Evaluating Non-Structural Measures of the Business Cycle

Timothy Cogley

Senior Economist, Federal Reserve Bank of San Francisco.
I am grateful to John Cochrane, Ken Kasa, Brian Motley, Glenn Rudebusch, and Simon van Norden for helpful comments, and to Desiree Schaan and Heather Royer for research assistance.

This paper evaluates a number of non-structural measures of the business cycle. It adopts a structural definition of the cycle, interprets non-structural measures as noisy approximations, and seeks a proxy that is reliable across a variety of plausible trend-cycle structures. The results favor a consumption-based measure proposed by Cochrane (1994). Across a variety of structures, it has the highest correlation and coherence with structural cycles, and best matches their dynamic properties. When applied to U.S. data, consumption-based measures conform closely to the dates of NBER recessions. They also yield a strong negative correlation between the cyclical components of productivity and hours, a fact that deepens the challenge to models which emphasize technology shocks as the primary source of business cycles.

The study of business cycles begins with the problem of measurement. The chief issue concerns how to separate macroeconomic data into trends and cycles. The empirical literature contains a wide variety of competing methods, some based on structural definitions of trends and others based on non-structural definitions. Among the former class, Blanchard and Fischer (1989) define the trend as that part of output which is driven by permanent shocks. This is a structural definition, as it requires a structural model in order to identify permanent shocks. In one respect, this approach is conceptually appealing because it defines the trend in terms of economic dynamics. But it also entangles the problem of measuring business cycles among all of the hard problems associated with specifying and estimating structural macroeconomic models.[1]

There are also a number of non-structural or statistical measures of trends. For example, Beveridge and Nelson (1981) define the trend in terms of the long-horizon forecast of the level of output, which can be derived from a reduced form forecasting model. Similarly, Cochrane (1994) argues that consumption provides a good estimate of the trend level of output, since consumers also try to distinguish between permanent and transitory movements in income. Alternatively, Hodrick and Prescott (1997) and Baxter and King (1995) define the trend as that part of output which lies in a low-frequency band, and they use linear filters to extract measures of the trend.

Non-structural approaches offer one important practical advantage. Because they do not condition on a structural model, they separate the problems of measuring and explaining business cycles. But the non-structural approach also suffers from an important drawback: there are infinitely many ways to decompose a growing time series into stationary and non-stationary components, and the choice of decomposition matters. For example, various authors have shown that stylized facts about periodicity, comovement,

---

1. See Blanchard and Quah (1989), Lippi and Reichlin (1994a,b), and Cogley and Nason (1995a) for applications of this approach. Hansen and Sargent (1991), Lippi and Reichlin (1993), Blanchard and Quah (1993), and Faust and Leeper (1997) discuss some of the difficulties that can arise.

and the relative importance of permanent and transitory fluctuations are all sensitive to the choice of statistical decomposition.[2]

Since the choice of decomposition matters, we need some kind of criterion in order to guide our selection. A convincing selection criterion must involve some economics, since a criterion that lacked substance would be arbitrary. Hence the non-structural approach to measuring business cycles ultimately leads back to the economic approach, at least to the extent of formulating a selection criterion.

This paper proposes a selection criterion and uses it to rank a number of popular non-structural measures. The paper adopts Blanchard and Fischer's structural definition of the business cycle and interprets non-structural measures as imperfect proxies or noisy indicators. The problem, then, is to look for a proxy that is reliable across a variety of plausible trend-cycle structures.

I start with a simple dynamic macroeconomic model in which output is driven by government spending and technology shocks. I follow the literature and alter the shock dynamics in order to generate a variety of trend-cycle structures. In one version, both shocks are transitory, which makes output "trend stationary" (TS). In other versions, government spending shocks are transitory but technology shocks are permanent, and this makes output "difference stationary" (DS). I also vary the innovation variances in the DS models in order to generate various degrees of trend reversion. In one DS specification, a stochastic trend dominates the variation in output growth; this is consistent with the interpretation of Nelson and Plosser (1982). Other DS specifications have moderate to substantial degrees of trend reversion and are more consistent with the interpretation of Blanchard and Quah (1989).

I use these models to generate artificial data, and then apply the methods of Baxter and King, Beveridge and Nelson, Cochrane, and Hodrick and Prescott. In each case, I compute a number of summary statistics to evaluate comovements between the proxies and the structural cycle and to compare the dynamic properties of the proxies with those of the structural cycle.

I find that Cochrane's idea dominates. Across the various structures, the consumption-based measure has the highest correlation and coherence with the structural cycle. It also best replicates the dynamics of the structural cycle, in the sense of matching its normalized spectrum. The consumption-based measure often overestimates the amplitude of the structural cycle, but a simple rescaling (based on the mean amplitude of the four indicators) helps to mitigate this problem. The rescaled consumption-based measure also has the highest R^2 statistics.

When applied to post-war U.S. data, consumption-based measures conform closely to the dates of NBER recessions. The troughs in output, investment, and hours are very close to NBER troughs, and there are no false signals of recession. Consumption-based measures also suggest that productivity and hours vary countercyclically. This fact is significant because it reinforces the challenge to models in which technology shocks are the primary source of business cycles.

The rest of the discussion is organized as follows. Section I describes various structural and non-structural measures of the business cycle and characterizes the errors that each of the proxies makes. Section II describes the data generating processes used in the simulations and illustrates their dynamic properties. Section III applies the non-structural methods to artificial data generated by those models and evaluates their performance as cyclical indicators. Section IV applies the consumption-based approach to postwar U.S. data and comments on the results. The paper concludes with some simple advice: researchers who desire a non-structural measure of the business cycle should detrend output by regressing it on consumption.

I. Structural and Non-Structural Measures of the Business Cycle

This section defines Blanchard and Fischer's structural measure of the business cycle and describes how it relates to a variety of non-structural measures. Each of the non-structural measures can be interpreted as noisy indicators or imperfect proxies for the structural cycle, and this section develops some intuition about the nature of the approximation error in each case.

Blanchard and Fischer's Structural Definition

Let x_t denote a vector of endogenous variables that includes output growth, $\Delta\ln(y_t)$, as well as other variables that have been transformed into a stationary form. For example, these might include error correction terms and growth rates of

2. Nelson and Kang (1981), Harvey and Jaeger (1993), and Cogley and Nason (1995b) show that periodic properties of the cycle, as measured by its autocorrelation function or power spectrum, are sensitive to the choice of statistical decomposition. King and Rebelo (1993) and Cogley and Nason (1995b) show that some important measures of comovements, such as the correlation between output and hours worked, are sensitive to the choice of statistical decomposition. And Quah (1992) shows that increments to the non-stationary component can be highly variable or arbitrarily smooth depending on assumptions about the degree of autocorrelation in trend growth. Hence, measures of the relative importance of permanent and transitory fluctuations are also sensitive to the choice of statistical decomposition.

other variables. Under weak regularity conditions, x_t has a structural MA representation:

$$x_t = \mu + S_P(L)e_{Pt} + S_T(L)e_{Tt} \, ,$$

where μ is the unconditional mean, $S_P(L)$ and $S_T(L)$ are matrices of polynomials in the lag operator, and e_{Pt} and e_{Tt} are vectors of mutually orthogonal, serially uncorrelated structural shocks. Permanent and transitory shocks are defined in terms of whether they have long-run effects on the level of output. The vector e_{Pt} represents shocks that have permanent effects on output, while e_{Tt} represents shocks that have only transitory effects. Hence, $s_yS_P(1) \neq 0$ and $s_yS_T(1) = 0$, where s_y is a row vector that selects $\Delta\ln(y_t)$ from x_t. This representation may or may not be fundamental for x_t.

Blanchard and Fisher define the structural trend as "that part of output which is due to permanent shocks." This is

$$\tau_t^{BF} = s_y\{\mu t + [S_P(L)/(1-L)]e_{Pt}\}.$$

Similarly, they define the cyclical component as "that part of output that comes from transitory shocks," and this is

$$c_t^{BF} = s_y[S_T(L)/(1-L)]e_{Tt} \, .$$

The variable c_t^{BF} is the object we seek to approximate.

Measuring it directly is difficult because a structural model is needed to identify permanent and transitory shocks. Various authors emphasize the difficulty of specifying convincing just-identifying restrictions. For example, Hansen and Sargent (1991), Lippi and Reichlin (1993), and Blanchard and Quah (1993) emphasize the importance of prior knowledge about the location of moving average roots. Similarly, Faust and Leeper (1997) emphasize the importance of prior knowledge about the number of shocks. Dynamic macroeconomic theory does not impose general restrictions on either of these features. Given the difficulty of overcoming these identification problems, it seems worthwhile to explore the reliability of some non-structural measures.

Beveridge and Nelson's Forecast-Based Measure

Beveridge and Nelson (1981) propose measures of trend and cycle that are based on long-horizon forecasts. They define the trend as the current level of output plus the undiscounted present value of future excess growth:

$$\tau_t^{BN} = \ln(y_t) + E_t \sum_{j=1}^{\infty} [\Delta\ln(y_{t+j}) - \mu_y] \, ,$$

where μ_y is the unconditional mean of $\Delta\ln(y_t)$. The Beveridge-Nelson cycle is output minus trend:

$$c_t^{BN} = -E_t \sum_{j=1}^{\infty} [\Delta\ln(y_{t+j}) - \mu_y] \, .$$

This definition formalizes the intuition that expected future output growth should be higher than average when output is below trend.

The Beveridge-Nelson method can be implemented by specifying a reduced form forecasting model for output growth and computing long-run forecasts using the Hansen-Sargent (1980) prediction formula. Suppose that an econometrician fits a reduced form model

$$x_t = \mu + \hat{R}(L)\hat{u}_t \, ,$$

where $\hat{R}_0 = I$ and $\hat{u}_t$ is approximately white noise. Then the Beveridge-Nelson cyclical component is

$$c_t^{BN} = s_y[\hat{R}(L) - \hat{R}(1)]\,\hat{u}_t/(1-L) \, .$$

One can rewrite this as the sum of the structural cycle plus some noise:

$$c_t^{BN} - c_t^{BF} = s_y\{[S_P(L) - S_P(1)]e_{Pt}/(1-L) + [S_P(1)e_{Pt} - \hat{R}(1)\hat{u}_t]/(1-L)\} \, .$$

The first noise term reflects the degree to which the structural trend differs from a random walk. The Beveridge-Nelson trend is a random walk, whereas the structural trend may not be. When a permanent shock hits the system, the Beveridge-Nelson method anticipates the total effect of the shock and immediately assigns it to the trend. In contrast, the structural trend incorporates the effects of permanent shocks when output actually moves. The first term vanishes if the structural trend is a random walk, and it is likely to be unimportant if increments in the structural trend are weakly autocorrelated or if the variance of permanent shocks is small.

The second noise term reflects the effects of approximation errors in the forecasting model. This term encompasses a variety of specification errors, including omitted variables, non-fundamental shocks, over-differencing, and premature truncation or other erroneous restrictions on lag polynomials. The literature on Beveridge-Nelson detrending suggests that c_t^{BN} is highly sensitive to alternative reduced form specifications (e.g., compare the univariate and bivariate results in Cochrane 1994). This sensitivity is evidence that the second noise term is often important in practice.

Cochrane's Consumption-Based Measure

Cochrane (1994) suggests using consumption as a measure of the trend in output. This is also a forecast-based approach, but it substitutes consumers' expectations of long-run movements in income for those of an econometrician. This approach is likely to be especially useful if the information available to consumers is superior to that available to econometric forecasters.

