The Detection of Cycles in Raw and Seasonally Adjusted Data

Tucker McElroy* Scott Holan†

Abstract
The detection and estimation of business cycles in economic time series is an important activity of econometricians, and typically involves the filtering of one or more seasonally adjusted time series. The community of econometricians favoring univariate model-based approaches to cycle estimation seeks to avoid the identification of spurious cycles via taking a data-driven approach, which is in contrast to nonparametric band-pass approaches. However, given that seasonal adjustment is a procedure that greatly affects all frequencies of the raw data, it is natural to ask the following questions: can cycles be adequately detected from raw data? If so, are the detection rates superior to those obtained from seasonally adjusted data, and does this question depend on the method of adjustment? Does seasonal adjustment generate spurious cycles? This paper seeks to provide statistical methodology that can be used to answer these queries. We introduce a diagnostic statistic for deciding the inclusion or exclusion of an unobserved component, such as a cycle, and determine its theoretical properties. We then describe how this can be used to address our research questions in a rigorous fashion, and how currently available tools are not adequate.

Key Words: ARIMA models, Goodness-of-fit, Signal extraction, Unobserved components

1. Introduction

The detection and estimation of business cycles in economic time series is an important activity of econometricians, and typically involves the filtering of one or more seasonally adjusted time series. The community of econometricians favoring univariate model-based approaches to cycle estimation seeks to avoid the identification of spurious cycles via taking a data-driven approach (Harvey (1989) and Harvey and Trimbur (2003)), which is in contrast to nonparametric band-pass approaches (Hodrick and Prescott (1997) and Baxter and King (1999)). However, given that seasonal adjustment is a procedure that greatly affects all frequencies of the raw data, it is natural to ask the following questions: can cycles be adequately detected from raw data? If so, are the detection rates superior to those obtained from seasonally adjusted data, and does this question depend on the method of adjustment? Does seasonal adjustment generate spurious cycles? This paper provides statistical tools to answer these queries.

We first observe that these questions are irrelevant if the analyst adopts the nonparametric framework, since by definition the cycle is merely the output of some user-specified filter, such as the Hodrick-Prescott (HP) or an ideal band-pass. In this paper we focus on the model-based approach to cycle identification and estimation, which has the advantage of not producing a cycle when the data does not warrant one (see Harvey and Jaeger, 1993). With this model-based philosophy, a cycle is viewed as a certain stochastic Unobserved Component (UC), and not merely a certain band of frequencies in the data's spectrum. We mention in passing the method of Kaiser and Maravall (2005), which marries the model-based approach to the HP filter, essentially finding an implicit cycle model through the use of a user-defined HP filter on seasonally adjusted data. Note that this method does not have the capacity to declare that a cycle is absent from the data.

*Statistical Research Division, U.S. Census Bureau, 4600 Silver Hill Road, Washington, D.C. 20233-9100, tucker.s.mcelroy@census.gov

†Department of Statistics,University of Missouri-Columbia, 146 Middlebush Hall, Columbia, MO, 65211-6100, holans@missouri.edu

On the topic of UC detection, one can reject the hypothesis that a certain UC is *not* present via considering whether it is plausible, given the data, that the UC's innovation variance is zero (see Harvey, 1989). Alternatively, one can use a likelihood ratio test; either way, this is a test of inclusion, since low p-values indicate rejection of the null hypothesis and hence inclusion of the UC (see the treatment of likelihood ratio tests and similar procedures in Taniguchi and Kakizawa, 2000). These tests are point hypotheses, i.e., they are concerned with parameter values, and not composite hypotheses such as goodness-of-fit (gof) tests, which are concerned with model specifications (although they do take parameter uncertainty into account). UC gof testing was briefly discussed in McElroy and Holan (2009), and is further developed here for tests of both inclusion and exclusion of UCs in Section 2. This additional flexibility, as well as the ability to test model specifications, makes these UC gof tests a natural tool for studying our questions.

