Seasonal adjustment, for example, by the widely known CENSUS-X11, invalidates the theory by construction, because it is a two-sided filter, which thus violates the key orthogonality condition between the data at time t and the available information at the same time. To avoid these flaws, one can use the original data and either control for a set of seasonal dummies or redefine the error term to incorporate the seasonal fluctuations. The first solution is weak because “data adjusted by the seasonal dummy technique will [...] tend to reject the model if it contains fundamental nonlinearities” (Miron 1986, 1260). The second solution is wrong because the error terms would be predictable to some extent, thus invalidating the rational expectation hypothesis (Osborn 1988). These simple facts have two important consequences: using seasonally adjusted data can have serious consequences on our results and the treatment of seasonality requires a serious systematic approach.

In this article, I am particularly interested in seasonality and unit roots at seasonal frequencies. I first review some basic theory about unit roots at seasonal frequencies (section 2); then I describe the new `sroot` command, which performs a formal test for unit roots in quarterly data (section 3); and then I give some advice for the applied researcher based on some Monte Carlo simulations (section 4). In section 5, I use the `sroot` command to detect seasonal unit roots in the original series of consumption in the UK for the years 1955–2006.

2 Unit roots at various frequencies

The spectrum of a seasonal series has distinct peaks at seasonal frequencies $\omega_s = 2\pi j/s$, where $j = 1, \dots, s/2$ and s is the number of periods within a year. In particular, we deal with $s = 4$ because it is the most common case.

While there is consensus on the importance of seasonality, there is little agreement on its treatment. Indeed, there are several ways to handle seasonality, each implicitly making different assumptions about the process, namely, as if it is

- a purely *deterministic* seasonal process,
- a *stationary* seasonal process, or
- an *integrated* seasonal process.

In applied work, the general (incorrect) belief is that the three methodologies are equivalent. In fact, they imply a very different data-generating process, as discussed below.

In a purely deterministic seasonal process, the reference model for the conditional mean of the dependent variable, y , can be written as

$$y = \mathbf{x}\beta + \sum_{i=1}^3 \delta_i D_i$$

where y is a vector of dimension n ; $\mathbf{x}$ is an $n \times k$ matrix with the first column containing only ones; β is a vector of length k ; and each δ_i is the coefficient attached to the vector D_i , a dummy vector equal to 1 only in season i . This notation will be employed throughout the article.

A stationary seasonal process can be written as an autoregressive model,

$$\phi(L)y_t = \epsilon_t \tag{1}$$

with all the roots of $\phi(L)$ outside the unit circle (but some come in complex pairs). If $s = 4$, then a stationary seasonal process is $y_t = \rho L^4 y_t + \epsilon_t$, where L is the lag operator and $L^4 y_t = y_{t-4}$. If some of the roots lie on the unit circle, the process is an integrated seasonal process.

Continuing with $s = 4$, a seasonally integrated series can be further decomposed into

$$\begin{aligned} (1 - L^4) y_t &= \epsilon_t \\ &= (1 - L)(1 + L)(1 + L^2)y_t \end{aligned} \quad (2)$$

which shows that in seasonal processes, the roots of modulus 1 can be four and not only one, as for the classical case. Also, two of the roots will be complex. Properties of each root are very similar to those at zero frequency; in particular, shocks have a permanent effect on the seasonal pattern, and their variances increase linearly with time, but shocks are asymptotically uncorrelated with unit-root processes of other frequencies. To see this more formally, consider the process in (2) as a stochastic difference process (details are in Hylleberg et al. [1990]), whose homogeneous solutions are

$$\begin{aligned} s_{1,t} &= \sum_{j=0}^{t-1} \epsilon_{t-j} && \text{for zero-frequency root;} \\ s_{2,t} &= \sum_{j=0}^{t-1} (-1)^j \epsilon_{t-j} && \text{for two-cycle-per-year root;} \\ s_{3,t} &= \sum_{j=0}^{\text{int}\{(t-1)/2\}} (-1)^j \Delta \epsilon_{t-2j} && \text{for one-cycle-per-year root} \end{aligned} \quad (3)$$

By expanding each single component of (3), we can show that the variance of each frequency increases linearly with time [specifically, $V(s_{1,t}) = V(s_{2,t}) = V(s_{3,t}) = t\sigma^2$]. Using the same technique, we can show that covariances are zero for complete years of data when the series are excited by the same ϵ_t and, thus, that the series are uncorrelated; for example,

$$\begin{aligned} \text{cov}(s_1, s_2) &= \overbrace{(\epsilon_t + \epsilon_{t-1} + \epsilon_{t-2} + \epsilon_{t-3} + \dots)}^{s_{1,t}} \overbrace{(\epsilon_t - \epsilon_{t-1} + \epsilon_{t-2} - \epsilon_{t-3} + \dots)}^{s_{2,t}} \\ &= \sigma^2 - \sigma^2 + \sigma^2 - \sigma^2 + \dots \\ &= 0 \end{aligned}$$

Differently from what has been suggested by many practitioners, I argue that using not seasonally adjusted (NSA) data should always be considered. At least as a robustness check, one should perform all the analysis with both seasonally adjusted (SA) and NSA data. However, we showed that NSA data have more involved processes than SA data, particularly because of unit roots at seasonal frequencies.

