When Can Trend-Cycle Decompositions Be Trusted?

Manuel Gonzalez-Astudillo[†]
manuel.p.gonzalez-astudillo@frb.gov
Federal Reserve Board
Washington, D.C., USA

John M. Roberts
john.m.roberts@frb.gov
Federal Reserve Board
Washington, D.C., USA

February 25, 2015

Abstract

In this paper, we examine the results of GDP trend-cycle decompositions from the estimation of bivariate unobserved components models that allow for correlated trend and cycle innovations. Three competing variables are considered in the bivariate setup along with GDP: the unemployment rate, the inflation rate, and gross domestic income. We find that the unemployment rate is the best variable to accompany GDP in the bivariate setup to obtain accurate estimates of the cycle. We show that the key feature of unemployment that allows for precise estimates of the cycle of GDP is that its nonstationary component is "small" relative to its stationary component. Using quarterly GDP and unemployment rate data from 1948:Q1 to 2013:Q4, we obtain the trend-cycle decomposition of GDP and find no evidence to reject orthogonality between trend and cycle components.

Keywords: Unobserved components model, trend-cycle decomposition, trend-cycle correlation

JEL Classification Numbers: C13, C32, C52

[*]The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System.

[†]Corresponding author. We thank Timothy Hills for speedy and accurate research assistance.

1 Introduction

When can we trust trend-cycle output decompositions? Univariate studies, such as Morley et al. (2003) (MNZ hereafter), typically find a large negative correlation between the innovations of trend and cycle as well as small and economically unimportant business cycles. However, the Monte Carlo studies by Basistha (2007) and Wada (2012) show that this strong correlation is spurious. In particular, they find that the correlation is estimated to be large even when it is zero in the data generating process assumed in the Monte Carlo experiment. As explained by Wada (2012), the reason the univariate estimation yields this (incorrect) result turns on the assumption of a positive definite variance-covariance matrix of the innovations, which makes the likelihood function of the unobserved components (UC) model have its global maximum at somewhere other than the true parameter values.

In contrast to the results from univariate studies, studies that include additional variables such as inflation or unemployment typically find estimates of the cyclical component of output that are largely conventional, closely resembling, for example, estimates published by the Congressional Budget Office (CBO). Moreover, estimated correlations between GDP trend and cycle are smaller than in MNZ and often statistically insignificant. Early examples include Clark (1989), who added the unemployment rate, and Roberts (2001), who added inflation and hours to the model and estimated separate trends for hours and productivity. Both these studies found the trend-cycle correlations to be small and statistically insignificant. Basistha and Nelson (2007) and Basistha (2007) focus on inflation as an additional variable. They find trend-cycle correlations that were larger than in Clark and Roberts but smaller than in MNZ; their estimated business cycles are conventional. Basistha (2007) performed a set of Monte Carlo simulations to examine the usefulness of a bivariate set up of the kind in Basistha and Nelson (2007). Results indicated that a univariate specification is not able to identify the correlation coefficient between trend and cycle innovations, whereas the bivariate setup yields an estimated correlation coefficient that is close on average to the true correlation.

The broad consensus among multivariate studies is encouraging, but the experience with univariate decompositions suggest that it is important to examine the small-sample properties of trend-cycle estimators. In this paper, we use Monte Carlo experiments to explore under what conditions an auxiliary variable will be helpful for identifying the business cycle. We examine three specific set-ups, one using an auxiliary variable designed to resemble unemployment, another with a variable resembling inflation, and a third that amounts to using two readings on output, meant to resemble the use of gross domestic income (GDI) as an auxiliary variable.

We find that there is considerable variation in the ability of these auxiliary variables to distinguish business cycles. In particular, for some auxiliary variables, the econometrician can obtain spurious estimates of the correlation between trend and cycle, similar to those obtained by MNZ. In particular, if both variables are nonstationary and resemble GDP—for example, if the bivariate model included both GDP and GDI—spurious results similar to those of univariate studies are a risk. If the other variable in the bivariate analysis is stationary—for example, resembling inflation—we show that spurious results can obtain as the auxiliary process becomes more persistent. Thus, simply adding an auxiliary variable appears not to be sufficient to allow the proper identification of trend and cycle, even if it is stationary.

We find, however, that if the auxiliary variable in the bivariate analysis resembles the unemployment rate, the estimation results can be trusted, even if that variable contains a unit root. Based on our experiments, it appears that the key reason the unemployment rate is well-suited to help distinguish the trend and the cycle is that the the variance of the unit root in the unemployment rate is relatively small compared to variance of its cyclical component.

Motivated by our Monte Carlo results, we use GDP and unemployment rate data to estimate a bivariate UC model. As in other studies using the unemployment rate, we find a conventional cyclical component of GDP, similar to that published by the CBO. The estimated cycle has a pronounced hump-shaped pattern and complex roots, with a period of 7.8 years. We also find that there is no evidence to reject the assumption of orthogonality between trend and cycle components of GDP. That result is suggested by standard statistical tests, and we find in our Monte Carlo work that the size and power of these tests are correct.

The paper is structured as follows: Section 2 presents a review of the literature on trend-cycle decompositions with UC models. In Section 3, we present the characteristics of the bivariate UC models we will examine. Section 4 presents the results of our Monte Carlo experiments. In Section 5, we estimate a bivariate UC model including GDP and unemployment data and test for significance of the correlation between trend and cycle components. Section 6 concludes.