This paper does not include numerical or empirical studies; instead we exposit the theoretical aspects of our approach. We have in mind seasonal economic time series that are thought to have salient cyclical effects (e.g., Housing Starts). Other series with a negligible cycle (for which economists can nevertheless produce a cycle with nonparametric filters such as the HP) could also be examined with our tests, but in this case the inclusion of a cycle UC is likely to be rejected. In the next section we develop the methodology and mathematical results that are needed, and explain how these can be used to assess the presence of cycles in seasonal data.

2. Statistical Identification of Unobserved Components

In this section we draw heavily on the concepts and notation of McElroy and Holan (2009). We have in mind a Null model described by the class of spectra $\mathcal{F} = \{f_\theta\}_{\theta \in \Theta}$, and an Alternative model described by another class $\mathcal{G} = \{f_\xi\}_{\xi \in \Xi}$. Also, we consider two UCs, a signal process S_t and a noise process N_t. The true spectral density is denoted $\tilde{f}$, but it is unknown whether $\tilde{f} \in \mathcal{F}$ or $\tilde{f} \in \mathcal{G}$. Then our hypotheses are

$$H_0 : \tilde{f} \in \mathcal{F}$$
$$H_a : \tilde{f} \in \mathcal{G}.$$

We assume that there are unique pseudo-true values $\tilde{\theta} \in \Theta$ and $\tilde{\xi} \in \Xi$, which are the respective minimizers of the Kullback-Leibler distance from $\tilde{f}$ to the model class. Then $\tilde{f}$ equals $f_{\tilde{\theta}}$ or $f_{\tilde{\xi}}$, depending on H_0 or H_a respectively. Parameters are obtained through maximum likelihood estimates (MLEs), which under certain conditions are asymptotically equivalent to the minimizers of the Kullback-Leibler distance from the periodogram to the model class. These MLEs are denoted $\widehat{\theta}$ and $\widehat{\xi}$.

2.1 UC Models and a General Result

The way we have formulated things, we have a signal process in mind and there is the question of whether to include or exclude a noise process, which is a second UC. Of course, there may be plenty of applications of UC testing that do not involve signal extraction, but the use of signal and noise as terminology is helpful for identification. We suppose that the signal process is potentially nonstationary with differencing operator $\delta^S(z)$, such that $U_t = \delta^S(B) S_t$ is stationary with mean zero (here B is the backshift operator). Likewise, the noise process (if it exists) has differencing operator $\delta^N(z)$ with $V_t = \delta^N(B) N_t$. Clearly if an operator is just the identity, we recover stationarity as a special case. As usual, it is assumed that no unit roots are common to $\delta^S(z)$ and $\delta^N(z)$; let $\delta(z) = \delta^S(z)\delta^N(z)$. Let $W_t = \delta(B) Y_t$, which is stationary.

Now in the case that the noise process is absent, we have $Y_t = S_t$ and $W_t = U_t$ (note that $\delta = \delta^S$ since $\delta^N = 1$). Then we have the simple relation

$$f_W(\lambda) = f_U(\lambda).$$

However, if the noise process is present, we have $Y_t = S_t + N_t$ and $W_t = \delta^N(B) U_t + \delta^S(B) V_t$, and now

$$f_W(\lambda) = |\delta^N(e^{-i\lambda})|^2 f_U(\lambda) + |\delta^S(e^{-i\lambda})|^2 f_V(\lambda). \quad (1)$$

When we refer to the model spectral density (either f_θ or f_ξ, depending on the model), this is the spectrum of the differenced data, i.e., f_W. If the noise process is absent, the parameters (either θ or ξ) enter simply through the spectral density f_U. But if the noise is present, then the parameters enter in through both f_U and f_V according to (1).

