Queen's Economics Department Working Paper No. 901

Measuring Business Cycles with Business-Cycle Models

Allan W. Gregory   Gregor W. Smith

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

5-1994

Measuring Business Cycles with Business-Cycle Models

Allan W. Gregory and Gregor W. Smith†

May 1994

Abstract

Business cycles may be defined or measured by parametrizing detrending filters to maximize the ability of a business-cycle model to match the moments of the remaining cycles. Thus a theory can be used to guide cycle measurement. We present two applications to U.S. postwar data. In the first application the cycles are measured with a standard, real business cycle model. In the second, they are measured using information on capacity utilization and unemployment rates. Simulation methods are used to describe the properties of the GMM estimators and to allow exact inference.

Keywords: business cycles, detrending

JEL Classification Codes: E32, C32, C15

† Department of Economics, Queen's University, Kingston Ontario Canada K7L 3N6. For helpful criticism we thank colleagues and seminar participants at Queen's, UBC, York, Princeton, the Bank of Canada, Carnegie Mellon, and the meetings of the Canadian Economics Association and of the Econometric Society. We acknowledge the financial support of the Social Sciences and Humanities Research Council of Canada and of the Foundation for Educational Exchange between Canada and the USA. Masao Ogaki graciously provided us with the Hansen-Heaton-Ogaki GMM code, written in GAUSS TM and supported by the National Science Foundation.

1. Introduction

While modern business-cycle theory has been used to answer various quantitative questions about cyclical fluctuations, it typically takes cycles themselves as given. Researchers generally study macroeconomic time series which have been detrended in some way (for example by linear detrending, calculation of growth rates, or application of the Hodrick-Prescott filter). Thus the model is not used to measure or define the business-cycle component of the series. This study offers an example of using business cycle theory for this purpose.

Many macroeconomists would argue that a cycle-trend decomposition should follow from a theoretical model which includes growth trends. But in practice much research focuses on fluctuations at business-cycle frequencies only, and many models deal with cycles without modeling growth. It seems reasonable to avoid arbitrariness in measuring those fluctuations, and using the cycle model under study to measure cycles is one way of doing this. For example, an important feature of regular fluctuations is the differential variability of inputs, output and its components; the statistical definition of cycles can be based on this feature.

Another view is that detrending is not a central issue as long as many researchers use the same detrending method. From this perspective, detrending is indeed arbitrary (as opposed to being based on a growth model) but business-cycle statistics are defined around some trend in a way which is consistent across studies. In practice, though, models are modified in response to discrepancies between theoretical moments and historical moments constructed from detrended data. Canova (1991), Cogley and Nason (1992a), and King and Rebelo (1993) have shown that these discrepancies may be very different, depending on the detrending method used. Two researchers might reach quite different conclusions about the same model, though each has studied well-defined statistics, because they use different detrending methods.

The main illustrative example in this paper calculates the trend or low-frequency component which minimizes the discrepancies (using a standard method-of-moments metric) between the properties of detrended, historical data and those of a business-cycle model. We estimate the trend which is optimal in this sense, and compare it to some conventional methods. We also use the calculated trend to measure postwar U.S. business cycles.

This method of detrending may be contrasted with methods which identify trends by restricting their correlation with cycles. To clarify the distinction, we suppose that a time series $\log(X_t)$ is decomposed additively as follows:

$$\log(X_t) = X_t^G + x_t, \tag{1.1}$$

where X_t^G is the trend or secular component and x_t is the cycle component. Often these components are identified by assuming something about their correlation. Some statistical models adopt a correlation of unity, as in Beveridge and Nelson (1981) and Stock and Watson (1988). Many others assume that the correlation is zero. In many cases these assumptions are unrelated to an economic model. For example, one might estimate the trend by least squares in:

$$\log(X_t) = \alpha + \beta \cdot t + \rho \log(X_{t-1}) + x_t, \tag{1.2}$$

where α, β, and ρ are unknown parameters. One criticism of this procedure is that there may be no economic reason to believe that local change in a trend is uncorrelated with the business cycle (see Zarnowitz (1992, chapter 6)). A second criticism is that this procedure chooses a trend to minimize the variance of the residual (cycle) component; here the trend would be defined as the linear combination of a time trend and lagged value which best explains the variation in $log(X_t)$. In contrast, the estimator we propose seeks to leave exactly as much variability in the cycle component as is predicted by a business-cycle model, and it also allows for correlation between the cycle and the trend. It does not treat the growth component X_t^G as the conditional expectation of $log(X_t)$.

Consider a $(n \times 1)$ vector time series X_t, $t = 1, \ldots, T$ of macroeconomic variables such as output, consumption, and investment. Business-cycle research focuses on a transformation of these data which induces stationarity but may also filter certain frequencies. For simplicity we shall call this detrending, even though inducing stationarity may not be the only aim. Denote by $x_t = f(X_t, \theta)$ the detrended series, where $\theta \in \Theta$ parametrizes a family of detrending filters. Investigators typically compare the moments of x_t with the population moments from a theoretical model. Denote the historical moments by $W_T(\theta)$, where the notation reflects the dependence of the historical, business-cycle properties on the detrending method. Denote by W the $q \times 1$ vector of population moments, for example from a numerically calibrated, stationary business-cycle model.