2 Contacts with the Literature

Clark (1987) was among the first to use the Kalman filter to decompose GDP into independent nonstationary trend and stationary cycle components. Clark's estimates implied that much of the quarterly variability in U.S. economic activity can be attributed to a stationary cyclical component. By contrast, Nelson and Plosser (1982), found that most of the variation in U.S. economic activity can be attributed to a nonstationary trend component. A central assumption of Clark's estimation was the orthogonality between trend and cycle components; by contrast, the method of Nelson and Plosser (1982) places no restrictions on the correlation between trend and cycle.

In a subsequent paper, Clark (1989) proposed considering GDP and the unemployment rate in a bivariate UC model to decompose GDP into trend and cycle components, and allowing a nonzero correlation between trend and cycle innovations. In this case, the trend and cycle disturbances are disentangled by assuming that the cyclical component of output can be estimated from unemployment data through an Okun's law relationship. The estimation of the model with U.S. data provides some evidence in favor of the hypothesis that innovations in the trend and cyclical components are independent: the 90 percent confidence interval for the correlation is [-0.4,0.3].

Kuttner (1994) pursued an alternative bivariate approach, adding inflation as an observable and relating inflation and the cycle through a Phillips curve relationship. Kuttner found a business cycle that was similar to Clark (1989). However, Kuttner did not allow for correlation between trend and cycle. Following Kuttner, Roberts (2001) also included a Phillips-curve relationship and further decomposed output into hours and productivity components. Hours and output per hour are each divided into trend and gap components, and the gap affects inflation through a Phillips curve. Both Roberts and Kuttner found that estimates of the trend-cycle decomposition were not much affected by the addition of inflation. In addition, Roberts found that the correlations between trend shocks and the cycle were not statistically significant at conventional levels.

MNZ carefully explored identification in the univariate UC model. They showed that the unrestricted ARIMA(2,1,2) representation of the UC model implies second moments that can be matched uniquely to the second moments of the UC model. The estimation of the cycle through both the Beveridge-Nelson decomposition of the ARIMA(2,1,2) model and a univariate UC model allowing for correlation between trend and cycle yield a decomposition where the cycle is mostly noise and most of the variability in GDP occurs through its trend component, similar to Nelson and Plosser (1982).

Basistha and Nelson (2007) estimate a bivariate model with inflation and GDP as observable variables and the output gap influences inflation in a way resembling the New Keynesian Phillips curve, allowing for a dense variance-covariance matrix of the shocks. They find that the GDP trend and cycle innovations are negatively correlated, as obtained by MNZ, albeit with a smaller correlation. Their estimated cycle is nonetheless conventional. The authors extend the model to include the unemployment rate as an additional observable and incorporate an Okun's law relationship. Results are similar to those when only inflation is included as an additional observable in the UC model. As noted in the introduction, Basistha (2007) performed a set of Monte Carlo simulations that showed that while a univariate specification is not able to identify the correlation coefficient between trend and cycle innovations, a bivariate setup similar to Basistha and Nelson (2007) yields an estimated correlation coefficient that is close on average to the true correlation.

Perron and Wada (2009) estimated a univariate UC model to decompose GDP into trend and cycle incorporating a break in the drift term of the trend in the year 1973. When the break is incorporated, the model that assumes a zero correlation between trend and cycle innovations yields a variance of the trend innovation equal to zero, whereas the model that allows for correlation between trend and cycle innovations yields a correlation equal to +1. The authors state that this last result is expected if the true value of the variance of the trend innovation is zero, since the covariance parameter between the shocks is not identified. In such cases, as discussed by Watson (1986), the trend-cycle decomposition is well identified but the fact that the Kalman filter minimizes the mean-squared error of the estimates of the state vector implies that a perfect correlation will result, since it allows for a perfect fit to the state vector. That is, when the break is allowed, the trend-cycle decompositions (including the Beveridge-Nelson decomposition) yield a cycle that is persistent and explains most of the variations in GDP, as opposed to MNZ's findings.

Building on the results of Perron and Wada (2009), Wada (2012) noted that a univariate UC model of the type proposed by MNZ, but for a stationary process, is likely to yield a correlation coefficient between trend and cycle innovations of -1 when positive definiteness of the variance-covariance matrix of the trend and cycle innovations is imposed, even when the innovations are uncorrelated. This occurs because, under no correlation between the trend and cycle shocks, the likelihood function of the UC model has its global maximum at somewhere other than the true parameter values.

3 UC Models for Trend-cycle Decompositions

We first present the basic structure of the unobserved-components trend-cycle decomposition and then examine three specific bivariate extensions. We explain the mechanism through which the models estimate the trend and cycle components, and the assumptions in terms of parameter configurations.

3.1 The basic UC model

$$y_t = \tau_{yt} + c_t \tag{1}$$
$$\tau_{yt} = \mu_y + \tau_{y,t-1} + \eta_{yt} \tag{2}$$
$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \varepsilon_t, \tag{3}$$

In this setup, $\{y_t\}$ is the log of GDP, $\{\tau_{y,t}\}$ is its unobserved trend, assumed to be a random walk with mean growth rate μ_y, and $\{c_t\}$ is the unobserved stationary cycle. The roots of $1 - \phi_1 z - \phi_2 z^2 = 0$ are outside the unit circle, and $\{\eta_{yt}\}$ and $\{\varepsilon_t\}$ are potentially correlated disturbances.