If the output-consumption ratio is stationary, output and consumption have the same mean growth rates, shocks that are long-run neutral for output are also long-run neutral for consumption, and $s_y S_P(1) = s_c S_P(1)$, where s_c is a row vector that selects $\Delta\ln(c_t)$ from x_t. In this case, the output-consumption ratio has the following structural MA representation:

$$\begin{aligned}\ln(y_t/c_t) &= (s_y - s_c)[S_P(L)e_{Pt} + S_T(L)e_{Tt}]/(1-L)\\ &= c_t^{BF} - s_c[S_T(L)/(1-L)]e_{Tt}\\ &\quad + (s_y - s_c)[S_P(L)/(1-L)]e_{Pt}\ .\end{aligned}$$

Thus, the output-consumption ratio can also be written as the sum of the structural cycle plus some noise. The first noise term reflects the response of consumption to shocks that are neutral in the long run, and the permanent income hypothesis suggests that this term should have low variance. The second noise term reflects the difference between the dynamic responses of output and consumption to permanent shocks. The permanent income hypothesis suggests that this term may be a problem, especially if a permanent shock represents news that income will be higher at some time in the future. For example, if technology shocks diffuse gradually through the economy, consumption is likely to rise in anticipation of the increase in income. The importance of this term is likely to be an increasing function of the variance of permanent shocks.

While the consumption-based measure is motivated by the permanent income hypothesis, it is important to realize that it is likely to be robust to certain departures from that theory. For example, some economists believe that liquidity or borrowing constraints are important for consumption.[3] Liquidity constraints might limit consumers' ability to smooth transitory shocks, thus increasing the variance of the first noise term, but they would also limit the ability of consumers to borrow in advance of a permanent increase in income, thus reducing the variance of the second. Since liquidity constraints would have offsetting effects on the two sources of approximation error, their net effect on the accuracy of the consumption-based measure is ambiguous.[4]

How does the consumption-based measure compare with the Beveridge-Nelson cycle? Subtracting one from the other yields

$$\begin{aligned}\ln(y_t/c_t) - c_t^{BN} &= -[s_c S_P(L) - s_y S_P(1)]\,e_{Pt}/(1-L)\\ &\quad - s_y[S_P(1)e_{Pt} - \hat{R}(1)\hat{u}_t]/(1-L)\\ &\quad - s_c[S_T(L)/(1-L)]e_{Tt}\ .\end{aligned}$$

The first term vanishes if consumption is a random walk. Since consumption is nearly a random walk, this term is likely to be unimportant. The second and third terms define a trade-off between the Beveridge-Nelson and consumption-based measures. The second term reflects the difference between the information sets available to consumers and econometricians. If consumers have superior knowledge about the structure or better information about recent shocks, their long-run forecasts are likely to be more accurate than those of an econometrician, and their superior knowledge might help mitigate reduced form approximation errors that plague Beveridge-Nelson detrending. Balanced against this advantage is the fact that the consumption-based measure is contaminated by the response of consumption to transitory shocks, which is reflected in the third term.

Linear Filters

Hodrick and Prescott (1997) and Baxter and King (1995) use linear filters to measure business cycles. The Hodrick-Prescott filter is

$$HP(L) = [\lambda(1-L)^2\,(1-L^{-1})^2\,]/[1+\lambda(1-L)^2\,(1-L^{-1})^2],$$

where λ controls the smoothness of the Hodrick-Prescott trend. When applied to quarterly data, λ is almost always set equal to 1600, in which case the Hodrick-Prescott filter approximates a high-pass filter.

Baxter and King propose a similar measure. Following Burns and Mitchell (1946), they impose upper and lower bounds on the duration of cycles: measured from trough to trough, a cycle must last at least 6 quarters and must be no longer than 8 years. Then they use band-pass filters to extract components whose periodicities range from 1.5 to 8 years per cycle.[5]

If we apply a filter, $F(L)$, to data in levels, the corresponding cyclical measure is

3. However, for evidence to the contrary, see Attanasio and Weber (1995), Shea (1995), or Meghir and Weber (1996).

4. In any case, because consumers who are subject to liquidity constraints are likely to engage in dynamic self-insurance, it seems likely that liquidity constraints would have little effect on either term.

5. This motivation should be regarded as heuristic. The Burns-Mitchell restrictions pertain to the duration between troughs, which are dated in terms of certain left-tail events for output growth. For example, one rule of thumb is that the NBER declares a recession whenever output falls for two or more consecutive quarters. Then, when output returns to its pre-recession level, the NBER dates the trough by locating the local minimum. Strictly speaking, the Burns-Mitchell restrictions on duration do not imply the Baxter-King restrictions on periodicity, in the sense that there are other trend-cycle decompositions which are consistent

$$c_t^F = c_t^{BF} - s_y[(1 - F(L))S_T(L)/(1 - L)]e_{Tt}$$
$$+ s_y[F(L)S_P(L)/(1 - L)]e_{Pt} .$$

Filtered cyclical measures can also be written as the sum of the structural cycle plus some noise. The noise terms are unimportant when two conditions are satisfied. First, the spectrum for the structural cycle should have little power outside frequencies for which $|1 - F(\omega)|^2$ is small. Second, the spectrum for growth in the structural trend should have little power outside frequencies for which $|F(\omega)/\Delta(\omega)|^2$ is small, where $\Delta(\omega)$ is the transfer function of the difference operator. If the first condition is violated, $F(L)$ deletes information about the structural cycle that should be included. If the second condition is violated, $F(L)$ includes information about the structural trend that should be deleted.[6] Unless the structural trend and cycle reside at non-overlapping frequency bands, which would make them nearly linearly deterministic, the filtering approach is likely to involve a trade-off between these two sources of noise.

II. DATA GENERATING PROCESSES

The previous section provides some intuition about the noise in each of the non-structural measures. This section and the next execute a number of simulations that are designed to quantify the importance of various approximation errors. This section describes the data generating processes (DGPs) that are used in those simulations.

In specifying DGPs, I have four goals in mind. First, I want to vary the relative importance of permanent and transitory structural components in order to encompass a range of hypotheses discussed in the literature. Second, I want consumption to satisfy a version of the permanent income hypothesis in order to investigate Cochrane's idea. Third, I want to approximate the sample variance and autocorrelations of per capita GDP growth. And fourth, consumption must be smoother than income.

All four goals can be achieved by using a simple real business cycle model to generate data. Many real business cycle models have weak internal propagation mechanisms, so one can generate just about any dynamic pattern in output by putting it in the shocks. Thus, the dynamic properties of output should be regarded as exogenous. The model is really just used to compute the consumption path that is optimal for various output paths. I make no claim that any of the specifications presented below are good models of the business cycle. I claim only that the DGPs generate output paths that describe some prominent hypotheses in the literature, that the associated consumption paths satisfy a version of the permanent income hypothesis, and that the dynamics of output and consumption are plausible.

Artificial data are generated using a two-shock real business cycle model developed by Christiano and Eichenbaum (1992). There is a representative agent whose preferences are:

$$E_t\{\sum_{j=0}^{\infty} \beta^j [\ln(c_{t+j}) + \gamma(N - n_{t+j})]\} ,$$

where c_t is consumption, N is the total endowment of time, n_t is labor hours, and β is the subjective discount factor. Following Christiano and Eichenbaum, I assume that $\beta = 1.03^{-0.25}$ and $\gamma = 0.0037$. There is also a representative firm that produces output by means of a Cobb-Douglas production function:

$$y_t = k_t^{\alpha}(a_t n_t)^{1-\alpha} ,$$

where y_t is the flow of output, k_t is the capital stock, and a_t is a technology shock. The capital stock obeys the usual law of motion:

$$k_{t+1} = (1 - \delta)k_t + i_t ,$$

where δ is the depreciation rate and i_t is gross investment. Christiano and Eichenbaum estimate $\alpha = 0.344$ and $\delta = 0.021$.

The model is driven by technology and government spending shocks. The DGPs differ according to whether technology shocks are trend or difference stationary and according to the relative magnitude of the two shocks. There are four versions.

The first DGP reflects a traditional approach to characterizing business cycles. It assumes that both shocks are neutral in the long run, which makes output trend stationary. TS representations for U.S. GDP have hump-shaped impulse response functions, and the shocks are calibrated to generate this pattern. Detrended technology shocks follow an ARMA(1,4) process:

$$(1) \qquad (1 - \rho_a L)[\ln(a_t) - \mu t] = (1 + \theta_1 L + \theta_2 L^2 + \theta_3 L^3 + \theta_4 L^4)e_{at} ,$$

where $\rho_a = 0.95$, $\theta_1 = 1.6$, $\theta_2 = 0.8$, $\theta_3 = 0.4$, $\theta_4 = 0.2$, and the standard error of e_{at} is 0.00275. The MA polynomial allows technology shocks to diffuse gradually, reaching a peak at a lag of about one year. The AR polynomial ensures that

---

with the Burns-Mitchell durations but which do not satisfy the Baxter-King periodicity restrictions. Section IV gives an example. In the language of Diebold and Rudebusch (1990), the Burns-Mitchell duration restrictions can be formalized in terms of the notion of stochastic weak periodicity, while the Baxter-King restrictions are closer to their definition of stochastic strong periodicity.

6. King and Rebelo (1993) emphasize the first source of noise, while Cogley and Nason (1995b) emphasize the second.

total factor productivity reverts gradually toward its trend. The result is a hump-shaped MA representation, which is shown in the top panel of Figure 1. Government spending shocks follow an AR(2) process around the technology trend:

$$\ln(g_t/a_t) = \bar{g} + e_{gt}/(1 - \rho_1 L)(1 - \rho_2 L)\ , \tag{2}$$

where $\rho_1 = 0.9$, $\rho_2 = 0.45$, $se(e_{gt}) = 0.015$, and $\mathrm{cov}(e_{at},e_{gt}) = 0$. The variable $\ln(g_t/a_t)$ also has a hump-shaped MA representation, which is shown in the bottom panel of the figure.

Since both shocks are transitory, they both drive the structural cycle in output, according to the Blanchard-Fischer definition. They combine to produce the traditional hump-shaped pattern in detrended output, which is shown in the top-left panel of Figure 2.

The second DGP represents a modest perturbation of the first. It sets the AR root in technology equal to one, so that output becomes difference stationary, but it scales the technology and government spending shocks so that most of the variation in output is due to transitory shocks. Technology shocks evolve according to

$$(1 - L)\ln(a_t) = \mu + (1 + \theta_1 L + \theta_2 L^2 + \theta_3 L^3 + \theta_4 L^4)e_{at}\ , \tag{3}$$

where μ and $\theta(L)$ are the same as above but $se(e_{at}) = 0.0005$. Government spending shocks still follow equation (2) but with $se(e_{gt}) = 0.065$. The MA representation for $\ln(a_t)$ is shown in the top panel of Figure 1. Total factor productivity still builds gradually, but it reaches a plateau after about one year and does not revert to trend. The MA representation for $\ln(g_t/a_t)$ has the same shape as before but is scaled up, and it is shown in the bottom panel of Figure 1.

Technology shocks now have a permanent effect on output and drive the structural trend. Government spending shocks have a transitory effect on output and drive the structural cycle. The innovation variances are set so that there is a great deal of trend reversion in output. For example, the upper right panel of Figure 2 illustrates the results of estimating structural VARs using data generated by this model. Most of the variation in output is due to transitory shocks, and permanent shocks are relatively unimportant. Hence, I call this the "Near TS" DGP.