In what follows, I analyze a formal test to study the presence of seasonal unit roots on a statistical basis, focusing on Hylleberg et al. (1990).

A general expression for seasonal processes combines the three seasonal processes and is compactly represented by

$$d(L)a(L)(y_t - \mu_t) = \epsilon_t$$

where the roots of $a(L) = 0$ lie outside the unit circle, the roots of $d(L) = 0$ lie on the unit circle, and $\mu_t = x\beta + \sum_{i=1}^3 \delta_i D_i$. It follows that stationary components of y are in $a(L)$, while deterministic seasonality is in μ_t when there are no seasonal unit roots in $d(L)$. The test by Hylleberg et al. (1990) studies this model and detects seasonal unit roots at different seasonal frequencies, as well as at zero frequency.

The methodology strongly relies on a Lagrangian polynomial expansion for $\phi(L)$ in (1). Applying this representation for quarterly data, Hylleberg et al. (1990) study

$$\phi(L)y_{4,t} = \pi_1 y_{1,t-1} + \pi_2 y_{2,t-1} + \pi_3 y_{3,t-2} + \pi_4 y_{3,t-1} + \epsilon_t \quad (4)$$

where

$$\begin{aligned} y_{1,t} &= (1 + L + L^2 + L^3) y_t \\ y_{2,t} &= -(1 - L + L^2 - L^3) y_t \\ y_{3,t} &= (1 - L^2) y_t \\ y_{4,t} &= (1 - L^4) y_t \end{aligned}$$

and π_i 's are coefficients for seasonal roots, which we test to establish the nature of seasonality. In particular, at root $1 - L$ the test is on coefficient $\pi_1 = 0$, at seasonal root $1 + L$ the test is on coefficient $\pi_2 = 0$, and finally, at seasonal roots $1 + L^2$ the test is joint on coefficients $\pi_3 = \pi_4 = 0$. For a unit root in a given frequency, the associated coefficient π_i is zero. If π_2 and either π_3 or π_4 are different from zero, there is no seasonal unit root. Similarly, if π_1 is also different from zero, the series has no unit roots at all. The natural alternative for these tests is stationarity, $\pi_1 < 0$ and $\pi_2 < 0$, respectively, for π_1 and π_2 , or that π_3 and π_4 are not jointly equal to zero. To consider all the possible cases in the three seasonal processes, (4) can be augmented in various directions, such as lagged values of y_4 or deterministic components, and consistently estimated by ordinary least squares.

Although I focus on quarterly data, this identical setup can be readily generalized to other cases frequently encountered in practice, like biannual data or monthly data (see Franses and Hobijn [1997]).

The asymptotic distribution of the estimator of the coefficients in (4) is nonstandard. Because the method is analogous to that of Dickey and Fuller (1979), the distribution theory for these tests can be extracted from Dickey and Fuller (1979) and Fuller (1976) for π_1 and π_2 , and from Dickey, Hasza, and Fuller (1984) for π_3 , if π_4 is assumed to be zero. The tests are asymptotically invariant with respect to nuisance parameters. According to Hylleberg et al. (1990, 224), the finite-sample results are well approximated by the asymptotic theory, and the tests have reasonable power against each of the specific alternatives. The intercept and trend in the model affect only the distribution of π_1 , whereas seasonal dummies affect only the distributions of π_2 , π_3 , and π_4 .

I would like to conclude this section with a natural extension of the seasonal unit root, i.e., a seasonal cointegration and seasonal vector error-correction model. There are several methods for testing and estimation of cointegration at seasonal frequencies (Lee [1995], Johansen and Schaumburg [1999], and Cubadda [2001], among others), and each deserves a specific treatment, which we leave for future extensions. However, a simpler approach goes back to Engle and Granger (1987) and is adapted to the seasonal case by Engle et al. (1993). This simpler approach is a two-step estimator that simply requires estimating the linear combination(s) of levels on data transformed to account for seasonality and, for seasonal vector error-correction model estimation, relies on the speed of convergence in the first step.

3 The `sroot` command

The increasing variety of time-series methods in Stata has increased the number of time-series users with Stata. Thanks to the simplicity of data management, the proposed command makes extensive use of Stata's routines for lag operators and the `regress` command. The syntax of `sroot` is

```
sroot varname [if] [in] [, noconstant trend season(varlist) regress
    lags(#) generate(string) residuals(string) ]
```

3.1 Options

`noconstant` suppresses the constant term (intercept) in the model and indicates that the process under the null hypothesis is a random walk without drift. `noconstant` may not be used with the `trend` or `season(varlist)` option.