In the univariate case, the variance-covariance matrix of the disturbances is:

$$\begin{bmatrix} \varepsilon_t \\ \eta_{yt} \end{bmatrix} \sim \text{iid } \mathbb{N}\left(\mathbf{0}_{2\times 1}, \begin{bmatrix} \sigma_\varepsilon^2 & \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon \\ \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \sigma_{\eta_y}^2 \end{bmatrix}\right). \tag{4}$$

In early work with this model, Clark (1987) and Watson (1986) assume that $\rho_{\eta_y\varepsilon} = 0$—that is, that trend and cycle were orthogonal. By contrast, MNZ allowed $\rho_{\eta_y\varepsilon}$ to be non-zero, and found that it was close to -1. As noted above, however, Wada (2012) and Basistha (2007) have shown that this estimate was spurious, raising questions about the ability of a univariate approach to estimate correctly $\rho_{\eta_y\varepsilon}$. We now consider several bivariate models that may allow trend and cycle to be decomposed even in the presence of nonzero correlation between trend and cycle disturbances.

3.2 Bivariate UC Model: GDP and the unemployment rate

Clark (1989) proposed extending the univariate UC model of GDP to include the unemployment rate in the following fashion:

$$u_t = \tau_{ut} + \theta_1 c_t + \theta_2 c_{t-1} \tag{5}$$
$$\tau_{ut} = \tau_{u,t-1} + \eta_{ut}, \tag{6}$$

and variance-covariance matrix:

$$\begin{bmatrix} \varepsilon_t \\ \eta_{yt} \\ \eta_{ut} \end{bmatrix} \sim \text{iid } \mathbb{N}\left(\mathbf{0}_{3\times 1}, \begin{bmatrix} \sigma_\varepsilon^2 & \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \rho_{\eta_u \varepsilon}\sigma_{\eta_u}\sigma_\varepsilon \\ \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \sigma_{\eta_y}^2 & \rho_{\eta_y\eta_u}\sigma_{\eta_y}\sigma_{\eta_u} \\ \rho_{\eta_u \varepsilon}\sigma_{\eta_u}\sigma_\varepsilon & \rho_{\eta_y\eta_u}\sigma_{\eta_y}\sigma_{\eta_u} & \sigma_{\eta_u}^2 \end{bmatrix}\right). \tag{7}$$

Here, $\{u_t\}$ is the unemployment rate and $\{\tau_{ut}\}$ is its unobserved trend, assumed to be a random walk with zero drift.

In Equation (5), the unemployment rate is decomposed into a trend and a cyclical component, with the cycle of output allowed to affect the unemployment rate both contemporaneously and with a lag, reflecting the well-known characterization of the unemployment rate as a lagging indicator (see Stock and Watson, 1998). The model allows the correlations between the cycle and trend innovations of GDP, $\rho_{\eta_y \varepsilon}$, and the unemployment rate, $\rho_{\eta_u \varepsilon}$, to be nonzero, as well as the correlation between the two trend shocks, $\rho_{\eta_y \eta_u}$.

There are a number of ways to interpret this model. One is that the system of Equations (1)-(3) and (5)-(7) forms a factor model, with $\{c_t\}$ the common factor, normalized so that its effect on $\{y_t\}$ is contemporaneous with a coefficient of one. Another interpretation of Equation (5) is Okun's Law, with the unemployment gap related to the output gap with a lag.

3.3 Bivariate UC Model: GDP and the inflation rate

Another alternative introduces inflation:

$$\pi_t = \tau_{\pi t} + \delta c_t \tag{8}$$

$$\tau_{\pi t} = \beta_0 + \beta_1 \pi_{t-1} + \eta_{\pi t}, \tag{9}$$

with $\beta_1 \in [0,1]$, $\beta_0 = 0$ if $\beta_1 = 1$, and

$$\begin{bmatrix} \varepsilon_t \\ \eta_{yt} \\ \eta_{\pi t} \end{bmatrix} \sim \text{iid } \mathbb{N} \left(\mathbf{0}_{3\times 1}, \begin{bmatrix} \sigma_\varepsilon^2 & \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \rho_{\eta_\pi \varepsilon}\sigma_{\eta_\pi}\sigma_\varepsilon \\ \rho_{\eta_y \varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \sigma_{\eta_y}^2 & \rho_{\eta_y \eta_\pi}\sigma_{\eta_y}\sigma_{\eta_\pi} \\ \rho_{\eta_\pi \varepsilon}\sigma_{\eta_\pi}\sigma_\varepsilon & \rho_{\eta_y \eta_\pi}\sigma_{\eta_y}\sigma_{\eta_\pi} & \sigma_{\eta_\pi}^2 \end{bmatrix} \right). \tag{10}$$

Here, $\{\pi_t\}$ is the inflation rate, $\{\tau_{\pi t}\}$ is its unobserved trend, which may be a stationary process when $\beta_1 < 1$, or a random walk without drift when $\beta_1 = 1$ and $\beta_0 = 0$.

This model incorporates a Phillips curve relationship (Equation (8)), where the cyclical component of output helps predict inflation. It also incorporates the notion of inflation expectations in the process $\{\tau_{\pi t}\}$, where expectations are specified as adaptive, hence, a function of lagged observed inflation.[1] Kuttner (1994) also introduced inflation, with correlations between innovations set to zero and inflation modeled as an invertible MA(3) process. Basistha (2007, 2009) assumed that inflation was stationary and allowed for correlated trend-cycle innovations, while Roberts (2001) assumed that inflation followed a random walk without drift and tested for correlations among the innovations.

[1] Another paper that incorporates a Phillips curve relationship in the estimation of the output gap is Basistha and Nelson (2007). In place of lagged inflation, these authors included a survey measure of inflation expectations in Equation (9) and so did not need to confront the question of the stationarity of inflation.