The idea of the UC gof tests of this paper is to adapt the framework of McElroy and Holan (2009), where a diagnostic statistic of the form $Q_n(I, g, \xi)$ is considered, where I is the periodogram (of the differenced data), g is a weighting function depending on ξ, and $Q_n(f, g, \xi) = n^{-1} \sum_\lambda f(\lambda) g_\xi(\lambda)$, with the summation extending over Fourier frequencies only. In Section 3 of that paper an extension of the main result to UC testing is discussed; here we restate the main theorem in this context:

Theorem 1 *Under the same conditions as Theorem 2 of McElroy and Holan (2009), we have*

$$\left\{ \sqrt{n} \left(Q_n(I^{j_i}, g_i, \hat{\xi}) - j_i! Q_n(f_{\hat{\theta}}^{j_i}, g_i, \hat{\xi}) \right) + \sqrt{n} \frac{j_i!}{2\pi} \int_{-\pi}^{\pi} g_{\tilde{\xi},i}(\lambda) (f_{\tilde{\theta}}^{j_i}(\lambda) - \tilde{f}^{j_i}(\lambda)) \, d\lambda \right\}_{i=1}^{L} \tag{2}$$
$$\stackrel{\mathcal{L}}{\Longrightarrow} \mathcal{N}(0, W(\widetilde{\theta}, \widetilde{\xi}))$$

as $n \to \infty$, with $W(\theta, \xi)$ an $L \times L$ variance matrix with klth entry

$$\begin{aligned} W_{kl}(\theta, \xi) = & \frac{(j_k + j_l)! - j_k! j_l!}{4\pi} \int_{-\pi}^{\pi} (g_{\xi,k}(\lambda) g_{\xi,l}(-\lambda) + g_{\xi,l}(\lambda) g_{\xi,k}(-\lambda) + 2 g_{\xi,k}(\lambda) g_{\xi,l}(\lambda)) \, \tilde{f}^{j_k + j_l}(\lambda) \, d\lambda \\ & + \frac{(j_k + 1)! - j_k!}{4\pi} \int_{-\pi}^{\pi} (g_{\xi,k}(\lambda) p_{\theta,\xi,l}(-\lambda) + p_{\theta,\xi,l}(\lambda) g_{\xi,k}(-\lambda) + 2 g_{\xi,k}(\lambda) p_{\theta,\xi,l}(\lambda)) \, \tilde{f}^{j_k + 1}(\lambda) \, d\lambda \\ & + \frac{(j_l + 1)! - j_l!}{4\pi} \int_{-\pi}^{\pi} (g_{\xi,l}(\lambda) p_{\theta,\xi,k}(-\lambda) + p_{\theta,\xi,k}(\lambda) g_{\xi,l}(-\lambda) + 2 p_{\theta,\xi,k}(\lambda) g_{\xi,l}(\lambda)) \, \tilde{f}^{j_l + 1}(\lambda) \, d\lambda \\ & + \frac{1}{4\pi} \int_{-\pi}^{\pi} (p_{\theta,\xi,k}(\lambda) p_{\theta,\xi,l}(-\lambda) + p_{\theta,\xi,l}(\lambda) p_{\theta,\xi,k}(-\lambda) + 2 p_{\theta,\xi,k}(\lambda) p_{\theta,\xi,l}(\lambda)) \, \tilde{f}^2(\lambda) \, d\lambda. \end{aligned}$$

These entries are defined in terms of the following quantities:

$$\begin{aligned} p_{\theta,\xi,i}(\lambda) &= -j_i! f_\theta^{-2}(\lambda) b_i'(\theta, \xi) M_f^{-1}(\theta) \nabla_\theta f_\theta(\lambda) \\ b_i(\theta, \xi) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f_\theta^{j_i}(\lambda) - \widetilde{f}^{j_i}(\lambda)) \nabla_\xi g_{\xi,i}(\lambda) + j_i g_{\xi,i}(\lambda) f_\theta^{j_i - 1}(\lambda) \nabla_\theta f_\theta(\lambda) \, d\lambda \\ M_f(\theta) &= \nabla_\theta \nabla_\theta' D(f_\theta, \tilde{f}). \end{aligned}$$