We suggest choosing a trend model by varying θ to minimize the distance between W_T and W. Thus

$$\hat{\theta} = \arg\min_{\theta} \|W_T(\theta) - W\|, \tag{1.3}$$

where the asymptotic distribution of the estimator will depend on the norm. We use the L^2 norm and GMM, surveyed by Ogaki (1992) and Davidson and MacKinnon (1993, chapter 17). This setup is quite general, and includes OLS estimation in (1.2) as a special case. In that case (with $n = 1$) $W_T(\theta)$ is:

$$\left\{ \frac{\sum_{t=1}^{T} x_t}{T}, \ \frac{\sum_{t=1}^{T} x_t \, t}{T}, \ \frac{\sum_{t=1}^{T} x_t \, log(X_{t-1})}{T} \right\}. \tag{1.4}$$

Equating these to $W = 0$ exactly identifies the three parameters. In contrast, if a business-cycle model is used to measure cycles then W will include moments such as the variance of output, investment, and consumption, and the covariance of consumption and investment with output. These moments may be calculated from a business-cycle model which has been fully calibrated and hence has no free parameters. Some researchers apply the transformation or filter f to the series from the business cycle model; we treat the transformation mainly as a detrending method and hence study the unfiltered implications of the model.

Section 2 applies this method using a standard, real business cycle model. It provides examples of cycle measurement using various moments and various parametric models of the growth component, X_t^G. Section 3 measures cycles using factor utilization rates which are stationary indicators of business cycles. Section 4 provides some simulation evidence on the properties of the procedure in various environments. Section 5 briefly concludes.

2. Measuring Cycles with a Real Business Cycle Model

2.1 Theoretical Model

The theoretical model used in the example is the baseline, one-sector real business cycle model of King, Plosser, and Rebelo (1988). Key features of the model include additively separable, logarithmic preferences, a Cobb-Douglas production function, a realistic rate of depreciation of capital, and persistent technology shocks (with first-order autocorrelation coefficient 0.9). In their analysis a stochastic growth model is solved approximately by linearizing first-order conditions around a steady-state path. The resulting difference equations can be solved analytically to give population moments for the approximated model. Those moments are listed completely in their Table 5, Panel B. A subset of them (an example of W) is listed in the top panel of Table 2.0.

2.2 Filters and Identification

We illustrate measuring cycles with this model by using several different parametric families of filters. The main filter family is:

$$x_t = f(X_t, \theta) = log(X_t) - \alpha - \beta \cdot t - \rho \cdot \log(X_{t-1}), \tag{2.1}$$

where $\theta = (\alpha \;\; \beta \;\; \rho)$ is a $n \times 3$ matrix. The filter thus includes as special cases: (i) log-linear detrending, when $\rho = 0$; (ii) log first-differencing (approximate calculation of growth rates), when $\beta = 0$ and $\rho = 1$. These detrending methods are widely used in economics. Importantly, there is now a lengthy list of studies (including those by Canova (1991), Blackburn and Ravn (1991) Cogley and Nason (1992a), and Harvey and Jaeger (1993)) which show that 'stylized facts' may vary with the detrending procedure used. That suggests that cycle models – which restrict those facts – will succeed in identifying the parameters θ of detrending filters.

In this family the parameter vector θ has $3n$ elements. Typically W consists of variances and of covariances with output, as well as zero-mean restrictions on x_{jt} $(j = 1, \ldots, n)$. In that case W_T has $q = n + 2n - 1$ distinct elements. A necessary condition for identification is that q be at least as great as the number of parameters estimated. Thus some restrictions on the parameters are needed, or else further moments (such as autocorrelations) must be added.

In the present application β and ρ are constrained to be the same across time series. Common trends across the n series might be included as a balanced-growth requirement, or simply as a restriction which aids identification and allows a richer family of transformations, given a small dimension of W. Cochrane (1994) also uses a common stochastic trend in consumption and output to decompose output into permanent and temporary components. It also might be possible to include cointegration restrictions in the multivariate filter used here.

The second filter involves a simple modification to the one in equation (2.1). The deterministic trend is now segmented in 1973:I, using a dummy variable D_t. That break point was suggested by Perron (1989). The filter is:

$$x_t = f(X_t, \theta) = log(X_t) - \alpha - (\beta - \delta \cdot D_t) \cdot t - \rho \cdot \log(X_{t-1}) \tag{2.2}$$

with $D_t = 0$ for $t < 1973$:I and $D_t = 1$ for $t \geq 1973$:I. This example is quite simple, and various extensions involving unknown breakpoints are possible. However, the method-of-moments framework does rule out filters with which we cannot define GMM residuals. Examples include the segmented trends model of Balke and Fomby (1991) and other trend models which are based on unobserved secular components such as those described by Watson (1986).

Filters may be applied to isolate business-cycle fluctuations, and not simply to induce stationarity. With this aim, many researchers use filters which smooth X_t with (possibly infinite) two-sided moving averages. Examples include the Kuznets filter studied by Howrey (1968) and Sargent (1987), the exponential smoothing filter adopted by Friedman (1957) and Lucas (1980) and studied by King and Rebelo (1993), and the Hodrick-Prescott (HP) (1980) filter studied by Singleton (1988), King and Rebelo (1993), Cogley (1990), Söderlind (1991), Cogley and Nason (1992a), and Harvey and Jaeger (1993). These filters typically are restricted so that their weights sum to one (which ensures that the filter has unit gain at the zero frequency and hence that the original series and the trend component have a common stochastic trend) and sometimes further restricted to be symmetric.