3.4 Bivariate UC Model: GDP and GDI

A third way to obtain trend-cycle decompositions involves GDP and GDI, as follows:

$$x_t = \tau_{xt} + c_t \tag{11}$$

$$\tau_{xt} = \mu + \tau_{x,t-1} + \eta_{xt}, \tag{12}$$

and

$$\begin{bmatrix} \varepsilon_t \\ \eta_{yt} \\ \eta_{xt} \end{bmatrix} \sim \text{iid } \mathbb{N}\left(\mathbf{0}_{3\times 1}, \begin{bmatrix} \sigma_\varepsilon^2 & \rho_{\eta_y\varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \rho_{\eta_x\varepsilon}\sigma_{\eta_x}\sigma_\varepsilon \\ \rho_{\eta_y\varepsilon}\sigma_{\eta_y}\sigma_\varepsilon & \sigma_{\eta_y}^2 & \rho_{\eta_y\eta_x}\sigma_{\eta_y}\sigma_{\eta_x} \\ \rho_{\eta_x\varepsilon}\sigma_{\eta_x}\sigma_\varepsilon & \rho_{\eta_y\eta_x}\sigma_{\eta_y}\sigma_{\eta_x} & \sigma_{\eta_x}^2 \end{bmatrix}\right). \tag{13}$$

In this specification, $\{x_t\}$ is the log of GDI, $\{\tau_{xt}\}$ is its unobserved trend, assumed to be a random walk with mean growth rate μ, which is assumed to be the same for GDP, and $\{c_t\}$ is the common unobserved stationary cycle. Fixler and Nalewaik (2007) and Nalewaik (2010) have been leading proponents of GDI as a measure of aggregate economic activity. Fleischman and Roberts (2011) include both GDP and GDI along with other indicators in their trend-cycle model. The correlation coefficient $\rho_{\eta_y\eta_x}$ captures the co-movements between the trends, and the coefficients $\rho_{\eta_y\varepsilon}$ and $\rho_{\eta_x\varepsilon}$ allow for separate trend-cycle correlations for GDP and GDI.

4 Monte Carlo Exercises

In our Monte Carlo experiments, we simulate each of the three bivariate models introduced in the previous section 3,000 times and estimate them by maximum likelihood using the Kalman filter. We focus mainly on two characteristics of the estimation: The properties of the estimated coefficient $\hat{\rho}_{\eta_y,\varepsilon}$, and the properties of the cycle, namely the percent of the variation in GDP growth that is explained by the variance of the cycle, as implied by the estimated autoregressive coefficients $\hat{\phi}_1$ and $\hat{\phi}_2$ and the estimated variances $\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_\eta^2$.

4.1 Monte Carlo Exercise: GDP and the unemployment rate

To perform the Monte Carlo experiment, we simulate the model in Equations (1)-(3) and (5)-(7) using the parametrization shown in Table 1. The central question is whether standard maximum-likelihood techniques can recover the correct values of the correlations between trend and cycle—in this case, $\rho_{\eta_y\varepsilon}$ and $\rho_{\eta_u\varepsilon}$. As emphasized by Wada (2012) and Basistha (2007), standard univariate techniques can find large estimates of these parameters even when the true values are zero. In the data generating process, we therefore begin by assuming that $\rho_{\eta_y\varepsilon}$ and $\rho_{\eta_u\varepsilon}$ are zero. The values of ϕ_1 and ϕ_2 imply that the cycle will display a hump-shaped pattern; in this case, there are complex roots implying a duration of the cycle of about 25 periods. Its standard deviation is 2.16 and the variance of the cycle explains about 90.5 percent of the variation of GDP growth. The calibration chosen for the process that represents GDP is close to the results obtained by MNZ in the uncorrelated setup of their paper (labeled UC-0). The calibration for the process that represents the unemployment rate is consistent with Fleischman and Roberts (2011).

Table 1: Parameter Values - Bivariate UC model: GDP and the unemployment rate

	Parameter Value
μ_y	0.8
ϕ_1	1.5
ϕ_2	-0.6
σ_ε	0.6
σ_{η_y}	0.7
$\rho_{\eta_y\varepsilon}$	0
θ_1	-0.40
θ_2	-0.25
σ_{η_u}	0.1
$\rho_{\eta_u\varepsilon}$	0
$\rho_{\eta_y\eta_u}$	0

With the data generated by the bivariate specification in Table 1, we repeatedly estimate the unconstrained UC model using maximum likelihood. Figure 1a shows the distribution of the estimated correlation coefficient $\hat{\rho}_{\eta_y\varepsilon}$ obtained from the simulations for sample sizes $n = 50, 100, 200, 1000$. The distribution of the maximum likelihood estimator of the correlation coefficient between the trend and cycle innovations of GDP has a conventional shape, with the mass of the distribution concentrated around the true coefficient value as the sample size increases. For typical U.S. quarterly sample sizes or around 200, results are reasonably precise.

The shape of the distribution of the estimator of the trend-cycle correlation coefficient under this bivariate setup contrasts sharply with the distribution of the same correlation coefficient when a univariate estimation that includes GDP only, as in MNZ, is used. In Figure 1b, results are reported based on data generated with the same parameters as in Table 1, but omitting the unemployment rate as an observable in the estimation. As can be seen, the maximum likelihood estimation of the univariate UC model implies an estimated trend-cycle correlation coefficient that has a distribution with masses close to -1 and +1. The results in Figures 1a and 1b suggest that estimation of the bivariate UC model that includes the unemployment rate substantially improves the small sample properties of the estimator of the trend-cycle correlation coefficient.