`trend` specifies that a trend term be included in the associated regression and that the process under the null hypothesis is a random walk, perhaps with drift. This option may not be used with the `noconstant` option.

`season(varlist)` indicates that the process under the null hypothesis is a random walk augmented for seasonal dummies. It is possible that `varlist` contains only one word (in which case the command builds the dummies) or that `varlist` contains the full set of dummies (in which case the command drops the last quarter because of multicollinearity). This option may not be used with the `noconstant` option.

`regress` specifies that the associated regression table appear in the output. By default, the regression table is not produced.

`lags(#)` specifies the number of lagged difference terms to include in the covariate list.

`generate(string)` generates a set of variables adjusted for seasonal filtering.

`residuals(string)` generates a variable containing the residual terms.

4 Some practical issues

In this section, I give some advice for applied research. I first explore distinctive features of `sroot` with respect to an existing similar command (section 4.1), and then I give some practical guidelines useful in empirical applications (section 4.2).

4.1 Why a new command?

The `hegy4` command in Stata performs the Hylleberg et al. (1990).¹ The two commands (`hegy4` and `sroot`) have key distinctions that I briefly explore in this section. I conclude

1. I thank C. Baum for bringing to my attention a very similar routine (Baum and Sperling 2001).

that **hegy4** and **sroot** are similar; thus suggestions in section 4.2 will be valid for both procedures.

First of all, the default in **hegy4** is to run a sequential test for the proper number of lags. A simple (unreported) simulation reveals that in some circumstances, it could be inappropriate. I designed the simulation for 48, 100, and 200 observations (12, 25, and 50 years), and for parameters $\beta_i \Delta_4 y_{t-i} = \{-0.8, -0.4, -0.2, -0.02, 0, 0.02, 0.2, 0.4, 0.8\}$ with $i = 1, 2$, along with their combinations (notice that coefficients are exactly equal to zero sometimes). The performance of a sequential test increases with sample size and with the absolute values of coefficients. However, two remarks are needed: First, as either β_1 or β_2 approaches but is different from zero, the practice is questionable and only in a bunch of cases, all with 200 observations, is the lag selection 100% correct. With 48 observations, the operational tool performs poorly. Second, undoubtable advantages are when β_1 or β_2 is indeed zero. However, **hegy4** offers the option “**notest** [that] may be specified to suppress the lag length test and utilize the lags specified in the option in generating the test statistic” (Baum and Sperling 2001).

Even though the Hylleberg et al. (1990) results are unaffected by nuisance parameters, according to the experiments in the next section, in cases of uncertainty about the correct number of lags, specifying the **notest** option seems a more convincing approach.

The second, most important, difference is in the **generate()** option. Engle et al. (1993) show that it is possible to study seasonal cointegration starting from transformed variables. The interested reader is referred to that article for further details, but for what matters here, we can build stationary combinations from transformed nonstationary variables in levels and study a vector autoregression augmented for these components in a seasonal error-correction model. I view this as a key distinction to push efforts toward the new **sroot** command.

A very minor difference is the **regress** option for **sroot**.

Of course, the tests, *ceteris paribus*, give the very same numbers, as shown with the following example:

```
. sroot x_nsa, lag(1) trend season(quarter)
```

HEGY test for SEASONAL unit roots			Number of obs	=	203
Test	1% Critical		5% Critical	10% Critical	
Statistic	Value		Value	Value	
Z(t) - Fr 0	0.180	-4.050	-3.490	-3.180	
Z(t) - Fr 1/2	2.522	-3.520	-2.910	-2.600	
Z(t) - L.Ann.	0.437	-4.040	-3.410	-3.100	
Z(t) - Annual	0.286	-2.650	-1.920	-1.480	
Joint Annual	0.133	8.960	6.570	5.560	
All SEAS. fr.	2.170	.	5.890	5.100	
All freq.	1.644	.	6.380	5.610	

(Continued on next page)

```
. hegy4 x_nsa, lag(1) det(strend)

HEGY Quarterly seasonal unit root test for x_nsa

Number of observations : 203
Deterministic variables : Seasonal dummies + constant + trend
Lags tested: 1
Augmented by lags : 1
```

	Stat	5% critical	10% critical
t[Pi1]	0.180	-3.490	-3.180
t[Pi2]	2.522	-2.910	-2.600
t[Pi3]	0.437	-3.410	-3.100
t[Pi4]	0.286	-1.920	-1.480
F[3-4]	0.133	6.570	5.560
F[2-4]	2.170	5.890	5.100
F[1-4]	1.644	6.380	5.610

4.2 Some practical guidelines

The most troublesome practical issues with this seasonal unit-root test are related to the deterministic terms and to the appropriate number of lags. In particular, with respect to the deterministic terms, the important question is whether they should be included in the model specification; with respect to the lags, the important question is how many lags should be considered. I try to answer these questions in this section by using Monte Carlo experiments based on 5,000 repetitions and designs specified below.