Our third filter is a simple, finite moving average with symmetric, geometrically declining weights. This allows for some smoothing, while identifying a small number of parameters in the time domain. For each X_j the filter family is:

$$x_{jt} = \log(X_{jt}) - \alpha_j - \beta \cdot t - \rho_0 \{ \log(X_{it}) + \sum_{k=1}^{3} \rho_1^k [\log(X_{jt+k}) + \log(X_{jt-k})] \} \tag{2.3}$$

This filter does not include (2.1) as a special case, because it is two-sided.

Finally, we also estimate by GMM the single parameter in the HP filter. In the literature on nonparametric smoothing this filter is known as the Whittaker-Henderson method, and its parameter has been estimated by cross validation (see Buja, Hastie, and Tibshirani (1989)).

2.3 Results

The estimation uses U.S. output (Y), consumption (C), and investment (I) for 1947:1 - 1991:3, quarterly, so that $n = 3$. All series are drawn from the CITIBASE database and are described at the end of the paper. Output is real GNP, consumption is expenditure on nondurables and services, investment is gross fixed investment. All three are quarterly, real, seasonally adjusted, and expressed in per capita terms by dividing by the total civilian non-institutional population 16 years of age and older. These series have been filtered with (approximately) a two-sided moving average: X-11.

In cases with overidentification, weights in the GMM minimization will affect the estimates. The instrument set is a constant, so as to match unconditional moments, as in the theory. The GMM estimator weights moments in inverse proportion to their sampling variability. We use the Hansen-Heaton-Ogaki GMM code, with Durbin's method used (where necessary) to calculate the weighting matrix.

The sample moments are:

$$W_T(\theta) = \sum_{t=1}^{T} w[f(X_t, \theta)] = \sum_{t=1}^{T} w(x_t). \tag{2.4}$$

For standard distribution theory to be used it is necessary that X_t be stationary so that $w[f(X_t, \theta)]$ is stationary for all admissible values Θ of θ and not just for the true values θ_0. While time trends can be estimated by GMM (see Eichenbaum and Hansen (1990) and Andrews and McDermott (1993)), and the point estimates below generally yield deterministic trends, integrated nonstationarity precludes standard inference. Because x_t is nonstationary for certain parameter values, so is w_t.

Moreover, asymptotic inference might be misleading because of the small sample sizes used. Andrews (1991) has shown, for example, that none of the HAC estimators he considers is reliable with less than 250 observations, when the GMM error terms have first-order autocorrelation of 0.9, which is almost exactly the case for the business-cycle terms x_t here. For these reasons, we view the method as providing consistent estimates (given an appropriate family of filters) regardless of the estimating environment, but suggest repeated simulation with fitted parameter values in order to gauge sampling variability.

We begin with the baseline filter family of equation (2.1) and illustrate measuring cycles with two sets of moment conditions W. Method I uses the three means, three variances, and two covariances with output, shown in Table 2.0. Thus there are 3 over-identifying restrictions. Parameter estimates are shown in Table 2.1. The implied moments of measured cycles (in percent deviations from trend, $100 \cdot x_{it}$) are shown in the bottom panel of Table 2.0.

For comparison, Method II uses the orthogonality conditions of OLS (1.4). These give 9 moment restrictions so that there are 4 over-identifying restrictions because a common trend is imposed. The estimates are shown in Table 2.2.

The two decompositions are strikingly different. The least-squares estimator (method II) finds a unit root in the series, whereas the cycle-theory estimator suggests a linear time trend. This trend is similar to the trend used by King, Plosser, and Rebelo (1988) and reported in their Table 6. The three over-identifying restrictions arise from the use of a common trend in all three variables. A test of these restrictions is a test of balanced growth, for if they hold then $log(C_t) - log(Y_t)$ and $log(I_t) - log(Y_t)$ will be stationary as long as $|\rho| < 1$. If $\rho = 1$ then this restriction is necessary for balanced growth but not sufficient, for the three series might not be cointegrated.

The rejection of these restrictions by the J-test in Table 2.1 seems to contradict the findings of King, Plosser, Stock, and Watson (1991) and Neusser (1991). Those authors found in trivariate systems (for output, consumption, and investment in the U.S.) that the consumption-income and investment-income ratios were stationary, an implication of balanced growth. However, the conflict is only apparent because King, Plosser, Stock, and Watson (1991) also found that the dynamics of these three series were not consistent with the predictions of one-sector real business cycle models. Our estimated trend is constructed so that the moments of the residual, cycle components resemble those of the real business cycle model as closely as possible. Thus any inadequacies of the model may show up in the trends rather than in the cycles. As in that study, we find that this model cannot explain both trends and cycles. Of course, rejections also could be due to the parametric filter used *or* to the use of the χ^2 distribution. Section 4 provides some simulation evidence on the distribution of the J-test statistic in this environment.

We also investigate the effect of using different sets of moments from the business-cycle model. For example, matching correlations instead of covariances gives slightly different results. It also seems sensible to include autocorrelation information from the theoretical model, in addition to the moments shown in Table 2.0 and used in Table 2.1. The theoretical first and second-order autocorrelations of output are 0.93 and 0.86. With these additional moment restrictions the estimates are very similar to those in Table 2.1.

Next we use the same 8 moments W from Table 2.0, but adopt the filter in equation (2.2), which allows for segmented trend. In this case $\hat{\delta} = 0.00114$, with an estimated standard error of 0.00147, so that the segmentation appears to be insignificant. Breaking the trend in 1973 does not make cycles before and after that time jointly resemble the theoretical cycles. Thus there is no evidence of a break, from the perspective of this business-cycle theory.