The third DGP has the same structure as the second but alters the innovation variances to increase the importance of permanent shocks and to decrease the importance of transitory ones. The goal is to generate an intermediate degree of trend reversion, which is more consistent with the results of Blanchard and Quah. Accordingly, this DGP increases the magnitude of permanent shocks and reduces the volatility to transitory shocks, so that $se(e_{at}) = 0.0012$

FIGURE 1

Shock Dynamics in the DGPs

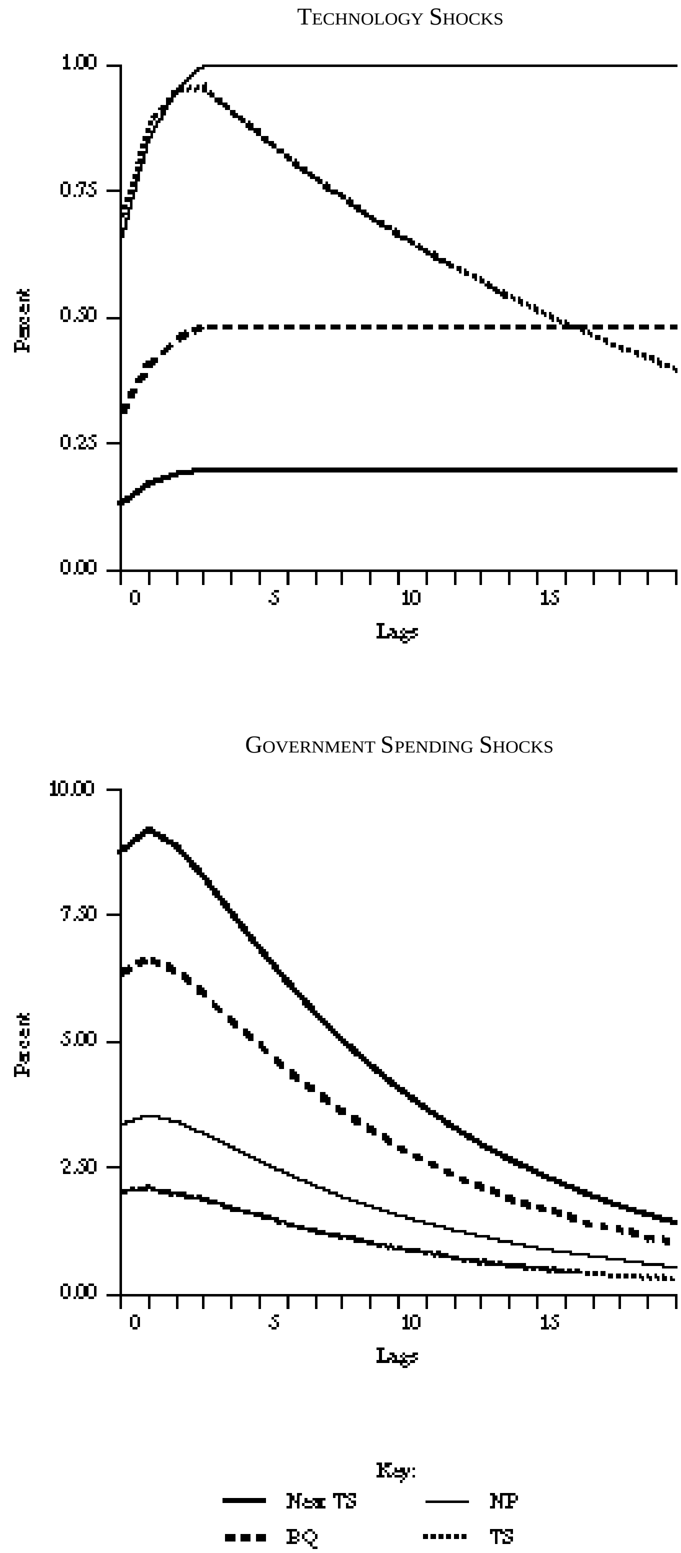

and $se(e_{gt}) = 0.047$. The MA representations for technology and government spending shocks have the same shape as those in the Near TS model but are scaled up and down, respectively (see Figure 1). Similarly, the impulse response functions for output have the same shape, but the permanent component is more variable and the transitory component less variable. For example, structural VAR results for this DGP are shown in the lower left-hand panel of Figure 2.

The fourth DGP retains the structure of the prior two models but scales the shocks to reflect the interpretation of Nelson and Plosser (1982), who argue that most variation

FIGURE 2

OUTPUT DYNAMICS IN THE DGPS

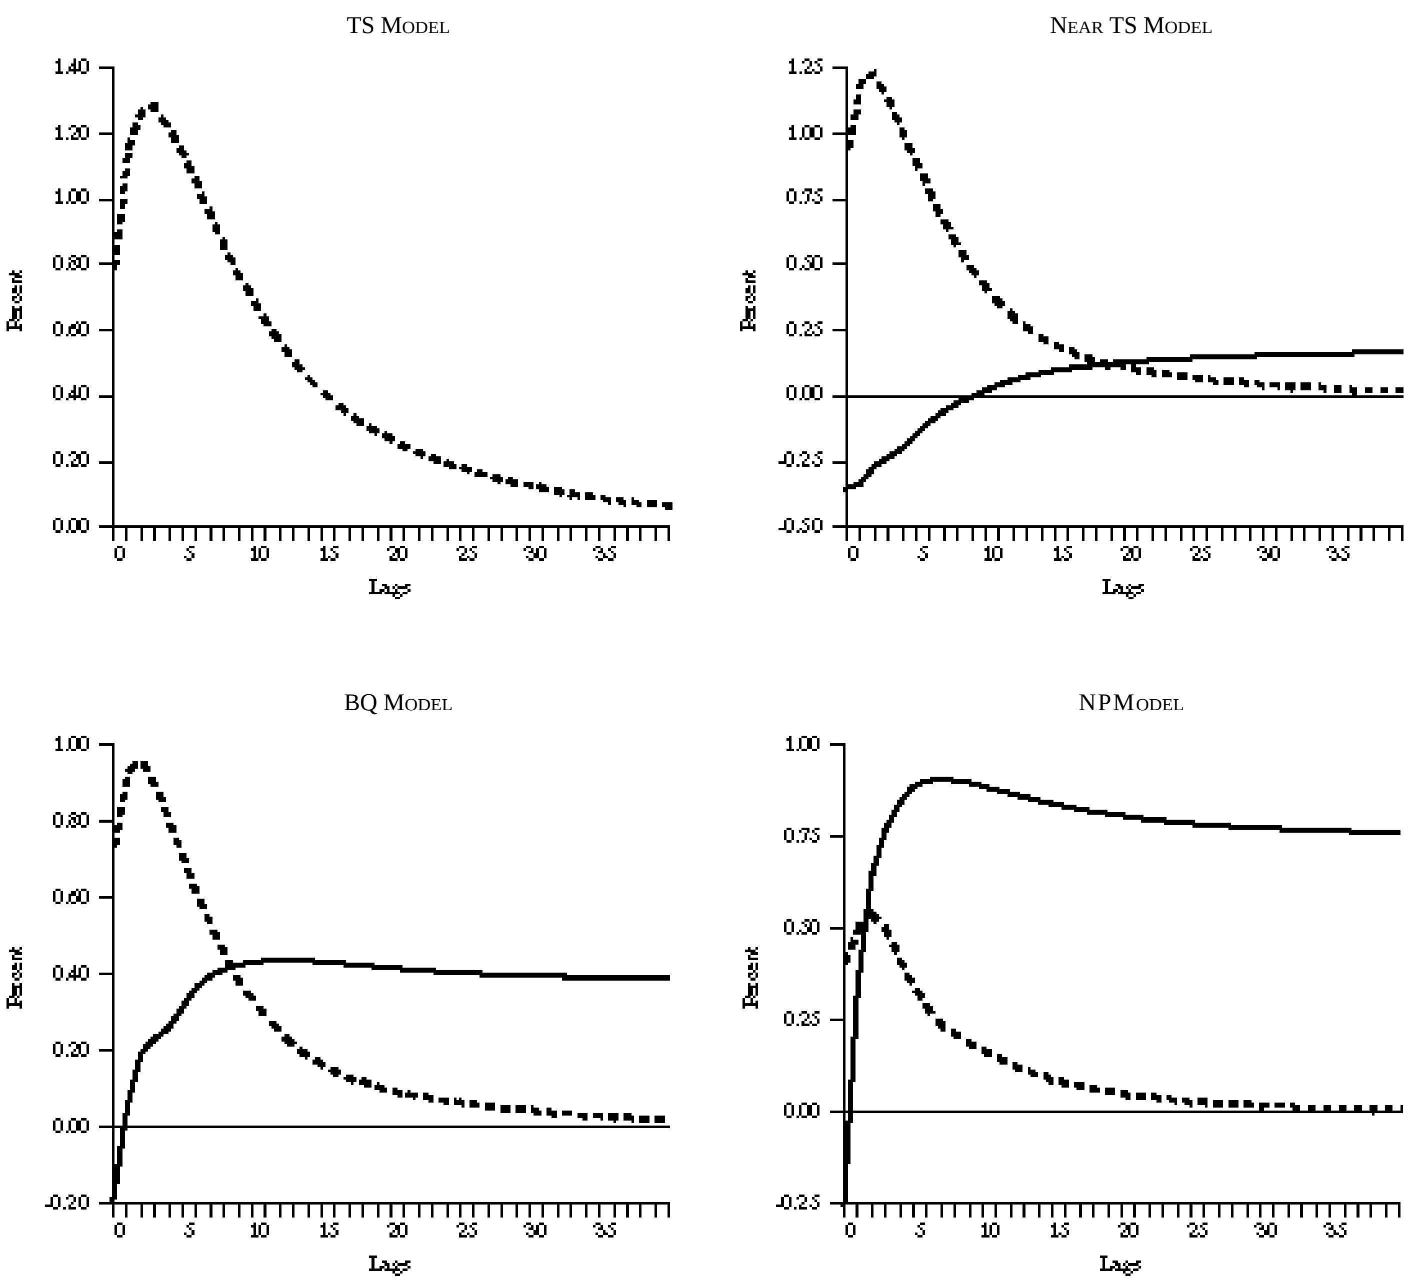

in output growth is due to permanent shocks. To generate data that are consistent with their interpretation, I set $se(e_{at})$ = 0.0025 and $se(e_{gt})$ = 0.025. The MA representations for the shocks are shown in Figure 1, and structural VAR estimates of the impulse response functions for output are shown in the bottom right panel of Figure 2. Consistent with the interpretation of Nelson and Plosser, most of the variation in output is due to the stochastic trend, and the cyclical component is relatively unimportant.

III. EVALUATING THE CYCLICAL INDICATORS

Each of the four DGPs was used to generate 1000 artificial samples consisting of 180 quarterly observations. For each sample, I computed the structural and non-structural measures and compared their properties. The structural cycle was computed by finding the path that output follows when all the permanent shocks are set equal to zero. The Beveridge-Nelson measure was derived from a fourth order VECM for output and consumption.[7] This appears to whiten the residuals in all the DGPs, but it is only an approximation to the true reduced form. I implemented Cochrane's idea by regressing the natural logarithm of output on the natural logarithm of consumption and collecting the residuals, but the simulation results for the output-consumption ratio are essentially the same. Finally, the filters were approximated by truncating at lead and lag 12 and rescaling so that the gain is zero at frequency zero. Since the filters use up the first and last 12 observations in each sample, summary statistics are computed using the middle 156 observations.

The proxies are evaluated along two dimensions. First, how closely do they covary with the structural measure? Second, how closely do their dynamic properties match those of the structural cycle? To evaluate comovements, I compute four statistics: the correlation, coherence, an R^2 statistic, and the fraction of observations in which the proxies and structural measure have the same sign. To evaluate the indicators' dynamic properties, I compare their normalized power spectra with that of the structural cycle.

Do the Weathermen Know Which Way the Wind Blows?

Table 1 reports the median probability that the indicators correctly signal whether the structural cycle is positive or negative.[8] This experiment is relevant to central bankers or other policymakers who wish to pursue countercyclical policies. As Milton Friedman (1953) emphasized many years ago, policymakers cannot effectively lean against the wind unless they know which way it blows. As a benchmark, note that the DGPs are linear and have Gaussian shocks, which implies that structural cycles are symmetric. Hence a strategy of always guessing positive (or negative) would be correct half the time.

The consumption-based indicator is a dominant strategy on this criterion, with the Beveridge-Nelson method coming in a close second. In data drawn from the polar DGPs (TS and NP), the consumption-based and Beveridge-Nelson indicators correctly signal the sign of the structural cycle about 71 to 76 percent of the time. Their success rates rise to about 86 percent when data are drawn from the intermediate DGPs (Near TS and BQ). The Baxter-King and Hodrick-Prescott filters have lower success rates in all DGPs. On average, the consumption-based indicator beats the filters by about 13 percentage points.