When all the $g_{\theta,i}$ functions are even and H_0 holds, the variance formulas simplify to

$$\begin{aligned} W_{kl}(\theta, \xi) &= \frac{(j_k + j_l)! - j_k! j_l!}{\pi} \int_{-\pi}^{\pi} r_{\theta,\xi,k}(\lambda) r_{\theta,\xi,l}(\lambda) \, d\lambda - 2 j_k! j_l! b_k'(\theta, \xi) M_f^{-1}(\theta) b_l(\theta, \xi) \\ h_\theta(\lambda) &= \nabla_\theta \log f_\theta(\lambda) = \frac{\nabla_\theta f_\theta(\lambda)}{f_\theta(\lambda)} \\ r_{\theta,\xi,i}(\lambda) &= g_{\xi,i}(\lambda) f_\theta^{j_i}(\lambda). \end{aligned}$$

Remark 1 The expression for b_i differs slightly ($\widetilde{f}$ and f_θ are swapped) from that in McElroy and Holan (2009), which is in error.

In order to form a normalized statistic under H_0, we see that the asymptotic variance matrix can be computed as if H_0 were true; denote this quantity by W_{kl}^0. Then the MLEs $\widehat{\theta}$ and $\widehat{\xi}$ can be substituted for the unknown parameters in all of the above formulas. As in McElroy and Holan (2009), this will provide a consistent estimate of the variance. In this article we focus on the linear case ($L = 1$ and $j_1 = 1$) and the quadratic case ($L = 2$ and $j_1 = 2, j_2 = 1$). The linear and quadratic diagnostics, appropriately normalized, are given by

$$\sqrt{n}\,\frac{Q_n(I, f_\xi^{-1}, \widehat{\xi}) - Q_n(f_{\widehat{\theta}}, f_\xi^{-1}, \widehat{\xi})}{\sqrt{W_{11}^0(\widehat{\theta}, \widehat{\xi})}}$$

$$\sqrt{n}\,\frac{Q_n(I^2, f_\xi^{-2}, \widehat{\xi}) - 2Q_n(I, f_\xi^{-1}, \widehat{\xi}) - 2Q_n(f_{\widehat{\theta}}^2, f_\xi^{-2}, \widehat{\xi}) + 2Q_n(f_{\widehat{\theta}}, f_\xi^{-1}, \widehat{\xi})}{\sqrt{W_{11}^0(\widehat{\theta}, \widehat{\xi}) - 4W_{12}^0(\widehat{\theta}, \widehat{\xi}) + 4W_{22}^0(\widehat{\theta}, \widehat{\xi})}}.$$

2.2 GOF Tests for UC Inclusion and Exclusion

The gof test for inclusion of a UC has the Null model corresponding to just the signal process, so that rejection of H_0 entails acceptance of $\mathcal{G}$, which corresponds to a model including the second UC (the noise process). Hence, the inclusion test provides p-values that represents the probability that the second UC should not have been included (i.e., Type I error). Since, as discussed above, the model classes correspond to a stationary spectrum of the differenced data process, we have

$$f_\theta(\lambda) = f_U(\lambda)$$
$$f_\xi(\lambda) = \left|\delta^N(e^{-i\lambda})\right|^2 f_U(\lambda) + \left|\delta^S(e^{-i\lambda})\right|^2 f_V(\lambda).$$

Similarly, the gof test for exclusion of a UC has the Null model corresponding to signal plus noise, so that rejection of H_0 entails rejection of the presence of the noise process. Then

$$f_\theta(\lambda) = \left|\delta^N(e^{-i\lambda})\right|^2 f_U(\lambda) + \left|\delta^S(e^{-i\lambda})\right|^2 f_V(\lambda)$$
$$f_\xi(\lambda) = f_U(\lambda).$$

In this way gof test for inclusion and exclusion of a UC can be handled. We next discuss two examples, which are central to the motivation for this article.