Finally, the linear trend is the only element filtered out by the family of filters in equation (2.3), which allows for a two-sided moving average in the secular component. The additional parameters ρ_0 and ρ_1 are zero to two decimal places and not different from zero at conventional significance levels. The results are essentially those of Table 2.1, a finding which is not surprising given findings of King and Rebelo (1993, section 1.2). For the HP filter they find that removing a two-sided moving average once data have been linearly detrended does not leave a cyclical component that resembles a business cycle (for example the first-order autocorrelation of output is 0.09). We use a simpler filter, but also let the estimation procedure choose whether or not a linear trend is significant. The same finding occurs when we estimate the parameter in the HP filter by GMM. The estimated growth component is the limiting case of a deterministic, linear time trend.

2.4 Discussion

These examples use the real business cycle model to measure cycles. In a sense this method may make it more difficult to *test* business-cycle models. If cycles are defined by varying θ to match empirical and theoretical moments then the closeness of the match cannot be used as a test. A test requires some overidentification restrictions (more moments than parameters), which constrain the coefficients in the filter. However, typical models make predictions for many moments, and so would allow both measurement and testing. Some implications of a theory may be used to isolate business cycles and then remaining differences between theory and evidence may be used to reformulate the theory, in confidence that these mismatches are not artefacts of the detrending procedure.

Ideally, the theoretical model is a distillation of evidence from past cycles (possibly in several countries). To serve as a measurement device, it must be quantitative, stationary, and not derived or calibrated after assuming some arbitrary trend. The examples use a model with no free parameters. But some real, business-cycle models are calibrated by assuming some form of trend, say parametrizing a technology shock process to mimic detrended Solow residuals. Moreover, King, Plosser, and Rebelo (1988) and Cogley and Nason (1992b) have shown that this model has little propagation; most of the output dynamics arise from shock dynamics. Persistence in shocks also affects relative volatilities in this model, and so may be identifiable. It would be conceptually straightforward (though computationally burdensome) to be agnostic about parameters of the shock process, such as its first-order autocorrelation, and estimate them jointly with those of the detrending filter using the X_{it} alone. Other free parameters also could be estimated, although jointly identifying the parameters of the filter and of the model may be challenging. Smith (1993) provides a simple example of such joint estimation.

The theoretical model described in section 2.1 is that of King, Plosser, and Rebelo (1988), which is consistent with trend-stationary time series behaviour for output, consumption, and investment. However, there is no circularity involved in using the moments from this model to estimate trends and then in finding a linear trend. As those authors note, the moments from a detrended growing version of the model (as in Table 2.0) are the same as those from a non-growing version with slightly different depreciation and discount rates.

While our illustration of measuring cycles with a real business cycle model may be new, the underlying idea is not. Koopmans (1947) argued precisely for using economic theory to help define business cycles. Lippi and Reichlin (1994) identify trend components as having S-shaped impulse response functions which they attribute to the diffusion of technical change. Recent studies which measure cycles using economic models of the cycle include those of Eichenbaum, Hansen, and Singleton (1988) (who use Euler equations to detrend), Cochrane (1994) (who uses the permanent-income hypothesis), and Laxton and Tetlow (1992) and Kuttner (1993) (who use Okun's Law and the Phillips Curve). The same principle has been applied to decompositions into seasonal and cyclical components, where parameters of an economic model and of a seasonal filter have been estimated jointly.

3. Measuring Cycles with Cycle Indicators

An alternative to adopting an economic model in order to define cycles is to restrict their correlation with business cycle indicators. This method fits in the same method-of-moments framework. Indicators which are themselves trending could be used to detrend output, by choosing filter parameters to maximize the correlations between the cyclical components. This section illustrates a simpler method, which is to use stationary indicators to measure cycles in output.

The indicators adopted are the rate of capacity utilization in U.S. manufacturing and the U.S. rate of unemployment. Both series are quarterly, 1948:I-1991:III and seasonally adjusted. The capacity utilization rate is that for manufacturing because it is available for the period beginning in 1948 and because it is based substantially on survey evidence. Raddock (1990) discusses the construction of this measure, while Shapiro (1989) outlines some of its shortcomings. We ignore issues arising from changes over time in the measurement of unemployment rates. Figure 3A shows these two stationary business cycle indicators, demeaned.

The filter used is again (1.2) but applied to output only:

$$\log(Y_t) = \alpha_y + \beta \cdot t + \rho \log(Y_{t-1}) + y_t. \tag{3.1}$$

We consider several sets of moment conditions to estimate the three parameters in equation (3.1). First, as a benchmark we use the OLS conditions. Second, we set the mean of y_t to zero, maximize the correlation between the output cycle y_t and the capacity utilization rate, and minimize (make close to -1) the correlation between y_t and the unemployment rate. Third, we combine these two sets of moment conditions so that there are two overidentifying restrictions. The three rows of Table 3 list the results of estimation with these three sets of moments.

Table 3: Output Filter Parameters

	Parameter			
	α_y	β	ρ	J
OLS	0.0137	0.00010	0.964	
	(0.00453)	(0.000064)	(0.0186)	
Indicators	0.223	0.00437	-0.193	
	(0.294)	(0.00117)	(0.782)	
Combination	0.268	0.00447	-0.292	55.98
	(0.0242)	(0.000199)	(0.0729)	(0.00)

Parentheses contain standard errors for coefficients and the p-value for the J-statistic.

From the first row, the OLS conditions again suggest first-differencing. From the second row, measuring cycles using correlations with cycle indicators leads to linear detrending. Unemployment may be a lagging indicator, but these estimates do not change very much when we lead the unemployment rate by two or four quarters. Figure 3B shows the cycles estimated in the first and second rows of Table 3. As in section 2, first-differencing does not produce a series with the persistence (or differential volatility, or correlation with cycle indicators) associated with business cycles.