Correlation and Coherence

Table 2 reports correlations between the proxies and the structural cycle. The consumption-based indicator is dominant on this criterion as well. It is slightly better than Beveridge-Nelson when the data are drawn from DGPs with complete to moderate degrees of trend reversion (TS, Near TS, and BQ), and it is superior when the stochastic trend dominates (NP). In the latter case, the median correlation for the consumption-based indicator is 0.734, while the median correlation for the Beveridge-Nelson proxy is 0.631. Both of these methods have significantly higher correlations than the Baxter-King and Hodrick-Prescott filters, and the margin of improvement is an increasing function of the importance of the stochastic trend. On average, the consumption-based indicator beats the filters by roughly 27 percentage points.

Coherence functions are shown in Figure 3, and they exhibit the same pattern as the correlation table. When data are drawn from the TS DGP, the Beveridge-Nelson, Hodrick-Prescott, and consumption-based indicators have high coherence with the structural cycle over most of the frequency domain. When data are drawn from the DS DGPs, the coherences fall, but the consumption-based indicator has the highest coherence at all frequencies. Thus the consumption-based indicator appears to be the best measure on this criterion as well.

7. Although other forecasting variables may be useful in practice (e.g., see Evans and Reichlin 1994), they would be redundant in the simulations because the DGPs have only two shocks.

8. Medians and standardized median absolute deviations are reported because there are outliers in some of the simulations, especially when data are drawn from the NP DGP.

TABLE 1

PROBABILITY OF CORRECTLY SIGNALING THE SIGN OF THE STRUCTURAL CYCLE

	TS	NTS	BQ	NP
YC	0.756	0.865	0.865	0.763
	(0.095)	(0.076)	(0.076)	(0.076)
BN	0.737	0.865	0.853	0.712
	(0.095)	(0.076)	(0.076)	(0.095)
HP	0.679	0.731	0.724	0.609
	(0.076)	(0.067)	(0.067)	(0.057)
BK	0.679	0.731	0.718	0.609
	(0.076)	(0.067)	(0.057)	(0.057)

NOTE: This table reports the median probability of correctly guessing the sign of the structural cycle, with standardized median absolute deviations in parentheses. YC denotes the residuals from a regression of output on consumption, BN denotes the Beveridge-Nelson measure, HP denotes the Hodrick-Prescott filter, and BK denotes the Baxter-King filter, while TS, NTS, BQ, and NP denote the various data generating processes.

TABLE 2

CORRELATIONS BETWEEN MEASURED AND STRUCTURAL CYCLES

	TS	NTS	BQ	NP
YC	0.735	0.923	0.912	0.734
	(0.137)	(0.072)	(0.069)	(0.111)
BN	0.696	0.922	0.899	0.631
	(0.167)	(0.072)	(0.072)	(0.183)
HP	0.568	0.674	0.628	0.349
	(0.113)	(0.089)	(0.089)	(0.114)
BK	0.570	0.670	0.624	0.349
	(0.113)	(0.087)	(0.089)	(0.118)

NOTE: This table reports median correlations, $\sigma_{c\hat{c}}/\sigma_c\sigma_{\hat{c}}$, for measured and structural cycles, with standardized median absolute deviations in parentheses. YC denotes the residuals from a regression of output on consumption, BN denotes the Beveridge-Nelson measure, HP denotes the Hodrick-Prescott filter, and BK denotes the Baxter-King filter, while TS, NTS, BQ, and NP denote the various data generating processes.

R^2 Statistics

The top panel of Table 3 reports median R^2 statistics for the proxies. These are defined as $1 - \sigma^2_{c-\hat{c}}/\sigma^2_c$ where $\sigma^2_{c-\hat{c}}$ denotes the error variance and σ^2_c is the variance of the structural cycle.[9] This metric yields a mixed bag: the consumption-based measure works best when data are drawn from the TS DGP, Beveridge-Nelson works best in the intermediate cases (Near TS and BQ), and the Baxter-King filter works best when the stochastic trend dominates (NP).

To understand why the consumption-based indicator is less successful on this metric, rewrite R^2 as

$$R^2 = (\sigma_{\hat{c}}/\sigma_c)(2\rho_{c\hat{c}} - \sigma_{\hat{c}}/\sigma_c)\ ,$$

where $\rho_{c\hat{c}}$ is the correlation between the structural cycle and the indicator. R^2 statistics reward correlation and penalize excessive relative amplitude. Since the consumption-based measure has the highest correlation with the structural cycle, it follows that its R^2 problem reflects errors in relative amplitude.

Table 4 confirms that this is the case. When data are drawn from the TS DGP, the consumption-based measure comes closest to matching the amplitude of the structural cycle. However, when data are drawn from DS DGPs, the consumption-based indicator overstates the structural amplitude by 45 to 75 percent. The Beveridge-Nelson measure does quite well in the intermediate cases (Near TS and BQ), and the filters are best in the NP case. Thus, no single estimator of amplitude is best across all DGPs.

This result is problematic because choosing the best measure of amplitude requires some knowledge of the structure. Some progress toward robustness can be made by using a combined estimator. For example, the last row of Table 4 reports the mean standard error of the four proxies, which is denoted $\sigma_{\tilde{c}}$. This estimator works reasonably well in the intermediate cases but is still off by 25 to 54 percent in the polar cases.

If we rescale the indicators using the combined estimator for amplitude, we can also achieve greater robustness on the R^2 metric. Multiply the indicators by $(\sigma_{\tilde{c}}/\sigma_{\hat{c}})$, so that they have amplitude $\sigma_{\tilde{c}}$. Obviously, this does not affect their sign, their correlation or coherence with the structural cycle, or their dynamic properties, so this does not alter any of the other simulation results. But the scale adjustment does affect R^2 statistics (see the bottom panel of Table 3). Since the rescaled proxies all have the same amplitude, it

9. Note that the error is not the residual from a projection, so this R^2 is not the square of the correlation coefficient. And since the error need not be orthogonal to the model value, this statistic can be negative.

follows that the ranking on this metric is the same as the ranking on the correlation metric. Hence, among the rescaled indicators, the consumption-based measure has the best fit across all four DGPs. More importantly, the rescaled consumption-based indicator compares favorably with the unscaled indicators (compare the top panel with the first row of the bottom panel). Its fit is slightly better than the Beveridge-Nelson measure when the data have complete to moderate degrees of trend reversion (TS, Near TS, and BQ), and it is superior when the stochastic trend dominates (NP). The rescaled consumption-based measure is also superior to the filtered measures across all DGPs.

FIGURE 3

Coherence between Measured and Structural Cycles

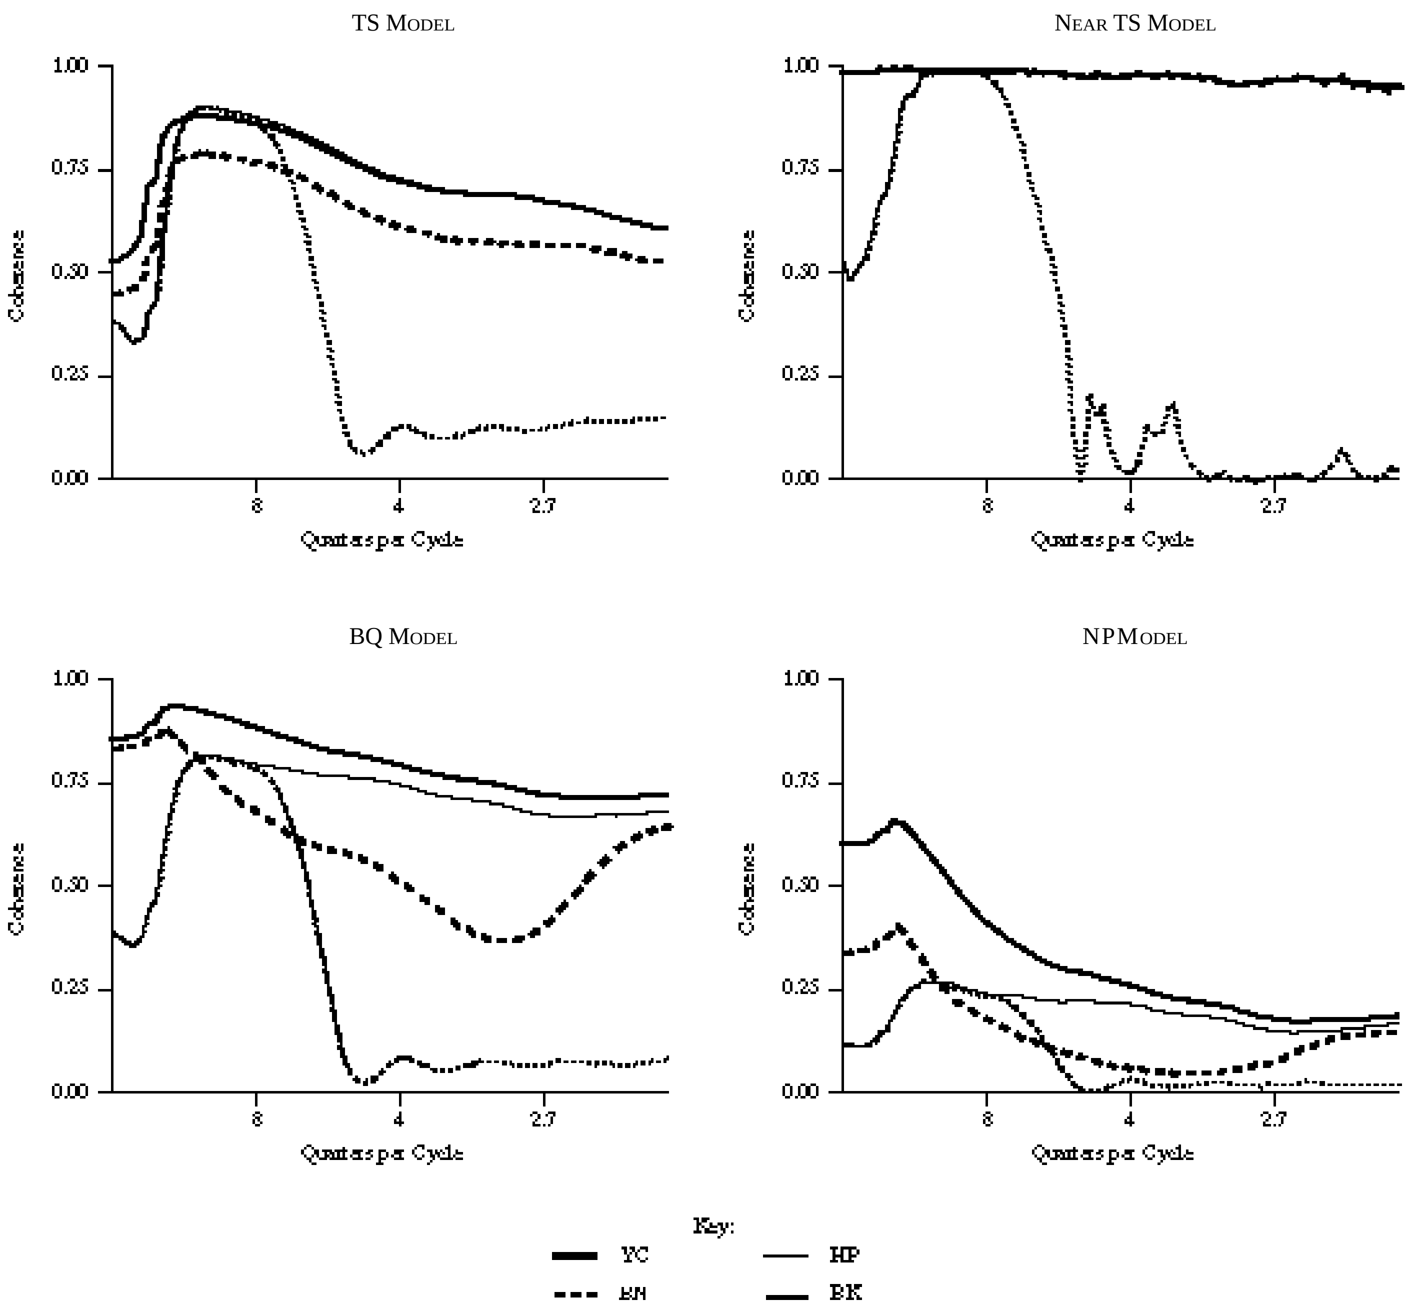

Dynamic Properties

The dynamic properties of the cycle and proxies are summarized by their power spectra. For reasons emphasized above, it is convenient to look separately at the variance and the normalized spectrum (i.e., the spectrum divided by the variance). The variance tells us whether the indicators have the right total amplitude, which we know is a problem. The normalized spectrum tells us whether the indicators have roughly the right autocorrelations.