We also simulate the bivariate UC model of this section assuming that there is important correlation between trend and cycle—that is, that the results of MNZ hold. In particular, we simulate the model with the same parametrization of Table 1, except that we also assume that $\rho_{\eta_y\varepsilon} = -0.9$ (MNZ's case) and that $\rho_{\eta_y\varepsilon} = 0.9$. Figure 2 reports the results of the bivariate and univariate UC model estimation. The plot shows that the bivariate model also does a good job at estimating the model when trend and cycle are correlated. The univariate model continues to give probability masses at the tails. Thus, if the trend and cycle were correlated, the bivariate model would be able to detected it.

Figure 1: Frequency Distribution of $\hat{\rho}_{\eta_y \varepsilon}$

(a) Bivariate UC model: GDP and the unemployment rate

(b) Univariate UC model: GDP

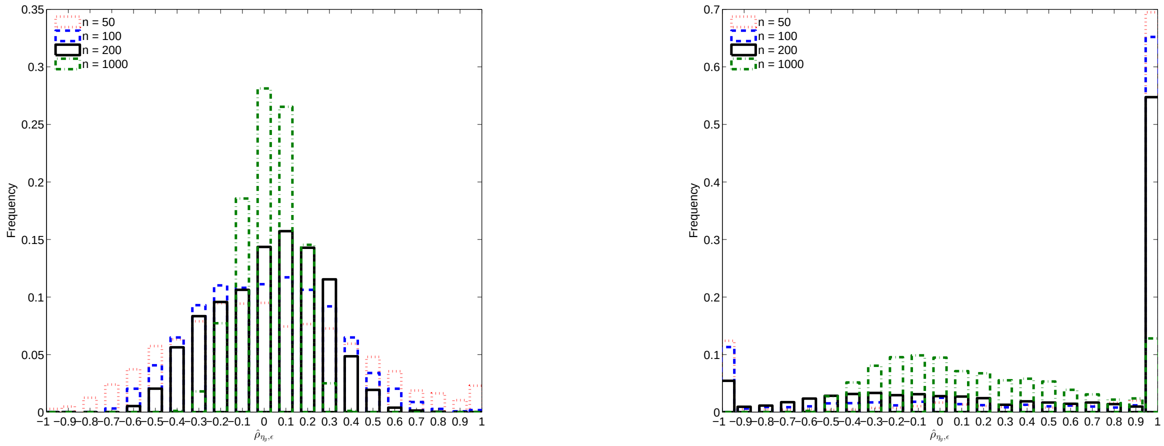

Figure 2: Frequency Distribution of $\hat{\rho}_{\eta_y \varepsilon}$ under Different True Values

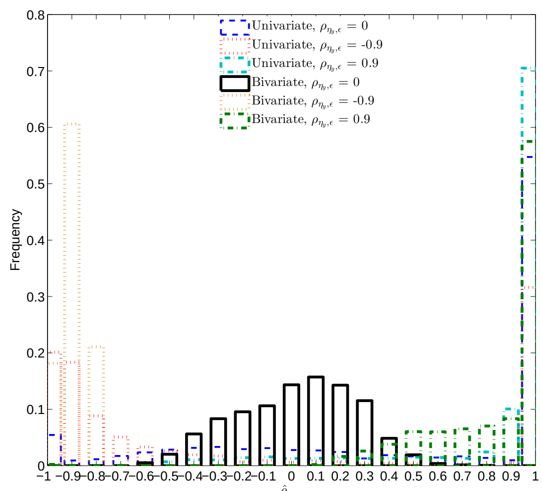

Figure 3: Frequency Distribution of the Estimated Period

(a) Bivariate UC model: GDP and the unemployment rate

(b) Univariate UC model: GDP

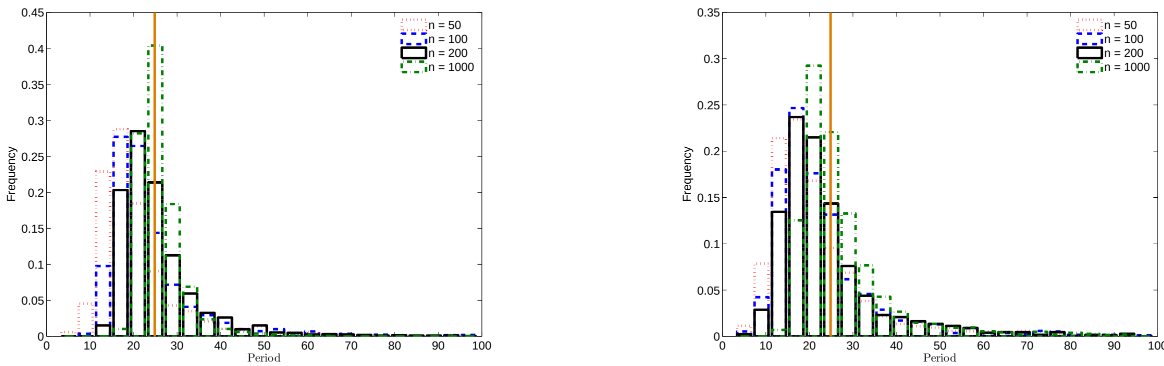

Note: The orange line determines the period implied by the parameters ϕ_1 and ϕ_2, which is 25. The implied ratio $\text{var}(c_t)/\text{var}(y_t)$ is 90.5%.