The third row shows that combining the restrictions results in more precision, but the overidentifying restrictions are rejected. The rejection implies that we cannot construct a cyclical component of output which is uncorrelated with the secular component and at the same time highly correlated with cycle indicators.

The cycle estimated from the correlation conditions may not necessarily strike the viewer of Figure 3B as the appropriate one in postwar U.S. history. That may be a consequence of the restrictive filter in equation (3.1). More flexible trends could be fitted using further moment conditions with these indicators. For example, we could match the autocorrelation properties in the output cycles with those in the indicators. The other way to add moment conditions would be to add indicators. Candidates include layoff rates, the ratio of price to unit cost in manufacturing, rates of business failure, and spreads between various interest rates. Indicators for the U.S. are described by Moore (1990), Frumkin (1990), Darnay (1992), and Zarnowitz (1992).

4. Simulation Examples

The applications have used samples containing fewer than 200 observations. They also involve a non-standard application of GMM to cases in which there potentially is a unit root in the series. We next present some Monte Carlo evidence on the properties of the procedure. The aim is to show some distributional properties of the estimators. We also calculate the empirical size of the standard J-test of overidentifying restrictions.

One possibility here would be to solve and then repeatedly simulate several parametrizations of a model of cycles and growth. One could then see whether the method correctly decomposed the nonstationary, simulated time series given various filters which could capture the true growth component. Our experimental design is in this spirit, but simpler.

We adopt a linear three-equation economy in which:

$$log(X_{jt}) = \alpha_j + \beta\, t + \rho\, log(X_{jt-1}) + x_{jt} \tag{4.1}$$

for $j = 1, 2, 3$. Also,

$$x_{jt} = \lambda_j x_{jt-1} + \omega_{jt}, \tag{4.2}$$

where ω_{jt} is serially uncorrelated and generated as pseudo-normal with zero mean. The business-cycle components x_{jt} have covariance matrix:

$$\Sigma_x = \begin{pmatrix} 0.0426^2 & & \\ 0.000953 & 0.0273^2 & \\ 0.00384 & 0.001 & 0.0982^2 \end{pmatrix}$$

which involves the same variances and covariances with output as in the business-cycle model used in section 2 (Table 2.0). These components also are persistent: $\lambda_j = 0.9$ for each series.

We consider two versions of this DGP, which differ only in their trend characteristics. And we consider two estimators: method I, which matches cycle moments from the theoretical model and method II which is multivariate OLS with a common trend. Version A has parameter values suggested by the application of method I to the historical data. That gives α_y, α_c, α_i, β, and ρ. Version A thus involves a linear time trend in the three series. The parameter values are listed in the top panel of Table 4A.

From version A we generate 1000 replications of length $T = 175$, which is the same as the number of quarterly observations. Then we estimate the filter parameters by the two methods. In each method we correctly impose common values for β and ρ across equations. Method I uses the three mean zero conditions, the three variances of the cycle components, and the covariances between output and consumption and between output and investment. The DGP is calibrated so that these conditions are satisfied given the trend in the series. Method II is OLS with these cross-equation restrictions. It uses the 9 OLS conditions to estimate the five parameters. These methods are exactly those used in the historical data in section 2.

Table 4.1 shows the results for estimates of β and ρ, their standard errors, and the J-statistic. The first column of the table gives the mean of the statistic across replications and the second column gives the standard error across replications.

Three general conclusions follow from Table 4.1. First, method I (BC) is consistent. Second, method II (OLS) is inconsistent. In particular, it dramatically underestimates the parameter on the time trend and overestimates the parameter on the lagged dependent variable in the detrending filter. This yields results similar to those found in the historical data and summarized in Table 2.2. Third, the asymptotic standard errors and critical values used in method I are very misleading. For $\hat{\beta}$ the average estimated standard error is 0.00017 while the Monte Carlo standard error is 0.00059. The J-test overrejects: its average value (27.0) has a p-value of 0.00 in the asymptotic distribution.

However, the simulations may be used to approximate exact inference, because they are calibrated to the estimates from Table 2.1. For example, the J-statistic in the application (Table 2.1) of method I is 65.6. Locating this in the simulated empirical distribution summarized in Table 4.1 gives a p-value of 0.175. Thus the business-cycle and common-trend restrictions are *not* rejected at conventional levels of significance. Likewise, the Monte Carlo experiment suggests that the standard errors in Table 2.1 be multiplied roughly by four.

Next, version B of the DGP has parameters set at the values fitted for this filter by method II in Table 2.2. The trend parameters in the DGP are listed in the top panel of Table 4.2, and involve a common unit root. In this case method II also is consistent, despite the autocorrelation in the residuals (cycles) because there is a unit root on the lagged dependent variable ($\rho = 1$).

The main conclusion from this experiment is that there is some efficiency loss in adopting method I rather than method II when the OLS orthogonality conditions hold. However, while method II is more efficient, inferences from the two methods would be identical.

The conditions used by the cycle-property estimator (method I) provide consistency whether or not there is a stochastic trend in the growth component. In contrast, least-squares (method II) tends to find near-unit-root behaviour when there is a linear trend. When the orthogonality conditions imposed under least-squares are appropriate (as in version B of the DGP) they provide only modest efficiency gains. This combination of findings suggests omitting the least-squares orthogonality conditions (and perhaps constructing a Durbin-Wu-Hausman test statistic).