Figure 4 reports mean normalized spectra for the various indicators and DGPs. Structural cycles have Granger's typical spectral shape: power is concentrated at low frequencies, and maxima occur at frequency zero. The consumption-based and Beveridge-Nelson measures also have this spectral shape, and they appear to match the structural spectra quite well. On the other hand, the filters produce hump-shaped spectra which have little power at low frequencies and which peak at around 7.5 years per cycle. They match the structural spectra less well.

The quality of the match is quantified in Table 5, which reports the root mean square error for $[f_{cc}(\omega) - f_{\hat{c}\hat{c}}(\omega)]$, where $f_{cc}(\omega)$ and $f_{\hat{c}\hat{c}}(\omega)$ denote the spectra for structural cycles and indicators, respectively. The consumption-based measure has the lowest root mean square error in all the DGPs. The Beveridge-Nelson measure is a close second in the intermediate cases (Near TS and BQ), but its fit drops off sharply in the polar cases (TS and NP). Both of these methods produce better results than the filters.

Summary

Collecting results, we can make the following case for the consumption-based measure. Across a variety of trend-cycle structures, it is the most reliable indicator of the sign

TABLE 3

R^2 STATISTICS

	TS	NTS	BQ	NP
		A. UNSCALED INDICATORS		
YC	0.525	0.600	0.525	–0.555
	(0.193)	(0.107)	(0.138)	(0.622)
BN	0.427	0.841	0.777	–0.214
	(0.222)	(0.134)	(0.138)	(0.544)
HP	0.283	0.418	0.384	–0.220
	(0.111)	(0.115)	(0.111)	(0.246)
BK	0.283	0.407	0.374	–0.176
	(0.110)	(0.110)	(0.107)	(0.237)
		B. RESCALED INDICATORS		
YC	0.458	0.840	0.824	0.247
	(0.169)	(0.135)	(0.128)	(0.349)
BN	0.421	0.840	0.801	–0.028
	(0.199)	(0.135)	(0.133)	(0.498)
HP	0.308	0.426	0.336	–0.687
	(0.123)	(0.127)	(0.128)	(0.379)
BK	0.311	0.418	0.329	–0.686
	(0.123)	(0.125)	(0.131)	(0.387)

NOTE: This table reports median R^2 statistics with standardized median absolute deviations in parentheses. Panel A reports statistics for unscaled indicators, and Panel B reports results for indicators that are scaled multiplying by $\sigma_{\bar{c}}/\sigma_{\hat{c}}$, where $\sigma_{\hat{c}}$ is the standard deviation of the individual indicator and $\sigma_{\bar{c}}$ is the mean standard deviation of the four indicators. YC denotes the residuals from a regression of output on consumption, BN denotes the Beveridge-Nelson measure, HP denotes the Hodrick-Prescott filter, and BK denotes the Baxter-King filter, while TS, NTS, BQ, and NP denote the various data generating processes.

TABLE 4

RELATIVE AMPLITUDE OF MEASURED AND STRUCTURAL CYCLES

	TS	NTS	BQ	NP
YC	0.63	1.41	1.45	1.75
	(0.122)	(0.155)	(0.152)	(0.317)
BN	0.50	0.98	0.97	1.37
	(0.156)	(0.112)	(0.174)	(0.409)
HP	0.38	0.49	0.53	0.94
	(0.095)	(0.089)	(0.101)	(0.206)
BK	0.37	0.48	0.52	0.90
	(0.094)	(0.088)	(0.095)	(0.202)
$\sigma_{\tilde{c}}/\sigma_c$	0.47	0.85	0.87	1.25
	(0.104)	(0.091)	(0.102)	(0.231)

NOTE: This table reports the median relative amplitude, $\sigma_{\hat{c}}/\sigma_c$, of the measured and structural cycles, with standardized median absolute deviations in parentheses. The last row reports the median relative amplitude of the combined estimator. YC denotes the residuals from a regression of output on consumption, BN denotes the Beveridge-Nelson measure, HP denotes the Hodrick-Prescott filter, and BK denotes the Baxter-King filter, while TS, NTS, BQ, and NP denote the various data generating processes.

FIGURE 4

Normalized Power Spectra

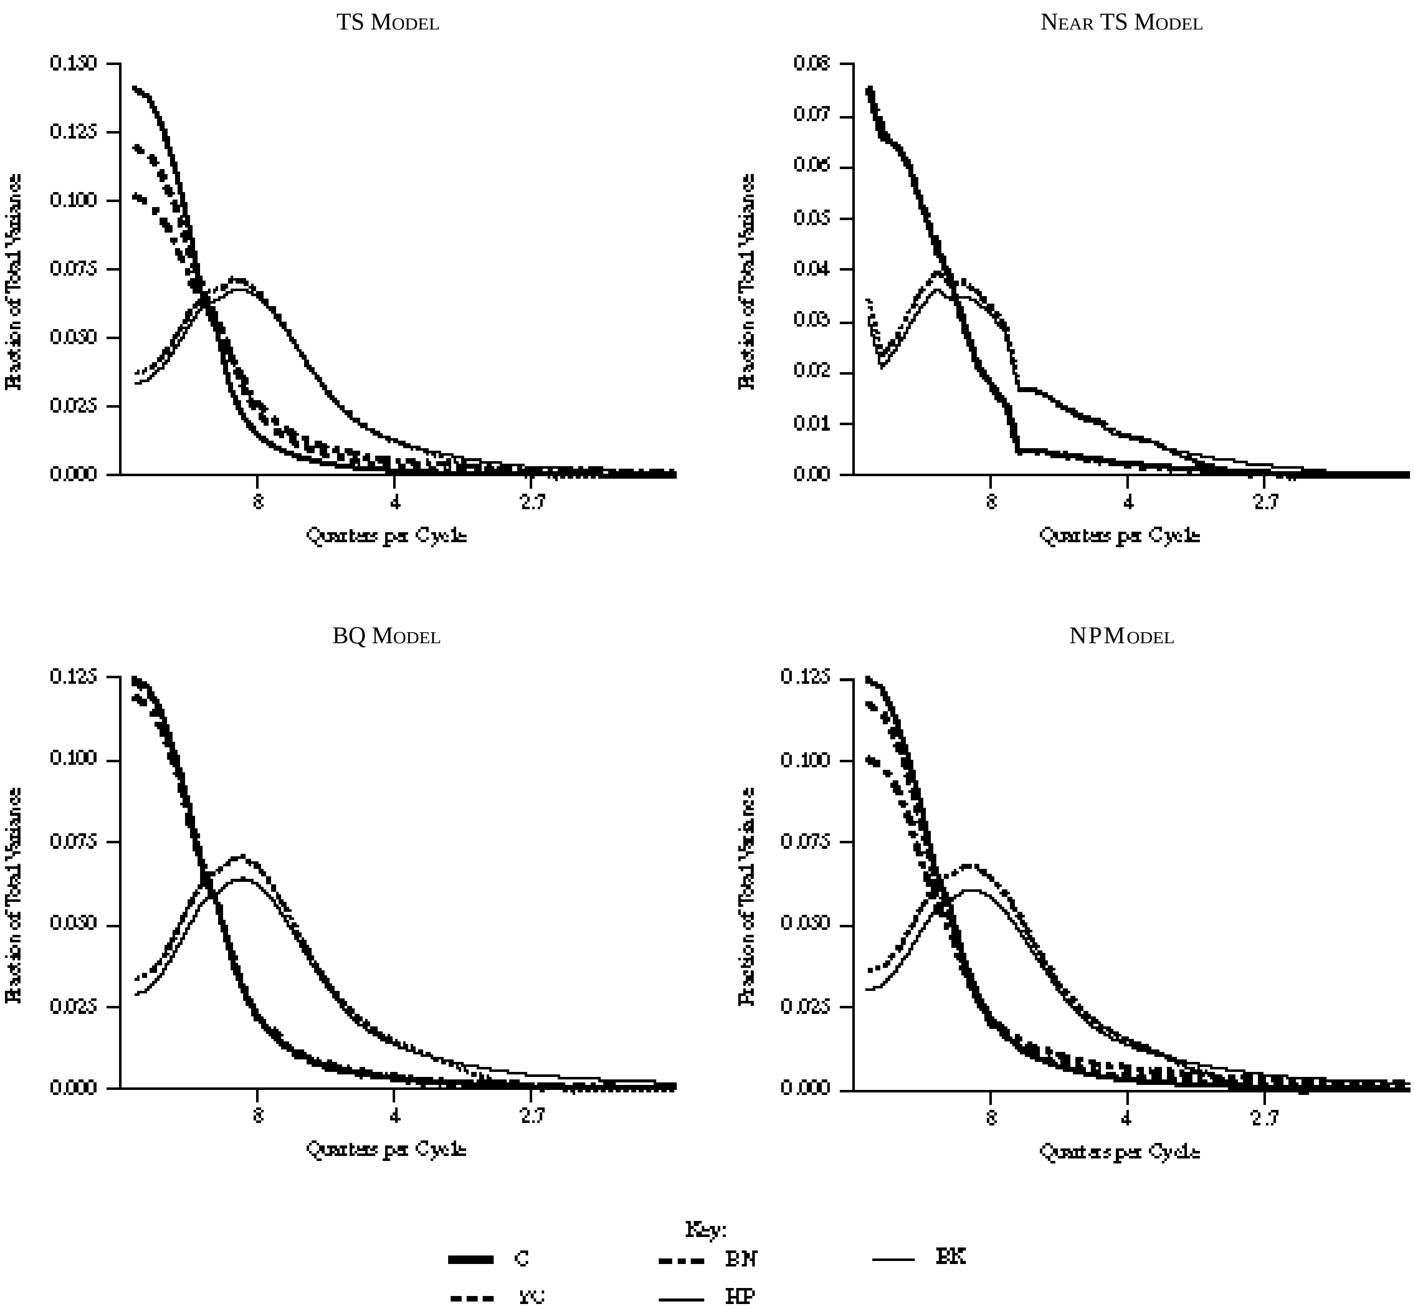

of the structural cycle. It has the highest correlation and coherence with the structural cycle, and in its scale-adjusted form its R^2 statistics are consistently higher than those of the other measures. Finally, the dynamics of the consumption-based indicator, as summarized by its normalized spectra, are the best match for those of the structural cycle.

The chief deficiency of the consumption-based measure is that it often provides poor estimates for the amplitude of the structural cycle. Some progress can be made by rescaling it using a combined estimate of the amplitude. Even so, we should probably take the amplitude measure with a grain of salt.

TABLE 5

DYNAMIC PROPERTIES OF MEASURED AND STRUCTURAL CYCLES

	TS	NTS	BQ	NP
YC	0.416 (0.278)	0.256 (0.128)	0.231 (0.147)	0.263 (0.161)
BN	0.663 (0.309)	0.257 (0.126)	0.256 (0.179)	0.642 (0.432)
HP	1.53 (0.164)	1.22 (0.157)	1.31 (0.155)	1.30 (0.147)
BK	1.50 (0.176)	1.18 (0.174)	1.29 (0.172)	1.26 (0.169)

NOTE: This table reports the median root mean square error (times 100) for $[f_{cc}(\omega) - f_{\hat{c}\hat{c}}(\omega)]$, with standardized median absolute deviations in parentheses. YC denotes the residuals from a regression of output on consumption, BN denotes the Beveridge-Nelson measure, HP denotes the Hodrick-Prescott filter, and BK denotes the Baxter-King filter, while TS, NTS, BQ, and NP denote the various data generating processes.