Scenario 1 Suppose that the signal process follows a Box-Jenkins Airline model, and the noise process is a (stationary) stochastic cycle C_t with dynamics

$$(1 - 2\rho\cos\omega B + \rho^2 B^2)C_t = \epsilon_t, \tag{3}$$

where ϵ_t is white noise with variance σ_C^2. The strength of the cycle is parametrized through ρ, while the frequency location is governed by ω. The airline process A_t has dynamics

$$(1 - B)(1 - B^{12})A_t = (1 - \eta_1 B)(1 - \eta_2 B^{12})\zeta_t, \tag{4}$$

with ζ_t white noise with variance σ_A^2. Then in the inclusion test, the null model is just the airline model so that $\theta = (\eta_1, \eta_2, \sigma_A^2)'$. The alternative model consists of the sum of the airline and cycle processes, so that

$$f_W(\lambda) = \left|1 - \eta_1 e^{-i\lambda}\right|^2 \left|1 - \eta_2 e^{-i12\lambda}\right|^2 \sigma_A^2 + \frac{\left|1 - e^{-i\lambda}\right|^2 \left|1 - e^{-i12\lambda}\right|^2 \sigma_C^2}{(1 - 2\rho\cos(\lambda + \omega) + \rho^2)(1 - 2\rho\cos(\lambda - \omega) + \rho^2)}. \tag{5}$$

In this case $\xi = (\eta_1, \eta_2, \sigma_A^2, \rho, \omega, \sigma_C^2)'$. But with the exclusion test, the role of θ and ξ are swapped.

Scenario 2 Now suppose that the signal process is a trend irregular T_t with an IMA(2,2) specification:

$$(1-B)^2 T_t = (1 + \vartheta_1 B + \vartheta_2 B^2)\kappa_t,$$

with white noise variance σ_T^2. The noise process is the cycle C_t described in Example 1 above. Then in the inclusion test we have the null spectral density given by

$$|1 + \vartheta_1 e^{-i\lambda} + \vartheta_2 e^{-i2\lambda}|^2 \sigma_T^2,$$

and $\theta = (\vartheta_1, \vartheta_2, \sigma_T^2)'$. But the alternative spectral density for the differenced data is

$$|1 - \vartheta_1 e^{-i\lambda} - \vartheta_2 e^{-i2\lambda}|^2 \sigma_T^2 + \frac{|1 - e^{-i\lambda}|^4 \sigma_C^2}{(1 - 2\rho\cos(\lambda+\omega) + \rho^2)(1 - 2\rho\cos(\lambda-\omega) + \rho^2)}.$$

Then we have $\xi = (\vartheta_1, \vartheta_2, \sigma_T^2, \rho, \omega, \sigma_C^2)'$. In the case of the exclusion test θ and ξ are swapped.

2.3 Seasonal Adjustment Methods

In this article we consider model-based Seasonal Adjustment (SA), which is a signal extraction problem where the noise process is the seasonal component. Since we are interested in testing for the presence of a cyle UC before and after SA, we consider the following two SA procedures:

1. Supposing that the cycle UC is present, we fit a model of the form Airline plus Cycle, where parameter estimates are obtained using a structural approach, i.e., the spectral density given in (5). A canonical decomposition is then performed on the airline model, so that a model for the seasonal is obtained; then model-based signal extraction is carried out, with a trend plus cycle plus irregular as the signal, and with the seasonal as the noise.

2. Supposing that no cycle UC is present, we fit an Airline model and perform the canonical decomposition in order to obtain a model for the seasonal. Then model-based signal extraction is carried out, with a trend plus irregular as the signal, and with the seasonal as the noise.