Section 2.2 mentioned the possibility of using simulation methods for testing, given the small sample sizes. We have done that here, using Monte Carlo methods at fitted values to construct standard errors and critical values. Of course, moment restrictions do not imply a unique DGP, so to use simulation to allow tests and to plot standard errors around measured cycles requires that one specify a complete DGP, as we have done here. The linear example here is simple, but the obvious choice is a complete business-cycle model. However, if that is adopted for inference the information contained in it also should be used for estimation. Maximum likelihood methods then could be used to estimate filter parameters and free parameters of the economic model.

5. Conclusion

This study has outlined a statistical method for defining or measuring cycles using business-cycle theory, and given some applications. Further issues to be explored include the effects on cycle measurement of using other business-cycle models and other parametric filters (although some filters cannot be implemented by GMM).

The statistical distribution theory for the estimators when there is stochastic nonstationarity remains to be studied, although inference can be undertaken using simulation around fitted values. Further over-identification could be established by requiring stability of moments across subsamples. Likewise, tests of parameter constancy could be applied, in addition to the J-test used here.

Sometimes detrending methods are criticized by showing that they may give misleading cycle properties if the true trend differs from that assumed in the detrending. An objection to this criticism is that the true trend is unknown. An advantage of the method described here is that it allows one to be agnostic (up to a parametric family) about the true trend, and yet to study cycles. Thus we can weaken considerably the auxiliary assumptions about trends which are made in studying cycle models. Of course, the method also can be applied when a model includes both growth and cycles. One would expect statistical efficiency gains in that case.

With overidentification, the cycle measurement may not be very sensitive to any one moment restriction. Nevertheless, the definition of cycles will change with changes to the cycle model. Though this seems natural to us, it contrasts with the program of confronting various models with a fixed set of business-cycle 'stylized facts'.

Data Sources

All series are drawn from the CITIBASE database. Output is real GNP, consumption is expenditure on nondurables and services, investment is gross fixed investment. All three are quarterly, real, seasonally adjusted, and expressed in per capita terms by dividing by the total civilian non-institutional population 16 years of age and older.

Capacity utilization rate in U.S. manufacturing, quarterly, SA. Source: Federal Reserve Board of Governors. U.S. unemployment rate (all workers older than sixteen years) SA. Source: Bureau of Labor Statistics.

Table 2.0: Model Properties W
(King, Plosser, and Rebelo (1988), unfiltered population moments)

Series	Std. Dev.	Std. Dev./Std. Dev.(output)	Correlation with Output
Output	4.26	1.00	1.00
Consumption	2.73	0.64	0.82
Investment	9.82	2.31	0.92

U.S. Data Properties $W_T(\hat{\theta})$
(detrended quarterly data 1948:1-1991:3)

Series	Std. Dev.	Std. Dev./Std. Dev.(output)	Correlation with Output
Output	5.09	1.00	1.00
Consumption	4.09	0.80	0.82
Investment	7.56	1.48	0.66

Table 2.1: Method I Parameter Estimates $\hat{\theta}$

Parameter	Estimate	Standard Error
α_y	0.187	0.0076
α_c	-0.415	0.015
α_i	-1.632	0.0536
β_{BC}	0.00360	0.000144
ρ_{BC}	0.00995	0.033
$J(3)$	65.6	0.00

Table 2.2: Method II Parameter Estimates $\hat{\theta}$

Parameter	Estimate	Standard Error
α_y	0.00804	0.0024
α_c	-0.000512	0.0048
α_i	-0.0196	0.018
β_{OLS}	0.0000335	0.000041
ρ_{OLS}	0.985	0.011
$J(4)$	5.75	0.22

Note: Method I uses business-cycle (BC) moments. Method II is multivariate ordinary least squares (OLS). The last column gives the estimated standard errors of the filter parameters and the p-value of the J-statistic from the χ^2 distribution.

Table 4.1: Simulated Deterministic Trend

DGP Trend parameters: $\alpha_y = 0.187$ $\alpha_c = -0.415$ $\alpha_i = -1.632$ $\beta = 0.0036$ $\rho = 0.00$.

	Mean	Standard Error
Method I		
$\hat{\beta}_{BC}$	0.00397	0.00059
se_β	0.00017	0.00010
ρ_{BC}	-0.10055	0.14380
se_ρ	0.04170	0.01048
$J(3)$	27.00749	23.2366
Method II		
$\hat{\beta}_{OLS}$	0.000398	0.000105
se_β	0.000106	0.000017
ρ_{OLS}	0.888948	0.029012
se_ρ	0.029354	0.004872
$J(4)$	4.311382	2.820464

Note: Method I uses business-cycle (BC) moments. Method II is multivariate ordinary least squares (OLS). The mean and standard error of each statistic are constructed from 1000 replications.

Table 4.2: Simulated Unit Root

DGP Trend parameters: $\alpha_y = 0.008$ $\alpha_c = 0.00$ $\alpha_i = 0.00$ $\beta = 0.00$ $\rho = 1.00$.

	Mean	Standard Error
Method I		
$\hat{\beta}_{BC}$	0.00000227	0.000374
se_β	0.0000737	0.0000487
ρ_{BC}	1.000076	0.0225
se_ρ	0.00435	0.00411
$J(3)$	23.74	19.88
Method II		
$\hat{\beta}_{OLS}$	-0.0000145	0.000239
se_β	0.0000416	0.0000138
ρ_{OLS}	1.00081	0.00789
se_ρ	0.00214	0.000839
$J(4)$	27.45	19.12

Note: Method I uses business-cycle (BC) moments. Method II is multivariate ordinary least squares (OLS). The mean and standard error of each statistic are constructed from 1000 replications.