IV. CONSUMPTION-BASED BUSINESS CYCLE MEASURES FOR THE UNITED STATES

This section applies the consumption-based method to post-war data for the U.S. The first part illustrates an additional virtue of the consumption-based measure, viz., that it automatically adjusts for trend or drift breaks in the data. The second part confirms that the consumption-based measure conforms with widely shared priors about the dates of U.S. business cycles. Then, following Prescott (1986), I summarize the properties of consumption-based measures in terms of their periodicity, comovement, and relative volatility. This yields a surprising result: according to this measure, productivity appears to be sharply *countercyclical.*

Trend/Drift Breaks

Perron (1989) hypothesizes that there are occasional, exogenous changes in the mean rate of growth in output.[10] For example, he reports evidence of a trend break in 1929 in the Nelson-Plosser data set and another in 1973.Q1 in post-war quarterly data. Dealing with a trend or drift break is straightforward if the date of the break is known a priori, but problematic if the date is unknown. For example, Christiano (1992) discusses the complications that arise when the date of the trend break must be estimated.

The consumption-based approach provides a very simple way to circumvent this problem. For, since consumption and income share a common trend, if there is a trend or drift break in one, there will be a break in the other at roughly the same time. Therefore, using consumption to detrend income automatically adjusts for any trend breaks that may be present and does not require prior knowledge about the dates of those breaks.

This virtuous property is illustrated in Figure 5, which compares consumption-based measures of the cyclical component of output with data that are detrended in accordance with Perron's results. The top panel reports results for the private sector, while the bottom panel illustrates results for data that include the government. Total output is measured by real per capita GDP, and private output consists of this measure minus real per capita government expenditures. Consumption is measured by real per capita expenditures on non-durables and services, but the results are robust to the addition of consumer durables.[11] The consumption-based measure was estimated by regressing the natural logarithm of output on the natural logarithm of consumption.[12] The scale-adjustment factors for output were nearly unity, so I chose to report unscaled indicators. The data are quarterly, and the sample period is 1954.Q1 to 1994.Q4. Shaded areas mark the dates of recessions, as determined by the NBER.

The figure illustrates two salient points. First, the consumption-based measure approximates the trend-break model rather well. The correlation between the two is 0.75 for private sector GDP and 0.68 for total GDP, and the Beveridge-Nelson variance ratios, $\text{var}(\Delta\tau_t^{BN})/\text{var}(\Delta \ln y_t)$, are 0.28 and 0.31 for private and total GDP, respectively. The latter suggests there is a significant degree of trend reversion in output. The close correspondence between consumption-based and trend-break measures of the business cycle follows from the fact that consumption is smoother than income and that the productivity slowdown begins to show up in consumption data around the same time as in output data.

Second, in one respect, the consumption-based measure appears to be superior to Perron's model. The latter misses

10. This is even more evident in data from other countries (e.g., see Cogley 1990).

11. I chose to exclude durables because they represent household investment rather than consumption. Non-durables and services also provide marginally better forecasts of future output growth.

12. The private sector cointegrating coefficient is close to 1, so this is essentially the same as the great ratio. The addition of the government sector pushes the cointegrating coefficient away from 1.

FIGURE 5

ADJUSTING FOR TREND BREAKS

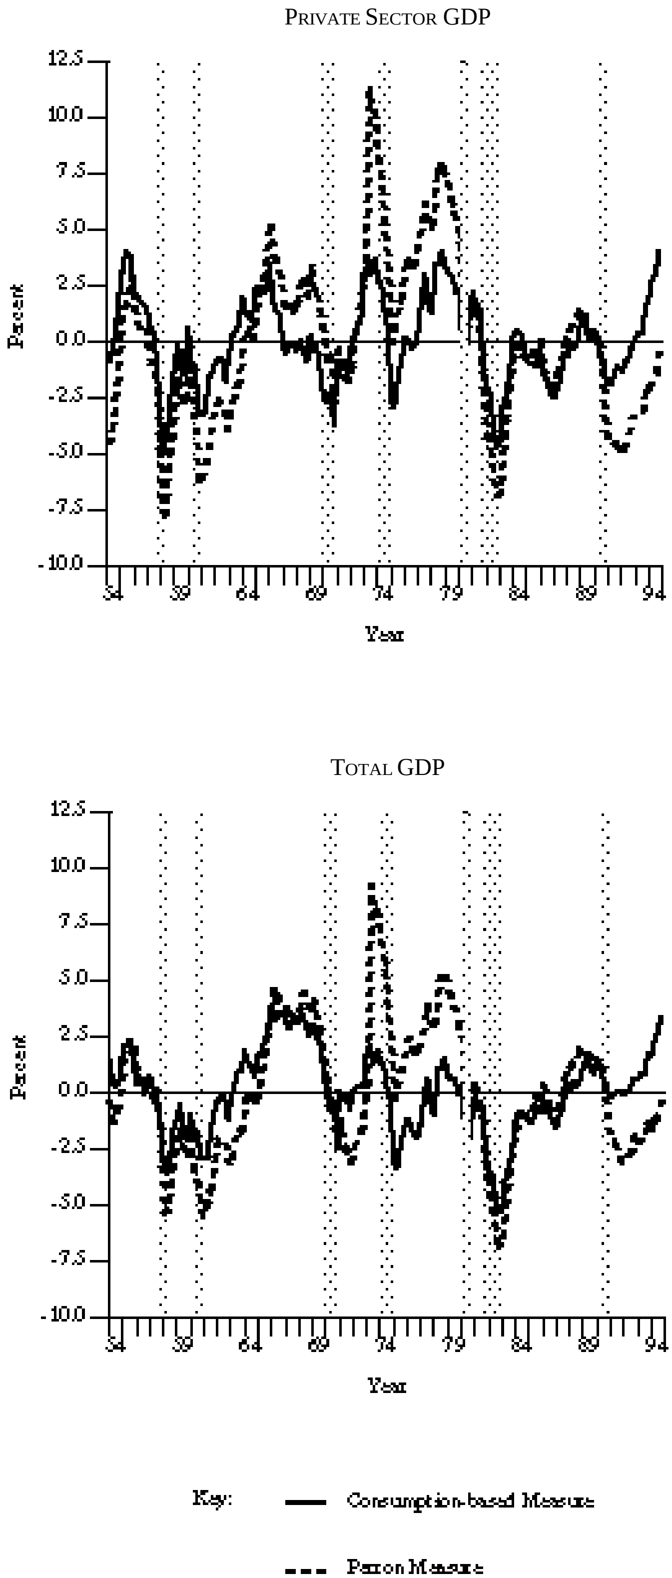

the 1974 recession, which by most accounts was one of the more severe recessions in the post-war period. This probably reflects uncertainty about the precise date or nature of the trend break. In contrast, the consumption-based measure fits this recession well. Because the consumption-based approach relies on less information about the date and nature of trend breaks, it is likely to be more reliable around the dates of those events.

Conformity with NBER Dates

One of the primary objectives of Baxter and King was to produce a measure that is broadly consistent with the findings of the NBER, and it is important to confirm that consumption-based measures also satisfy this criterion.

Figure 6 reproduces the cyclical components of output and also shows consumption-based measures of the cyclical components of investment, average labor productivity, and hours. Investment is measured by gross private domestic investment, and productivity and hours are each measured in two ways, in accordance with the two measures of GDP discussed above. The solid lines report results for the private sector, while the dotted lines are for data that include the government. Total hours are measured by hours worked in non-agricultural establishments, and private hours subtract out the government sector. Average productivity is non-farm output (total and private) divided by the relevant measure of hours. All the variables are real and are measured in per capita units.

The cyclical components of investment and average productivity were estimated by regressing their natural logarithms on the natural logarithm of consumption. The rationale for regressing investment and productivity on consumption is a belief that these variables share a common trend with output. Hence, if consumption measures the trend in output, it also measures the trend in investment and productivity. The cyclical component of hours was estimated by subtracting the cyclical component of productivity from the cyclical component of output, so that the cyclical components of the various series satisfy the productivity identity.

The shaded areas in Figure 6 mark the dates of NBER recessions. The cyclical components of output, investment, and hours generally conform quite closely to the NBER dates, and there do not appear to be any false signals of recession. The troughs in output and investment often occur in the same quarter as the NBER trough and always occur within plus or minus one quarter. In 6 of 8 recessions, the trough in hours also occurs within one quarter of the NBER date. The two exceptions are the 1970 recession, when the trough in hours lagged the NBER date by three quarters, and the 1990–91 recession, when it lagged by six.

FIGURE 6

CONSUMPTION-BASED MEASURES FOR THE U.S.