	Bivariate		Univariate	
Sample Size	Median Period	% $\text{var}(c_t)/\text{var}(y_t)$	Median Period	% $\text{var}(c_t)/\text{var}(y_t)$
$n = 50$	18	82.0%	18	83.0%
$n = 100$	21	87.4%	20	80.0%
$n = 200$	23	89.2%	21	82.0%
$n = 1,000$	25	90.1%	24	88.0%

Another dimension along which the results of the univariate and multivariate UC models' results differ significantly concerns the properties of the business cycle. While the univariate setup in MNZ implies a cycle for the U.S. economy with a period of about 10 quarters when correlation is allowed between trend and cycle innovations, the periodicity of the cycle is considerably longer in bivariate specifications that allow for correlation between trend and cycle. For example, the period of the cycle is estimated to be about 28 quarters in Clark (1989), about 20 quarters in the trivariate setup of Basistha and Nelson (2007), and infinite in their bivariate setup that includes GDP and inflation only. Figure 3 reports the frequency distribution of the periods implied by the simulation results, the median estimated period and the percent of the variation in GDP growth that is explained by the variance of the cycle. The results in Figure 3a show that, as the sample size increases in the bivariate setup, the estimated duration of the cycle approaches the implied true duration of about 25 time periods and an amplitude very similar to that implied by the true parameters.

Figure 3b reports the results from the univariate estimation of the UC model that includes GDP data only. The results show that the univariate estimation tends to deliver a shorter cyclical period. It also underestimates the fraction of the variation of GDP growth explained by the cycle. In particular, with a sample size of 200, the bivariate model is very close to the correct cyclical contribution while the univariate model understates the cyclical contribution by about 10 percent. Thus, the addition of the unemployment rate also helps reduce the bias in the estimation of the period and the amplitude of the cycle. However, benefits are

Table 2: Parameter Values - Bivariate UC model: GDP and the inflation rate

	Parameter Value
μ_y	0.8
ϕ_1	1.5
ϕ_2	-0.6
σ_ε	0.6
σ_{η_y}	0.7
$\rho_{\eta_y \varepsilon}$	0
δ	0.4
β_0	1.5
β_1	{0.5,0.75,0.95,1}
σ_{η_π}	1
$\rho_{\eta_\pi \varepsilon}$	0
$\rho_{\eta_y \eta_\pi}$	0

not as great as for the estimates of the correlation between trend and cycle.

4.2 Monte Carlo Exercise: GDP and the inflation rate

We next consider the implications for trend-cycle decomposition of adding a variable that resembles inflation. Here, the model is that specified in Equations (1)-(3) and Equations (8)-(10); the parameters are shown in Table 2. Parameters that are common with the previous model—notably, those associated with the trend and cycle in GDP—are the same as before. We also start by assuming that the correlations between trends and cycle in the data generating process are zero. The calibration for the process that represents inflation is consistent with the Monte Carlo exercise in Basistha (2007). The table shows a variety of inflation persistence coefficients, β_1, in the set $\{0.5, 0.75, 0.95, 0.99\}$. We simulate inflation with each of these values keeping the unconditional variance of inflation constant when inflation is assumed to be stationary.

Figure 4 reports the frequency distribution of $\hat{\rho}_{\eta_y \varepsilon}$. The results indicate that the distribution of the estimated trend-cycle correlation coefficient has masses close to -1 and +1 even for low values of the persistence of inflation, although the distortion of the distribution is less severe when inflation is less persistent. These results suggest that if inflation is even mildly persistent, it may not be the best variable to disentangle the correlation between trend and cycle components of GDP in a bivariate setup.[2]

Figure 5 reports the frequency distribution of the periods implied from the results of the simulations, the median estimated period, and the percentage of the variation in GDP

[2] Basistha (2007) also considered a bivariate set-up that included inflation. Basistha's results appear to indicate considerably greater success using inflation than in Figure 4. However, Basistha reported the mean of the distribution of the estimated correlation coefficient only and not the full sampling distribution, obscuring the large masses at the extremes.

Figure 4: Frequency Distribution of $\hat{\rho}_{\eta_y \varepsilon}$ - Bivariate UC model: GDP and the inflation rate

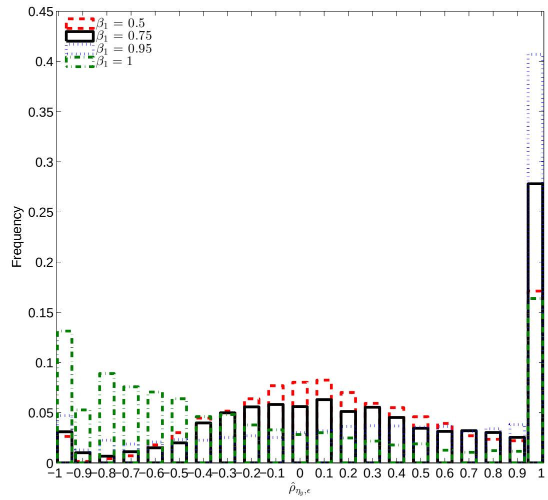

growth that is explained by the variance of the cycle. The results show that the duration of the cycle declines as the degree of persistence of the inflation process increases. The percent of the variation in GDP growth that is explained by the variance of the cycle is significantly below the theoretical percent of 90.5 percent, decreasing to about 38 percent when inflation is stationary but highly persistent. If inflation has a unit root, the percent of the variance of GDP explained by variations of the cycle is about 73 percent, still well below the theoretical value.