References

Andrews, D.W.K. (1991) Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica* 59, 817-858.

Andrews, D.W.K. and McDermott, C.J. (1993) Nonlinear Econometric Models with Deterministically Trending Variables. Cowles Foundation Discussion Paper 1053, Yale University.

Balke, N.S. and Fomby, T.B. (1991) Shifting Trends, Segmented Trends, and Infrequent Permanent Shocks. *Journal of Monetary Economics* 28, 61-85.

Beveridge, S. and Nelson, C. (1981) A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle.' *Journal of Monetary Economics* 7, 151-174.

Blackburn, K. and Ravn, M.O. (1991) Univariate Detrending of Macroeconomic Time Series. memo 1991-22, Aarhus Universitet.

Buja, A., Hastie, T. and Tibshirani, R. (1989) Linear Smoothers and Additive Models. (with discussion) *Annals of Statistics* 17, 452-555.

Canova, F. (1991) Detrending and Business Cycle Facts. European University Institute Working Paper in Economics 91/58.

Cochrane, J.H. (1994) Permanent and Transitory Components of GNP and Stock Prices. *Quarterly Journal of Economics* 109, 241-265. University of Chicago.

Cogley, T. (1990) Spurious Business Cycle Phenomena in HP Filtered Time Series. Department of Economics, University of Washington, working paper 90-21.

Cogley, T. and Nason, J.M. (1992a) Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research. Discussion Paper No. 92-23, Department of Economics, University of British Columbia.

Cogley, T. and Nason, J.M. (1992b) Do Real Business Cycle Models Pass the Nelson-Plosser Test? Discussion Paper No. 92-24, Department of Economics, University of British Columbia.

Darnay, A.J. ed. (1992) *Economic Indicators Handbook.* Detroit: Gale Research.

Davidson, R. and MacKinnon, J.G. (1993) *Estimation and Inference in Econometrics.* Oxford University Press.

Eichenbaum, M.S. and Hansen, L.P. (1990) Estimating Models with Intertemporal Substitution using Aggregate Time Series Data. *Journal of Business and Economic Statistics* 8, 53-69.

Eichenbaum, M.S., Hansen, L.P., and Singleton, K.J. (1988) A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty. *Quarterly Journal of Economics* 103, 51-78.

Friedman, M. (1957) *A Theory of the Consumption Function.* Princeton NJ: Princeton University Press.

Frumkin, N. (1990) *Guide to Economic Indicators.* Armonk, NY: M.E. Sharpe.

Harvey, A.C. and Jaeger, A. (1993) Detrending, Stylized Facts, and the Business Cycle. *Journal of Applied Econometrics* 8, 231-247.

Hodrick, R.J. and Prescott, E.C. (1980) Postwar U.S. Business Cycles: An Empirical Investigation. Discussion Paper 451, GSIA, Carnegie Mellon University.

Howrey, E.P. (1968) A Spectrum Analysis of the Long-Swing Hypothesis. *International Economic Review* 9, 228-252.

King, R.G., Plosser, C.I., and Rebelo, S.T. (1988) Production, Growth and Business Cycles I: The Basic Neoclassical Model. *Journal of Monetary Economics* 21, 195-232.

King, R.G. and Rebelo, S.T. (1993) Low Frequency Filtering and Real Business Cycles. *Journal of Economic Dynamics and Control* 17, 207-231.

King, R.G., Plosser, C.I., Stock, J.H., and Watson, M.W. (1991) Stochastic Trends and Economic Fluctuations. *American Economic Review* 81, 819-840.

Koopmans, T.C. (1947) Measurement without Theory. *Review of Economics and Statistics* 29, 161-172.

Kuttner, K.N. (1993) An Unobserved-Components Model of Constant-Inflation Potential Output. mimeo, Federal Reserve Bank of Chicago.

Laxton, D. and Tetlow, R. (1992) A Simple Multivariate Filter for the Measurement of Potential Output. Bank of Canada Technical Report No. 59.

Lippi, M. and Reichlin, L. (1994) Diffusion of Technical Change and the Decomposition of Output into Trend and Cycle. *Review of Economic Studies* 61, 19-30.

Lucas, R.E. (1980) Two Illustrations of the Quantity Theory of Money. *American Economic Review* 70, 1005-1014.

Moore, G.H. (1990) *Leading Indicators for the 1990s.* Homewood, Il.: Dow Jones-Irwin.

Neusser, K. (1991) Testing the Long-Run Implications of the Neoclassical Growth Model. *Journal of Monetary Economics* 27, 3-37.

Ogaki, M. (1992) Generalized Method of Moments: Econometric Applications. In G.S. Maddala, Rao, C.R., and Vinod, H.D., eds. *Handbook of Statistics, volume 11: Econometrics.* Amsterdam: North-Holland.

Perron, P. (1989) The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis. *Econometrica* 57, 1361-1401.

Raddock, R.D. (1990) Recent Developments in Industrial Capacity and Utilization. *Federal Reserve Bulletin* 76 (June), 411-435.

Sargent, T.J. (1987) *Macroeconomic Theory.* second edition. Academic Press.

Shapiro, M. (1989) Assessing the Federal Reserve's Measures of Capacity Utilization. *Brookings Papers on Economic Activity* 181-225.

Singleton, K.J. (1988) Econometric Issues in the Analysis of Equilibrium Business Cycle Models. *Journal of Monetary Economics* 21, 361-386.