Nevertheless, I strongly suggest that the researcher verify, case by case, that residuals have desired properties, through the `residuals()` option. For example, Baum and Sperling (2001) suggest the regression of the (generated) residuals on four lags and the original regressors under the rationale that if all the information has already been considered, the null hypothesis that all coefficients are jointly equal to zero should not be rejected; other useful checks could be performed on specific moments of the distribution of the residuals, like the third moment (skewness) and the fourth moment (kurtosis).

Finally, it should be clear that here I adopt an empirical approach; theoretical consequences can be found in Ghysels, Lee, and Noh (1994).

More deterministic terms is better than fewer deterministic terms

I first examine the importance of deterministic terms in the model specification. According to the common wisdom, I will conclude that in the empirical applications, in case of uncertainty, it is safer to include deterministic terms even when they are not in the true data-generating process (DGP), rather than vice versa, neglecting deterministic terms that, in fact, are in the true DGP.

I use the following experiment:

$$y_t = \rho y_{t-4} + \sum_{i=1}^4 \delta_i D_i + \gamma t + u_t \quad t = 1, \dots, T$$

where ρ determines the (non)stationarity of the model; D_i is for seasonal dummies, associated to parameters $\delta_i = (-0.1, 0.05, -0.05, 0.1)$, which correspond to a -10% annualized drop, followed by a 5% increase, a 5% decrease, and a 10% increase, respectively; γ is set equal to 0.05 ; and u_t is white noise. The correct lag length in the DGP is zero, but we carried on the test for different lags, from zero to four. Additional lags will be indicative of the consequences of controlling for more lags than are needed in the presence of deterministic components. This experiment, except for the time trend, is used in Ghysels, Lee, and Noh (1994).

In table 1, I report the share of rejection of the null hypothesis of the unit root. The table has two main parts: on the left-hand side, the true DGP does in fact contain unit roots at all frequencies ($\rho = 1$), and the right-hand side is stationary ($\rho = 0.85$; see Ghysels, Lee, and Noh [1994] for further details on this specific value). Each side is further differentiated: in one case, the model is misspecified because we neglect the presence of seasonal dummies, and in the other case, it is correctly specified because all the deterministic components are controlled for. For easier readability, we report only the share of stationary roots at zero frequency and jointly at all the frequencies. When unit roots are in the true DGP, the entry should be zero, whereas when the model is stationary, the entry should be one. I discuss these measures.

When the data contain unit roots at all frequencies, neglecting seasonal dummies has serious consequences on the share of rejection. Because we control, by default of the command, for the intercept, the consequences on frequency 0 will be attenuated (as expected from section 2). Nevertheless, when we consider the whole set of frequencies, the conclusion will be seriously biased. As suggested in Ghysels, Lee, and Noh (1994), the reason is that test statistics under this misspecification are functions of the unknown seasonal dummy coefficients.

When we consider stationary data, two main conclusions can be drawn from the table. First, when the data are stationary but we neglect the deterministic terms in the model specification, our conclusions about unit roots at zero frequency are biased toward the nonrejection of the null hypothesis, which (however) is false (technically, the power of the test against this misspecification is low). Second, once deterministic terms are considered, a correct specification of lag length is less important, as shown by the comparison of the first panel (labeled “Lag: 0”) with respect to lower panels. This specific aspect is elaborated upon in the next subsection.

Aside from the correct specification, the number of observations plays a critical role. A performance with fewer than 100 observations is unsatisfactory, whereas a performance with 200 or more observations is good.

Table 1. Consequences of neglecting deterministic components

Obser.	Unit root				Stationary			
	Misspec.		Correct		Misspec.		Correct	
	Fr 0	All	Fr 0	All	Fr 0	All	Fr 0	All
Lag: 0								
48	0.000	0.701	0.001	0.019	0.000	0.070	0.018	0.085
100	0.000	1.000	0.001	0.007	0.000	0.388	0.064	0.175
200	0.000	1.000	0.000	0.006	0.000	0.846	0.340	0.700
300	0.000	1.000	0.000	0.006	0.000	0.983	0.684	0.974
Lag: 1								
48	0.000	0.426	0.002	0.019	0.000	0.058	0.017	0.080
100	0.000	1.000	0.001	0.010	0.000	0.311	0.063	0.169
200	0.000	1.000	0.000	0.007	0.000	0.784	0.321	0.676
300	0.000	1.000	0.000	0.007	0.000	0.973	0.657	0.967
Lag: 2								
48	0.000	0.284	0.002	0.024	0.000	0.057	0.015	0.083
100	0.000	0.994	0.002	0.011	0.000	0.279	0.059	0.159
200	0.000	1.000	0.001	0.006	0.000	0.731	0.292	0.650
300	0.000	1.000	0.000	0.006	0.000	0.961	0.619	0.961
Lag: 3								
48	0.000	0.183	0.004	0.025	0.000	0.049	0.012	0.074
100	0.000	0.891	0.002	0.009	0.000	0.235	0.061	0.157
200	0.000	0.990	0.000	0.008	0.000	0.679	0.272	0.614
300	0.000	0.985	0.000	0.006	0.000	0.946	0.574	0.956
Lag: 4								
48	0.000	0.149	0.002	0.019	0.000	0.045	0.015	0.057
100	0.000	0.660	0.002	0.008	0.000	0.218	0.055	0.127
200	0.000	0.772	0.000	0.007	0.000	0.640	0.235	0.553
300	0.000	0.553	0.000	0.006	0.000	0.922	0.536	0.931