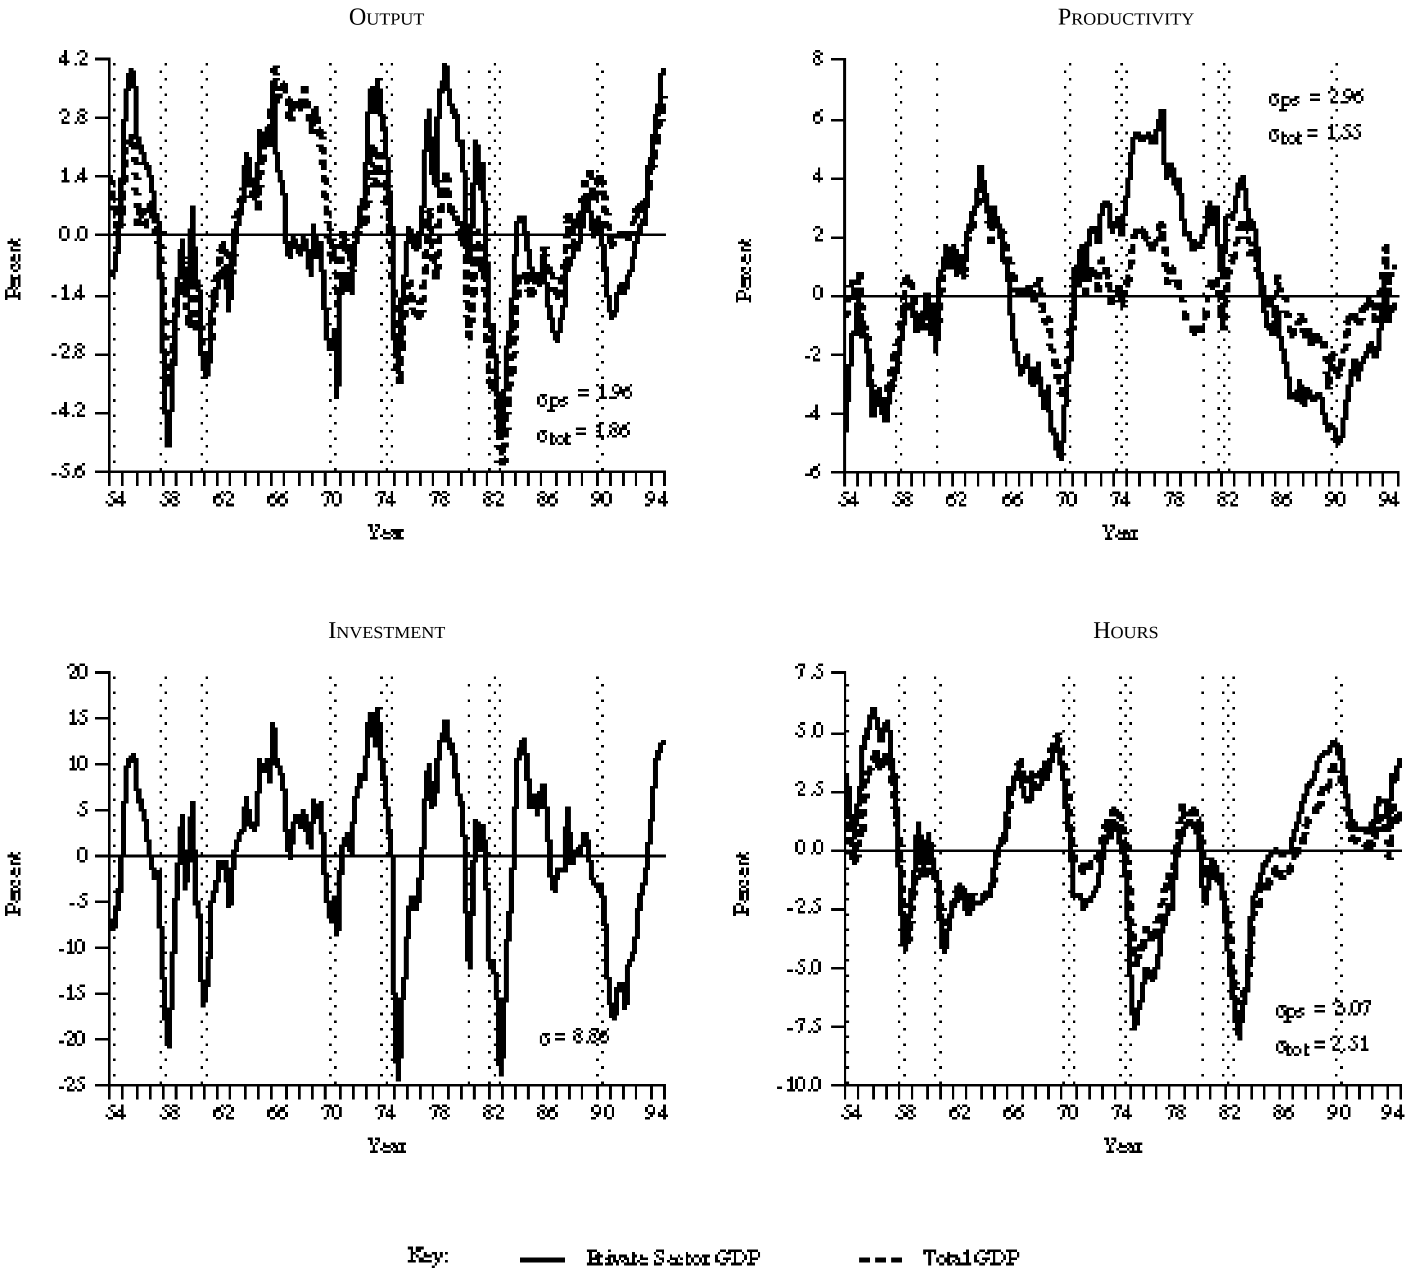

The peaks in output, investment, and hours systematically lead the NBER dates by one or two years. I do not regard this as an embarrassment because the peaks in detrended data should lead the NBER peaks, which are dated using the level of the series (see Romer 1994). Growth that is positive but slower than trend continues to raise the level of output but reduces the deviation from trend. If growth tends to slow at the end of expansions, peaks in detrended data will systematically lead NBER peaks.

Average labor productivity varies countercyclically. Productivity tends to grow more slowly than trend during the latter half of expansions, often reaching a local trough within plus or minus two quarters of an NBER peak. This illustrates the end-of-expansion productivity effect, first noted by Gordon (1979). Then, during recessions, productivity tends to rise, usually reaching a peak within a year after the NBER trough.

Relative Volatility, Periodicity, and Comovement

Figure 6 shows that investment is by far the most volatile series, with a standard error that is roughly 4.75 times larger than the standard error for output. Hours are approximately 35 to 50 percent more variable than output, depending on whether total or private sector measures are used. Average productivity is about 50 percent more variable than output in the private sector, and it is about 20 percent less variable than output when government is included. This is a fairly standard ranking for relative volatility.

The periodic properties of the four measures are summarized by their autocorrelation functions, which are shown in Figure 7. Consumption-based measures are highly per-

FIGURE 7

AUTOCORRELATIONS FOR CONSUMPTION-BASED MEASURES

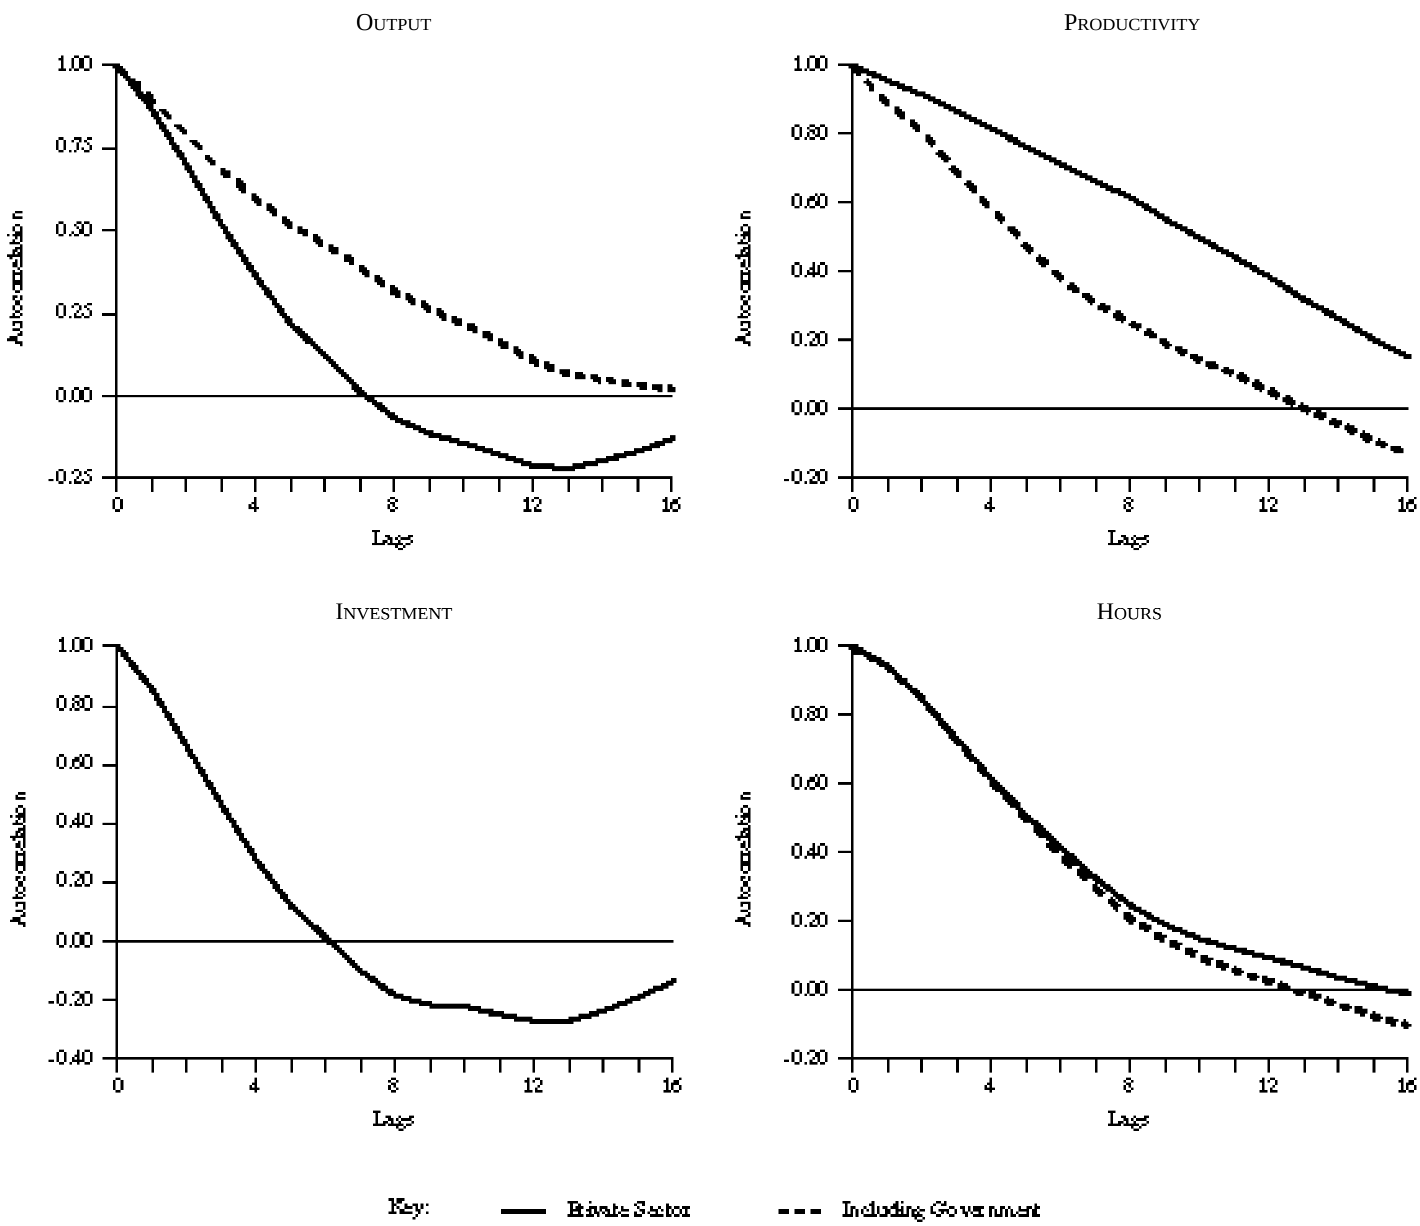

sistent. Their first-order autocorrelations are around 0.85 to 0.9, and higher-order autocorrelations decay slowly. The cyclical components of output, productivity, and hours have Granger's typical spectral shape: their power spectra are concentrated at low frequencies and have peaks at frequency zero. The spectrum for investment is concentrated at medium frequencies, with a peak at around 6.8 years per cycle. Low or medium frequency filtering would eliminate much of the cyclical variability in output, productivity, and hours, but it would do less harm to cyclical movements in investment.

Comovements among the four series are summarized by their cross-correlation functions, which are shown in Figure 8. These are defined as $Ex_t y_{t-k}$, so positive values of k