4.3 Monte Carlo Exercise: GDP and GDI

In this Monte Carlo experiment, we simulate the model in Equations (1)-(3) and (11)-(13) using the parametrization shown in Table 3. Parameters that are common with the previous models—notably, those associated with the trend and cycle in GDP—are the same as before. The correlations between trends and cycle, as captured by $\rho_{\eta_y \varepsilon}$ and $\rho_{\eta_x \varepsilon}$, are set to zero. The calibration for $\rho_{\eta_y \eta_x}$ is consistent with Fleischman and Roberts (2011).

Figure 6 reports the frequency distribution of $\hat{\rho}_{\eta_y \varepsilon}$. The results show that the distribution of the estimated trend-cycle correlation coefficient has masses close to -1 and +1. While these masses are smaller than in the univariate case, they nonetheless indicate that there is considerable risk of obtaining biased results in this set-up. Thus, as with inflation, these results suggest that GDI would not be as successful as the unemployment rate in disentangling the correlation between trend and cycle components of GDP in a bivariate setup.

We also explore how the inclusion of GDI affects the estimation of the period and amplitude of the cycle. Figure 7 reports the frequency distribution of the periods implied from the results of the simulations, the median estimated period, and the percent of the variance of GDP growth explained by the variance of the cycle. The results show that the median

Figure 5: Frequency Distribution of the Estimated Period - Bivariate UC model: GDP and the inflation rate

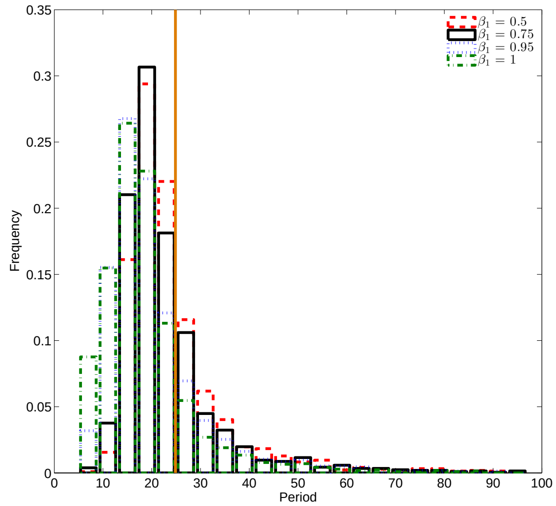

Note: The orange line determines the period implied by the parameters ϕ_1 and ϕ_2, which is 25. The implied ratio $\text{var}(c_t)/\text{var}(y_t)$ is 90.5%.

Persistence	Median Period	% $\text{var}(c_t)/\text{var}(y_t)$
$\beta_1 = 0.5$	21	69.7%
$\beta_1 = 0.75$	20	63.3%
$\beta_1 = 0.95$	18	38.3%
$\beta_1 = 1$	17	73.2%

Table 3: Parameter Values - Bivariate UC model: GDP and GDI

	Parameter Value
μ	0.8
ϕ_1	1.5
ϕ_2	-0.6
σ_ε	0.6
σ_{η_y}	0.7
$\rho_{\eta_y \varepsilon}$	0
σ_{η_x}	0.7
$\rho_{\eta_x \varepsilon}$	0
$\rho_{\eta_y \eta_x}$	0.4

Figure 6: Frequency Distribution of $\hat{\rho}_{\eta_y \varepsilon}$ - Bivariate UC model: GDP and GDI

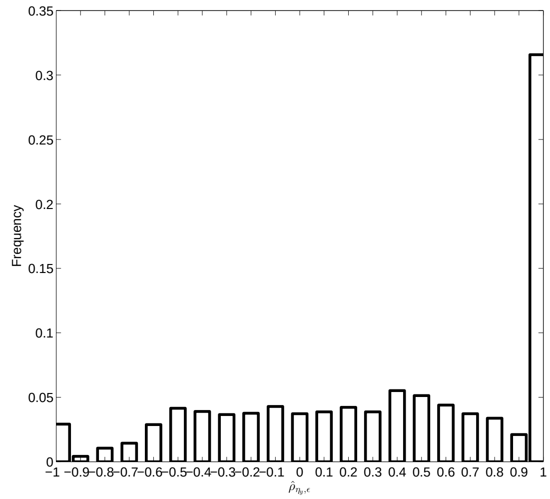

duration of the cycle is below the true period and that the cycle does not account for as much variation of GDP growth as in the data generating process.

4.4 Understanding The Effect of Incorporating the Unemployment Rate

In this section, we explore why using a variable that resembles the unemployment rate in a bivariate model leads to improved GDP trend-cycle decompositions whereas adding variables that resemble inflation and GDI do not. Our conjecture is that the advantage of the unemployment rate is in the relatively small variance of its nonstationary component relative to its cycle. To test this conjecture, we conduct a Monte Carlo exercise in which the standard deviation of the the unemployment rate trend, σ_{η_u}, varies in the set $\{0.01, 0.1, 0.6, 1.2\}$ and the calibration is otherwise the same as in Table 1. Figure 8 shows the results.