Of course, one can wonder what the consequences are of controlling for undue deterministic terms. The same experiment from above, without deterministic terms, supports the view that the impact of undue deterministic components in the model specification is rather limited. Indeed, the shares of rejection of nonstationary root are close between the misspecified model that controls for deterministic terms and the correctly specified models (table 2).

Table 2. Consequences of imposing deterministic components

Obser.	Unit root				Stationary			
	Misspec.		Correct		Misspec.		Correct	
	Fr 0	All	Fr 0	All	Fr 0	All	Fr 0	All
Lag: 0								
48	0.022	0.060	0.035	0.048	0.030	0.106	0.051	0.130
100	0.031	0.048	0.051	0.056	0.052	0.157	0.088	0.362
200	0.040	0.046	0.051	0.055	0.107	0.491	0.211	0.907
300	0.041	0.045	0.052	0.052	0.202	0.865	0.410	0.999
Lag: 1								
48	0.015	0.061	0.032	0.046	0.015	0.090	0.047	0.117
100	0.027	0.047	0.051	0.052	0.042	0.147	0.083	0.351
200	0.036	0.045	0.053	0.054	0.097	0.480	0.200	0.901
300	0.040	0.044	0.050	0.052	0.187	0.861	0.392	0.998
Lag: 2								
48	0.018	0.060	0.035	0.043	0.016	0.091	0.044	0.115
100	0.033	0.048	0.049	0.051	0.041	0.146	0.084	0.350
200	0.035	0.046	0.051	0.050	0.096	0.479	0.193	0.900
300	0.039	0.044	0.050	0.053	0.177	0.861	0.383	0.998
Lag: 3								
48	0.015	0.060	0.038	0.050	0.010	0.076	0.043	0.107
100	0.027	0.047	0.049	0.048	0.031	0.132	0.081	0.334
200	0.036	0.048	0.050	0.050	0.087	0.465	0.181	0.901
300	0.036	0.043	0.050	0.053	0.158	0.853	0.356	0.998
Lag: 4								
48	0.016	0.048	0.032	0.045	0.020	0.067	0.036	0.080
100	0.030	0.042	0.045	0.046	0.043	0.114	0.073	0.272
200	0.037	0.044	0.050	0.048	0.096	0.385	0.176	0.826
300	0.038	0.044	0.049	0.050	0.172	0.766	0.338	0.994

For these reasons, I strongly suggest controlling for deterministic terms when performing a seasonal unit-root test.

More lags is better than fewer lags

The second important issue is related to the appropriate specification of the number of lags. I will conclude that it is a less important decision than that about deterministic terms, but in case of uncertainty about the true DGP, it could be safer to control for more lags than for fewer lags.

This conclusion is based on the following experiment:

$$y_t = \rho y_{t-4} + \sum_{j=1}^2 \alpha_j \Delta y_{t-j} + u_t \quad t = 1, \dots, T$$

for various combinations of α_1 and α_2 . The correct lag length would be two, but we carried on the test for different lags, from zero to four. In table 3, we select only lag zero (i.e., fewer lags than needed), lag two (i.e., correctly specified), and lag four (i.e., more lags than needed).

The parameter ρ determines the (non)stationarity of the model. In table 3, on the left, the true DGP is nonstationary ($\rho = 1$), while on the right-hand side, it is stationary ($\rho = 0.85$).

The main message from table 3 is that a correct lag specification is less important than a correct specification of deterministic terms. As expected, the best performances are achieved when the lag is correctly specified, overall with 300 observations. However, under lag misspecification, controlling for more lags than are needed could be safer than controlling for fewer. Indeed, based on experiments in Ghysels, Lee, and Noh (1994, 425), adding lags beyond what is necessary could be understood as an attempt to control for possible moving-average components whose “bias shrinks as additional lags of the autoregressive terms are included in the model”. Finally, from table 3, the trade-off in the number of lags is clear, because adding lagged values reduces the *power* of the test, while the *size* suffers if too few parameters are included (Engle et al. 1993).