Note that UC gof tests for inclusion/exclusion of the cycle can be done on the original (linearized) data, as described in Scenario 1; then we can seasonally adjust according to method 1 or 2 above, as indicated by the results of the test; and finally we can again test for existence of the cycle in the SA data, following the approach of Scenario 2. This provides several different scenarios, so that we can investigate the following questions: are cycles present in the raw data also present in the SA? Or is cycle detection easier after SA? How does this depend on the method of SA, i.e., whether the presence of the cycle is taken into account?

We briefly justify our hybrid approach to modeling the UCs, which really combines the structural approach and the canonical decomposition approach. A standard canonical decomposition approach would place an AR(2) factor in the model for the raw data, and attempt to decompose into four components; however, in practice it is difficult to get good fits for such SARIMA models (in our experience), and the decomposition is no longer guaranteed to exist. On the other hand, decompositions frequently exist for airline models, at least within the scope of parameter values typically encountered with economic time series. As for a fully structural approach, this would postulate four UCs (trend, seasonal, irregular, and cycle) with potentially 19 parameters (assuming an IMA(2,2) specification for trend and an MA(11) model for the differenced seasonal). Although this number can be reduced by selecting more parsimonious models for the seasonal (one extreme dictates that the differenced seasonal should be white noise, although this produces a very chaotic seasonal pattern that is unlikely to fit data with stable seasonality), there are still more parameters than the case of the hybrid approach. But with the hybrid approach, the model for the seasonal is essentially governed by the parameter η_2 of the airline UC. We do not argue that the hybrid method is the best, only that it seems to be an interesting and workable alternative to two widely used techniques.

In the practice of seasonal adjustment we would like to perform signal extraction with simpler models, if possible. Therefore we would like to exclude the cycle component in the modeling stage, if this is statistically justifiable. Models without a cycle correspond (essentially) to current practice at the U.S. Census Bureau, and therefore there is an institutional incentive to utilize the tests of exclusion rather than those of inclusion. Note that in the event that the presence of a cycle is rejected with significance, it is not saying that no cycle exists in the nonparametric sense of the HP filter, but rather that the particular ARIMA model for that UC has no statistical support. Any minor cyclical effects will then be incorporated into the models for the other components, but in any event have little impact on seasonal adjustment results.

3. Conclusion

This paper has set out some questions regarding the investigation of cycles in seasonal economic data, utilizing a model-based perspective. We have presented tests for the inclusion and exclusion of unobserved components in a mathematically rigorous fashion, and specialized to the case of seasonal-trend-cycle data in our examples. These new tests allow for the exclusion of a cycle component, which cannot be done using likelihood ratio tests (or similar statistics such as the Wald test). This is accomplished by directly comparing the fits of both models in the frequency domain, although all quantities can actually be computed in the time domain using simple formulas (omitted to save space). Future work will focus on numerical studies of these new tests, and how they can be used to answer our leading questions.

REFERENCES

Baxter, M. and King, R. (1999), "Measuring business cycles: approximate bandpass filters for economic time series," *Review of Economics and Statistics*, 81, 575–593.

Harvey, A. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter,* Cambridge, UK: Cambridge University Press.

Harvey, A. and Jaeger, A. (1993), "Detrending, stylized facts and the business cycle," *Journal of Applied Econometrics*, 8, 231–247.

Harvey, A. and Trimbur, T. (2003), "General model-based filters for extracting cycles and trends in economic time series," *Review of Economics and Statistics*, 85, 244–255.

Hodrick, R. and Prescott, E. (1997), "Postwar U.S. business cycles: an empirical investigation," *Journal of Money, Credit, and Banking*, 29, 1–16.

Kaiser, R. and Maravall, A. (2005), "Combining filter design with model-based filtering: an application to business-cycle estimation," *International Journal of Forecasting*, 21, 691–710.

McElroy, T. and Holan, S. (2009), "A local spectral approach for assessing time series model misspecification," *Journal of Multivariate Analysis*, 100, 604–621.

Taniguchi, M. and Kakizawa, Y. (2000), *Asymptotic Theory of Statistical Inference for Time Series,* New York: Springer-Verlag.