Smith, G.W. (1993) Method-of-Moments Measurement of U.K. Business Cycles. mimeo, Department of Economics, Princeton University.

Söderlind, P. (1991) Is There a Cycle in Real Business Cycle Models? mimeo, Institute for International Economic Studies, Stockholm University.

Stock, J.H. and Watson, M.W. (1988) Testing for Common Trends. *Journal of the American Statistical Association* 83, 1097-1107.

Watson, M.W. (1986) Univariate Detrending Methods with Stochastic Trends. *Journal of Monetary Economics* 18, 49-75.

Zarnowitz, V. (1992) *Business Cycles: Theory, History, Indicators, and Forecasting.* NBER, University of Chicago Press.

FIGURE 3A: CAPACITY UTILIZATION AND UNEMPLOYMENT RATES

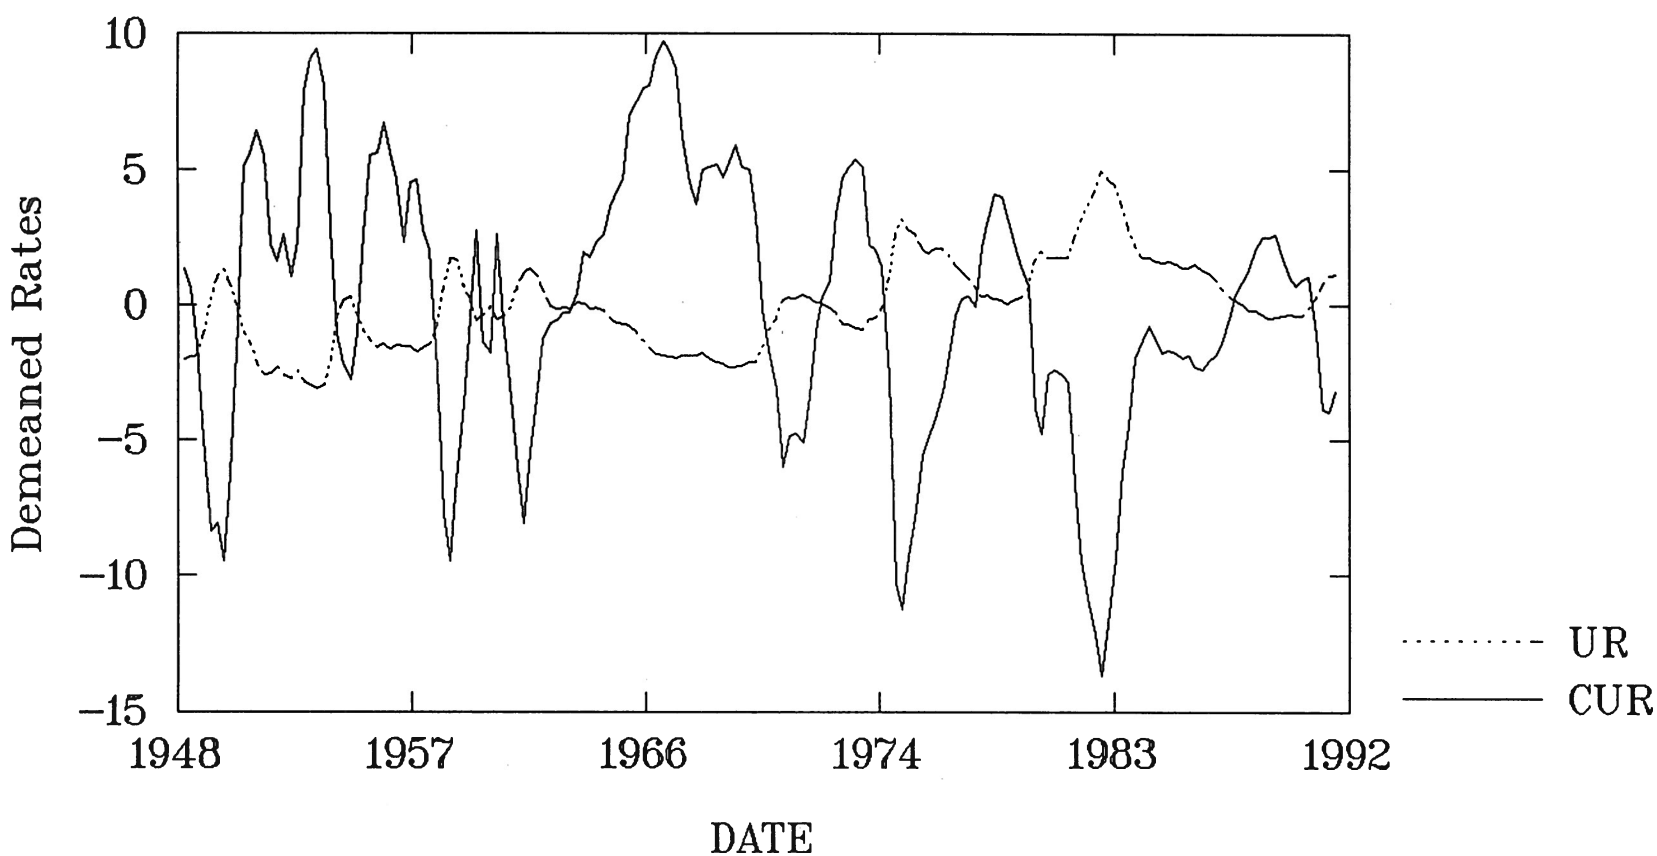

FIGURE 3B: OUTPUT CYCLES