Although the evidence is not clear-cut, based on theoretical considerations, the practical guidance for the applied researcher in case of uncertainty is that it is safer to control for more lags than are needed.

Table 3. Consequences of neglecting lags; the model is stationary

Observ.	α_{t-1}	α_{t-2}	Unit root		Stationary	
			Fr 0	All	Fr 0	All
Lag: 0						
100	0.000	0.000	0.063	0.066	0.106	0.405
100	0.020	0.000	0.063	0.066	0.106	0.405
100	0.020	0.020	0.055	0.059	0.085	0.391
100	0.400	0.000	0.063	0.066	0.106	0.405
100	0.400	0.020	0.055	0.059	0.085	0.391
100	0.400	0.400	0.086	0.403	0.000	0.885
300	0.000	0.000	0.058	0.059	0.422	0.999
300	0.020	0.000	0.058	0.059	0.422	0.999
300	0.020	0.020	0.050	0.051	0.373	0.999
300	0.400	0.000	0.058	0.059	0.422	0.999
300	0.400	0.020	0.050	0.051	0.373	0.999
300	0.400	0.400	0.100	0.487	0.000	1.000
Lag: 2						
100	0.000	0.000	0.052	0.051	0.086	0.346
100	0.020	0.000	0.052	0.051	0.086	0.346
100	0.020	0.020	0.052	0.052	0.089	0.346
100	0.400	0.000	0.052	0.051	0.086	0.346
100	0.400	0.020	0.052	0.052	0.089	0.346
100	0.400	0.400	0.054	0.056	0.658	0.759
300	0.000	0.000	0.050	0.051	0.374	0.998
300	0.020	0.000	0.050	0.051	0.374	0.998
300	0.020	0.020	0.050	0.051	0.401	0.999
300	0.400	0.000	0.050	0.051	0.374	0.998
300	0.400	0.020	0.050	0.051	0.401	0.999
300	0.400	0.400	0.056	0.052	1.000	1.000
Lag: 4						
100	0.000	0.000	0.046	0.051	0.078	0.267
100	0.020	0.000	0.046	0.051	0.078	0.267
100	0.020	0.020	0.046	0.050	0.078	0.264
100	0.400	0.000	0.046	0.051	0.078	0.267
100	0.400	0.020	0.046	0.050	0.078	0.264
100	0.400	0.400	0.054	0.055	0.507	0.584
300	0.000	0.000	0.048	0.051	0.331	0.993
300	0.020	0.000	0.048	0.051	0.331	0.993
300	0.020	0.020	0.048	0.051	0.357	0.994
300	0.400	0.000	0.048	0.051	0.331	0.993
300	0.400	0.020	0.048	0.051	0.357	0.994
300	0.400	0.400	0.050	0.052	1.000	1.000

5 Example

In this section, I use the `sroot` command to test for the presence of the unit root at seasonal frequency for the series of consumption in the UK. The data are from the National Institute of Statistics for the years 1955–2006 on a quarterly basis. I first test for the presence of a unit root for NSA data:

```
. sroot x_nsa,lag(4) trend season(quarter) regress gen(pi1 pi2 pi3)
```

HEGY test for SEASONAL unit roots				Number of obs = 200	
Test	1% Critical	5% Critical	10% Critical		
Statistic	Value	Value	Value		
Z(t) - Fr 0	0.330	-4.050	-3.490	-3.180	
Z(t) - Fr 1/2	2.739	-3.520	-2.910	-2.600	
Z(t) - L.Ann.	1.014	-4.040	-3.410	-3.100	
Z(t) - Annual	-0.023	-2.650	-1.920	-1.480	
Joint Annual	0.514	8.960	6.570	5.560	
All SEAS. fr.	2.842	.	5.890	5.100	
All freq.	2.166	.	6.380	5.610	

x_nsa	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
x_nsa						
Freq.0	.0001847	.0005594	0.33	0.742	-.0009189	.0012882
Freq.1/2	.0461743	.0168582	2.74	0.007	.0129175	.0794311
L.Annual	.0154886	.0152743	1.01	0.312	-.0146435	.0456206
Annual	-.0003434	.0152415	-0.02	0.982	-.0304109	.0297241
LD.	.7711865	.0777132	9.92	0.000	.6178793	.9244938
L2D.	.0894235	.0952005	0.94	0.349	-.0983814	.2772285
L3D.	.1412389	.0953463	1.48	0.140	-.0468536	.3293314
L4D.	-.2178033	.0753951	-2.89	0.004	-.3665376	-.069069
_trend	.0117578	.0031245	3.76	0.000	.005594	.0179216
Q1	-.1432412	.1811381	-0.79	0.430	-.5005779	.2140954
Q2	-.0581299	.1837385	-0.32	0.752	-.4205965	.3043368
Q3	.1162713	.1815115	0.64	0.523	-.2418021	.4743447
_cons	-.3112752	.1896996	-1.64	0.103	-.6855015	.0629511