FIGURE 8

CROSS-CORRELATIONS AMONG CONSUMPTION-BASED MEASURES

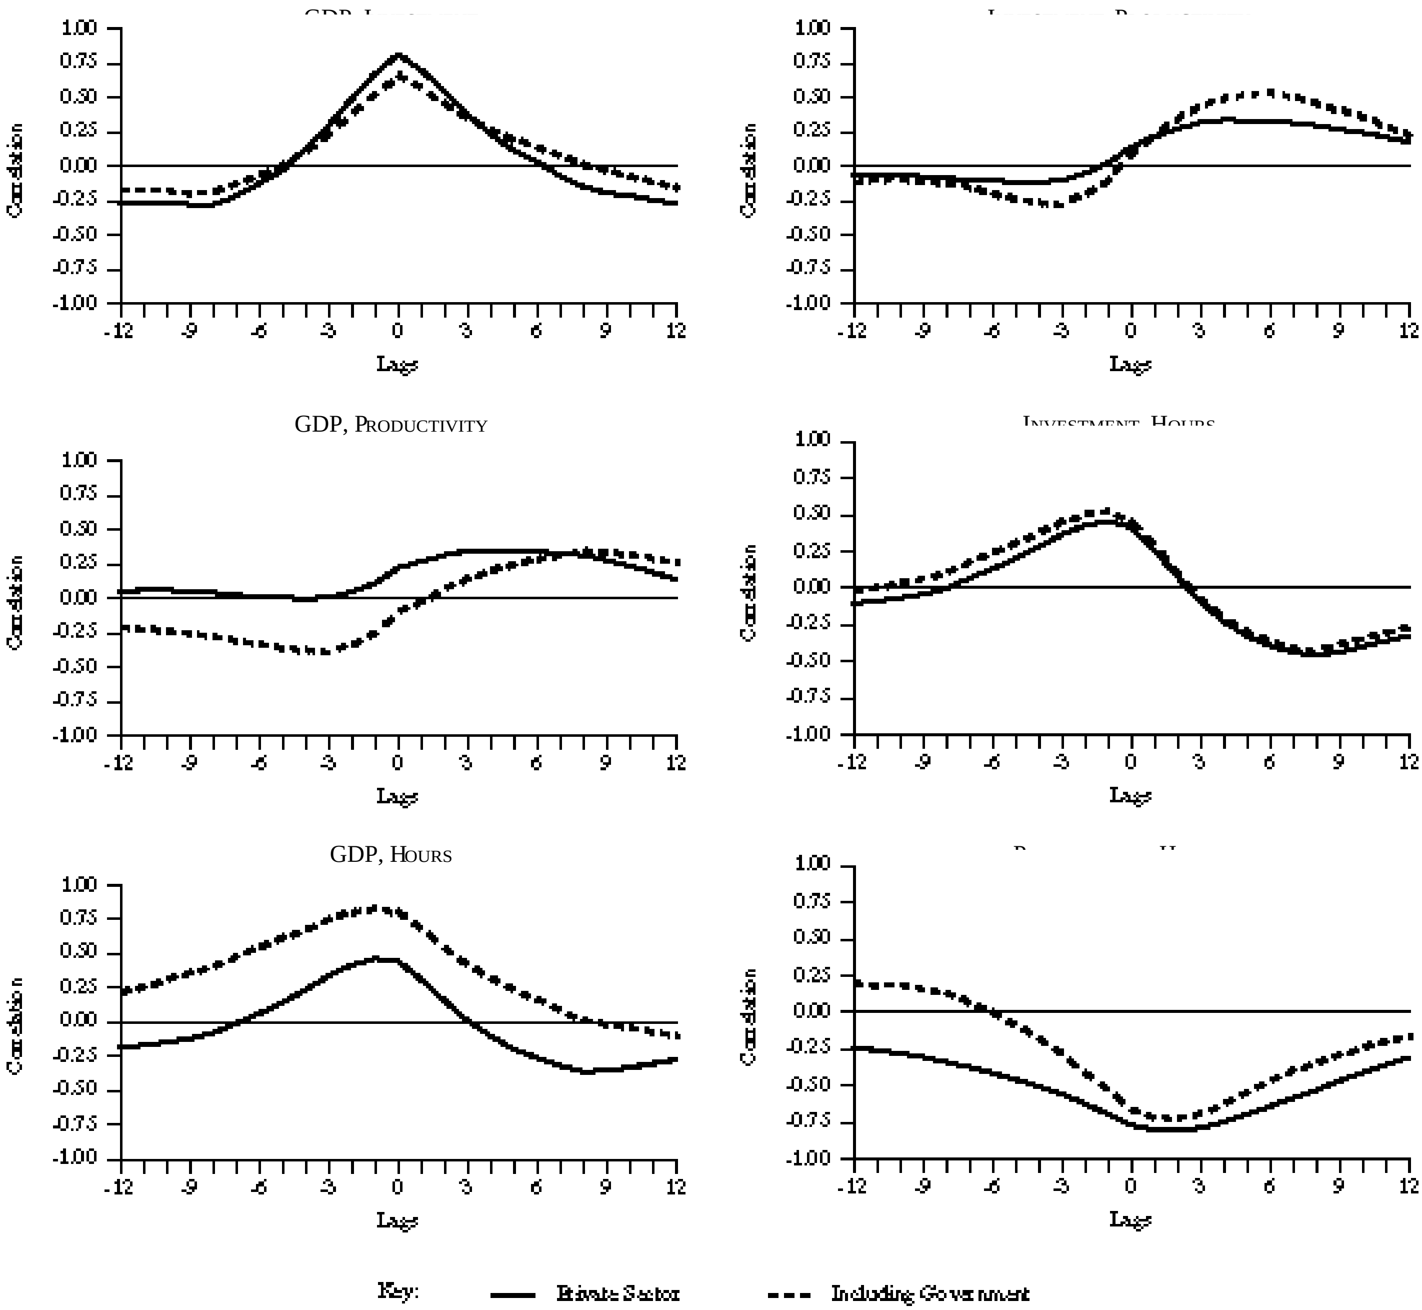

refer to the correlation between current values of y_t and future values of x_t, while negative values of k refer to the correlation between current values of x_t and future values of y_t. Investment and hours covary positively with output and are nearly in phase, while productivity is out of phase with output. Cyclical increases in productivity are nearly uncorrelated with contemporaneous movements in output and investment, but they tend to lead increases in output and investment by roughly one to two years.

The most interesting fact in Figure 8 is the negative correlation between average productivity and hours: the contemporaneous correlations are –0.66 and –0.77 for total and private measures, respectively, and high levels of hours forecast low productivity a few quarters hence. This fact is significant because it deepens the challenge to models in which technology shocks are the dominant source of cyclical fluctuations.

One-shock real business cycle models predict that average productivity should covary positively with hours, since a positive technology shock shifts labor demand outward. Christiano and Eichenbaum (1992) show that this prediction is counterfactual: while one-shock models imply a correlation in excess of 0.9 in HP filtered data, the sample correlation is around zero.[13] Productivity is less procyclical in multi-shock real business cycle models. Roughly speaking, if there are shocks that shift labor supply as well as labor demand, the productivity-hours correlation could be positive, negative, or zero depending on the relative magnitude of the shocks and on the dynamic responses of output and hours. For example, in the Christiano-Eichenbaum model, technology shocks still generate procyclical movements in productivity and hours, but government spending shocks generate countercyclical movements. The latter effect reduces the unconditional correlation to about 0.58 in HP filtered data. The labor hoarding model of Burnside, Eichenbaum, and Rebelo (1993) generates endogenous dynamic responses in output and hours, and this reduces the unconditional correlation even further, to about 0.04 in HP filtered data.

When the consumption-based approach is applied to data generated by the Burnside-Eichenbaum-Rebelo model, the hours-productivity correlation is 0.25 with a standard error of 0.16. Thus the model still suggests that productivity is weakly procyclical. But the sample correlation is around –0.7, which suggests that productivity is actually strongly countercyclical. Hence the consumption-based measure deepens the productivity puzzle. This fact suggests that shocks which generate movements along a neoclassical production function are more important for business cycles than shocks which shift production functions. It reinforces Christiano and Eichenbaum's message that it is important to find measurable economic impulses which generate countercyclical variation in productivity.

V. Conclusion

This paper adopts a structural definition of the business cycle and interprets non-structural measures as noisy indicators. It calibrates a variety of trend-cycle structures and looks for an indicator that is robust. A consumption-based measure, originally proposed by Cochrane, seems to work best. Across a variety of structures, it has the highest coherence and correlation with the structural cycle, and its dynamics are the best match for those of the structural cycle. Its chief weakness is that it often over-estimates the amplitude of the cycle, and this can be problematic when the stochastic trend is the dominant source of output fluctuations.

When applied to post-war U.S. data, the consumption-based indicators conform well with the dates of NBER recessions. They also suggest that productivity and hours covary negatively over the business cycle. This fact is consistent with diminishing marginal productivity, and it reinforces the challenge to models in which business cycles are driven primarily by technology shocks.

Finally, this analysis could be extended by adding to the lists of structures and indicators. Naturally, as our thinking about plausible structures evolves, so too will our thinking about useful indicators.

13. It is either weakly positive or weakly negative depending on how hours are measured and on the sample period.

REFERENCES

Attanasio, Orazio P., and Gugliemo Weber. 1995. “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey.” *Journal of Political Economy* 103(6) (December) pp. 1121–1157.

Baxter, Marianne, and Robert G. King. 1995. “Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series.” NBER Working Paper No. 5022, February.

Beveridge, Stephen, and Charles R. Nelson. 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(2), pp. 151–174.

Blanchard, Olivier Jean, and Stanley Fischer. 1989. *Lectures on Macroeconomics.* Cambridge: MIT Press.

__________, and Danny Quah. 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79(4) pp. 655–673.

__________, and __________. 1993. “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Reply.” *American Economic Review* 83(3) pp. 653–658.

Burns, Arthur M., and Wesley C. Mitchell. 1946. *Measuring Business Cycles.* New York: National Bureau of Economic Research.

Burnside, Craig, Martin Eichenbaum, and Sergio T. Rebelo. 1993. “Labor Hoarding and the Business Cycle.” *Journal of Political Economy* 101(2) pp. 245–273.

Christiano, Lawrence J. 1992. “Searching for a Break in GNP.” *Journal of Business and Economic Statistics* 10(3) pp. 237–250.

__________, and Martin Eichenbaum. 1992. “Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations.” *American Economic Review* 82(3), pp. 430–450.

Cochrane, John H. 1994. “Permanent and Transitory Components of GNP and Stock Prices.” *Quarterly Journal of Economics* 109(1) pp. 241–265.

Cogley, Timothy. 1990. “International Evidence on the Size of the Random Walk in Output.” *Journal of Political Economy* 98 (3) pp. 501–518.

__________, and James M. Nason. 1995a. “Output Dynamics in Real Business Cycle Models.” *American Economic Review* 85(3) pp. 492–511.

__________, and __________. 1995b. “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research.” *Journal of Economic Dynamics and Control* 19(1–2) pp. 253–278.

Diebold, Francis X., and Glenn D. Rudebusch. 1990. “A Nonparametric Investigation of Duration Dependence in the American Business Cycle.” *Journal of Political Economy* 98(3) pp. 596–616.

Evans, George, and Lucrezia Reichlin. 1994. “Information, Forecasts, and Measurement of the Business Cycle.” *Journal of Monetary Economics* 33(2) pp. 233–254.

Faust, Jon, and Eric Leeper. 1997. “When Do Long-Run Identifying Restrictions Give Reliable Results?” *Journal of Business and Economic Statistics,* forthcoming.

Friedman, Milton. 1953. “The Effects of a Full Employment Policy on Economic Stability: A Formal Analysis.” In *Essays in Positive Economics.* Chicago: University of Chicago Press.

Gordon, Robert J. 1979. “The ‘End-of-Expansion’ Phenomenon in Short-Run Productivity Behavior.” *Brookings Papers on Economic Activity* 2, pp. 447–61.

Hansen, Lars P., and Thomas J. Sargent. 1991. “Two Difficulties in Interpreting Vector Autoregressions.” In *Rational Expectations Econometrics.* Boulder, CO: Westview Press.

__________, and __________. 1980. “Formulating and Estimating Dynamic Linear Rational Expectations Models.” *Journal of Economic Dynamics and Control* 2(1) pp. 7–46.

Harvey, Andrew C., and Albert Jaeger. 1993. “Detrending, Stylized Facts, and the Business Cycle.” *Journal of Applied Econometrics* 8(3) pp. 231–247.

Hodrick, Robert, and Edward C. Prescott. 1997. “Postwar U.S. Business Cycles: An Empirical Investigation.” *Journal of Money, Credit, and Banking* 29(1) pp. 1–16.

King, Robert G., and Sergio T. Rebelo. 1993. “Low Frequency Filtering and Real Business Cycles.” *Journal of Economic Dynamics and Control* 17 (1–2) pp. 207–231.

Lippi, Marco, and Lucrezia Reichlin. 1994a. “Diffusion of Technical Change and the Decomposition of Output into Trend and Cycle.” *Review of Economic Studies* 61(1) pp. 19–30.

__________, and __________. 1994b. “Common and Uncommon Trends and Cycles.” *European Economic Review* 38(3–4) pp. 624–635.

__________, and __________. 1993. “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment.” *American Economic Review* 83 (3) pp. 644–652.

Meghir, Costas, and Gugliemo Weber. 1996. “Intertemporal Nonseparability or Borrowing Restrictions? A Disaggregate Analysis Using a U.S. Consumption Panel.” *Econometrica* 64(5) pp. 1151–1181.

Nelson, Charles R., and H. Kang. 1981. “Spurious Periodicity in Inappropriately Detrended Time Series.” *Econometrica* 49(3) pp. 741–751.

__________, and Charles I. Plosser. 1982. “Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications.” *Journal of Monetary Economics* 10(2) pp. 139–162.

Perron, Pierre. 1989. “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis.” *Econometrica* 57(6) pp.1361–1402.

Prescott, Edward C. 1986. “Theory Ahead of Business-Cycle Measurement.” *Carnegie-Rochester Conference Series on Public Policy* 25 (Autumn) pp. 11–44.

Quah, Danny. 1992. “The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds.” *Econometrica* 60(1) pp. 107–118.

Romer, Christina. 1994. “Remeasuring Business Cycles.” *Journal of Economic History* 54(3) pp. 573–609.

Shea, John. 1995. “Union Contracts and the Life-Cycle/Permanent Income Hypothesis.” *American Economic Review* 85(1) pp. 186–200.