Because I specified `regress`, the result has two main pieces, i.e., the test in the upper panel and the regression table in the lower panel. Let's start from the lower panel for clarity. It is helpful to have a look at regression results because there are four important components. The first four regressors are crucial for the test statistics. The second component is the set of lagged values, which are included in an attempt to remove serial correlation in ϵ_{it} . Third are the deterministic components, namely, a `_trend` and a set of seasonal dummies. The set of seasonal dummies automatically drops the last quarter because of multicollinearity. The user may either specify the *varname* for quarter or specify directly for the complete set of dummies. Fourth, there is the constant term.

We are mainly interested in the upper panel, which is intrinsically tied to the lower panel. In particular, the t statistics of the variables `Freq.0–Annual` from the lower panel will be the very same numbers that we find in test statistics in the upper panel. However, because the distribution is nonstandard, we also report the critical values at

some sensible confidence level, namely, 1, 5, and 10%. Further, for frequency zero and frequency 1/2, we can rely on the significance level for single coefficients, whereas for **L. Ann.** and **Annual**, we should test that their coefficients are jointly equal to zero, as we do in the line **Joint Annual**, along with their own critical values. The last two lines test the joint hypotheses that all *seasonal* coefficients are zero, i.e., the presence of seasonal unit roots, and that all relevant unit root coefficients are zero, i.e., full set of unit roots at all frequencies. In these cases, critical values are available only for 5 and 10% confidence levels.

In what follows, I interpret these numbers.

According to the t statistics from **Freq.0**, we do not reject that π_1 is different from zero at a conventional confidence level. Equivalently, we cannot reject that the time series has a unit root at frequency zero. According to section 2, test statistics and critical values for this frequency could have been obtained from those already tabulated from the Dickey–Fuller test, and most importantly, the decision is based on the same rule.

For frequency $\pi/2$, we do not reject the presence of a (seasonal) unit root at, say, the 95% confidence level. This is because the alternative hypothesis concerning π_2 in (4) is stationarity, or $\pi_2 < 0$; thus values of the t statistic smaller than the critical values at the preferred confidence level reject the null hypothesis of unit root. Vice versa, values of the t statistic larger than the critical values at the preferred confidence level do not reject the null hypothesis of unit root. Here the t statistic is 2.739 against a critical value of -2.910 at a 5% confidence level, and thus we cannot reject the presence of a (seasonal) unit root.

In (4), we have the annual frequency and its lag, and in principle they can return contrasting results. However, from section 2 we know that results depend on the joint test on coefficients. Being an F -type statistic, we reject the null hypothesis in cases where the test statistic is larger than the critical value. For the example at hand, we cannot reject the unit root at the annual frequency based on the line **Joint Annual**.

The test for unit roots at all seasonal frequencies and the test for unit roots at all frequencies are also F -type; thus the decision is based on the same rule of the annual frequency. From the line **All SEAS. fr.**, we do not reject the joint significance of seasonal unit roots, and from the line **All freq.**, the joint significance of the full set of unit roots, at seasonal and nonseasonal frequencies.

The evidence indicates that UK consumption has a unit root at frequency zero, as could be inferred from the classical Dickey–Fuller test. The new **sroot** command indicates that there are two more roots, one at frequency 1/2 (or biannual) and the other at annual frequency. **hegy4** returns the same qualitative conclusions. In general, **hegy4** and **sroot** test statistics need not be equal because **hegy4** uses an automatic lag selection method unless **notest** is specified. In the case at hand, the sequential tests on lags of the dependent variable select only lags 1 and 4:

```
. hegy4 x_nsa, lag(1 2 3 4) det(strend)

HEGY Quarterly seasonal unit root test for x_nsa

Number of observations : 200
Deterministic variables : Seasonal dummies + constant + trend
Lags tested: 1 2 3 4
Augmented by lags : 1 4
```

	Stat	5% critical	10% critical

t[Pi1]	0.354	-3.490	-3.180
t[Pi2]	2.812	-2.910	-2.600
t[Pi3]	0.589	-3.410	-3.100
t[Pi4]	0.414	-1.920	-1.480
F[3-4]	0.254	6.570	5.560
F[2-4]	2.731	5.890	5.100
F[1-4]	2.084	6.380	5.610

For comparison purposes, we repeat the Hylleberg et al. (1990) test for SA data. As expected, they have only one unit root, found at frequency zero:

```
. sroot x_sa, lag(4) trend season(quarter)

HEGY test for SEASONAL unit roots

      Test      1% Critical      Number of obs      =      200
      Statistic      Value      5% Critical      10% Critical
                                Value      Value      Value
-----
Z(t) - Fr 0      0.290      -4.050      -3.490      -3.180
Z(t) - Fr 1/2    -4.784      -3.520      -2.910      -2.600
Z(t) - L.Ann.    -5.310      -4.040      -3.410      -3.100
Z(t) - Annual    -4.676      -2.650      -1.920      -1.480
Joint Annual     28.976      8.960      6.570      5.560
All SEAS. fr.    32.859      .      5.890      5.100
All freq.        24.734      .      6.380      5.610
```

The actual existence of seasonal unit roots in the series of consumption sheds more light on the potentially dramatic impact that a suboptimal econometric technique can have on a test of an economic theory. In this sense, the results from `sroot` are important per se. However, we can go a step further.

In particular, the `generate()` option is crucial to consider one possible extension of the unit root at seasonal frequencies, namely, cointegration at seasonal frequencies. Indeed, the option stores three different variables obtained from the transformation employed in the procedure. We just plot the transformed series in figure 1 as they are generated by `sroot` (i.e., with no editing adjustment). Although not pursued here, notice that the option allows the replication of the procedure by Engle et al. (1993) to fit a seasonal vector error-correction model.

Figure 1. Series of consumption and its transformations

6 Conclusion

In this article, I presented the new `sroot` command, which implements a test to detect unit roots at frequencies other than zero, in quarterly data. The motivation for the new command is that many time series may have seasonal unit roots. Although the usual practice is to work with seasonally adjusted data, I view this as a weak solution because fluctuations do contain information and because adjustments can be responsible for rejection of economic theories even though the underlying model is correct. I argue that one should always consider using seasonally unadjusted data, which can be characterized by seasonal unit roots. It is important to go beyond the classical test at frequency zero, as I propose with `sroot`, paying much attention to the model specification. Finally, a promising extension is cointegration at seasonal frequencies that can be studied by exploiting the `generate()` option in `sroot`, even though more efficient methods are available in the literature.

7 Acknowledgments

I would like to thank Christopher Baum for stimulating suggestions and comments and Lucia Corno for comments on earlier drafts of the article. I also thank an anonymous referee and the editor for their suggestions. The article was written while I was completing my PhD at the Università degli Studi di Roma “Tor Vergata”, where I benefited from the stimulating environment. All errors are mine.

8 References

- Baum, C. F., and R. Sperling. 2001. HEGY4: Stata module to compute Hylleberg et al. seasonal unit root test. Statistical Software Components, Boston College Department of Economics. <http://ideas.repec.org/c/boc/bocode/s416502.html>.
- Cubadda, G. 2001. Complex reduced rank models for seasonally cointegrated time series. *Oxford Bulletin of Economics and Statistics* 63: 497–511.
- Dickey, D. A., and W. A. Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74: 427–431.
- Dickey, D. A., D. P. Hasza, and W. A. Fuller. 1984. Testing for unit roots in seasonal time series. *Journal of the American Statistical Association* 79: 355–367.
- Engle, R. F., and C. W. J. Granger. 1987. Co-integration and error correction: Representation, estimation, and testing. *Econometrica* 55: 251–276.
- Engle, R. F., C. W. J. Granger, S. Hylleberg, and H. S. Lee. 1993. Seasonal cointegration: The Japanese consumption function. *Journal of Econometrics* 55: 275–298.
- Franses, P. H., and B. Hobijn. 1997. Critical values for unit root tests in seasonal time series. *Journal of Applied Statistics* 24: 25–48.
- Fuller, W. A. 1976. *Introduction to Statistical Time Series*. New York: Wiley.
- Ghysels, E., H. S. Lee, and J. Noh. 1994. Testing for unit roots in seasonal time series: Some theoretical extensions and a Monte Carlo investigation. *Journal of Econometrics* 62: 415–442.
- Hylleberg, S., R. F. Engle, C. W. J. Granger, and B. S. Yoo. 1990. Seasonal integration and cointegration. *Journal of Econometrics* 44: 215–238.
- Johansen, S., and E. Schaumburg. 1999. Likelihood analysis of seasonal cointegration. *Journal of Econometrics* 88: 301–339.
- Lee, L.-F. 1995. Semiparametric maximum likelihood estimation of polychotomous and sequential choice models. *Journal of Econometrics* 65: 381–428.
- Miron, J. A. 1986. Seasonal fluctuations and the life cycle—permanent income model of consumption. *Journal of Political Economy* 94: 1258–1279.
- Osborn, D. R. 1988. Seasonality and habit persistence in a life cycle model of consumption. *Journal of Applied Econometrics* 3: 255–266.

About the author

Domenico Depalo received his PhD in econometrics and empirical economics from the Università degli Studi di Roma “Tor Vergata”. After attending a postdoctoral program at the Sapienza–Università di Roma, he is currently a junior researcher at the Bank of Italy. The views expressed in the article are those of the author and do not involve the responsibility of the bank.