As can be seen, increasing the standard deviation of the trend innovation of the variable that resembles the unemployment rate makes the estimates of $\rho_{\eta_y \varepsilon}$ less accurate, and pushes them in the direction of the univariate estimates. This result helps explain why the correlation coefficient is poorly estimated in the GDP-GDI bivariate UC model, where GDP and GDI have the same standard deviation of the trend innovation, which is much higher than the standard deviation of the cycle. This also explains why, as inflation becomes more persistent, the correlation coefficient is imprecisely estimated in the GDP-inflation bivariate UC model. Thus, adding a variable that has a high trend variability compared with the variability of the cycle would not help to recover the correlation coefficient $\rho_{\eta_y \varepsilon}$ or to give accurate estimates of the cycle.[3]

[3] We also explored the importance of the lagging-indicator properties of the unemployment rate by conducting a Monte Carlo exercise. We set $\theta_2 = 0$ in the bivariate GDP-unemployment model and boosted

Figure 7: Frequency Distribution of the Estimated Period - Bivariate UC model: GDP and GDI

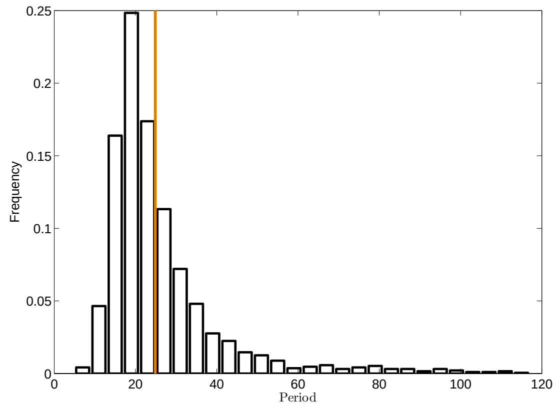

Note: The orange line determines the period implied by the parameters ϕ_1 and ϕ_2, which is 25. The implied ratio $\mathrm{var}(c_t)/\mathrm{var}(y_t)$ is 90.5%.

Median Period	% $\mathrm{var}(c_t)/\mathrm{var}(y_t)$
22	84.8%

Figure 8: Frequency Distribution of $\hat{\rho}_{\eta_y \varepsilon}$ under Different Values of σ_{η_u}

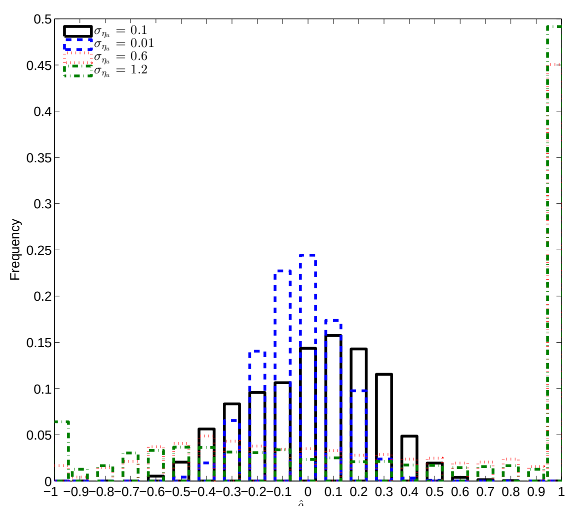

4.5 Size and Power of the Likelihood Ratio Test of Hypotheses with Respect to $\rho_{\eta_y\varepsilon}$

In this subsection, we explore the size and power of hypotheses tests with respect to $\rho_{\eta_y\varepsilon}$. The aim of the section is to show that the likelihood ratio (LR) test has the wrong size and no power in the univariate case, while the size and power improve substantially in the bivariate case that uses the unemployment rate. This finding supports the view that likelihood ratio tests conducted to determine the statistical significance of the trend-cycle correlation coefficient in the estimation of a bivariate UC model that includes the unemployment rate yield correct inferences.

We first consider the ability of the likelihood-ratio test to reject various false hypothesized values for $\rho_{\eta_y\varepsilon}$—that is, the power of the test. Two exercises are performed in which we simulate the bivariate model 1,000 times. First, we simulate the model according to the benchmark parameter specification in Table 1, that is, assuming that $\rho_{\eta_y\varepsilon} = 0$, and set the null hypotheses in the univariate and bivariate estimations as $H_0 : \rho_{\eta_y\varepsilon} = \rho^o_{\eta_y\varepsilon}$, where $\rho^o_{\eta_y\varepsilon} \in \{-0.95, -0.9, -0.8, -0.7, \ldots, -0.1, 0.1, \ldots, 0.7, 0.8, 0.9, 0.95\}$. Second, we simulate the model setting $\rho_{\eta_y\varepsilon} = -0.9$ and test the null hypotheses as $H_0 : \rho_{\eta_y\varepsilon} = \rho^o_{\eta_y\varepsilon}$, where $\rho^o_{\eta_y\varepsilon} \in \{-0.99, -0.98, -0.975, -0.95, -0.925, -0.85, -0.8, -0.7, \ldots, -0.1, 0, 0.1, \ldots, 0.7, 0.8, 0.9, 0.95\}$.

Figure 9 reports the fraction of times the likelihood ratio test rejects the false hypothesized value, for the univariate and bivariate models, in the two exercises performed. When the true value of $\rho_{\eta_y\varepsilon}$ is zero, Figure 9a shows that the test based on univariate estimation has virtually no power to reject hypothesized values of $\rho_{\eta_y\varepsilon}$ greater than -0.5. The power increases as the value of the hypothesized correlation coefficient approaches the left tail but still reaches only 33 percent for the null hypothesis $H_0 : \rho_{\eta_y\varepsilon} = -0.9$. Hence, the test based on the univariate estimation has very low power to reject the false null hypothesis of negative correlation between trend and cycle.[4] The test based on the bivariate estimation, on the other hand, has power increasing to almost 100 percent as the hypothesized value moves away from the true correlation coefficient of zero.

When the true value of $\rho_{\eta_y\varepsilon}$ is -0.9, Figure 9b shows that the univariate estimation lacks the ability to reject almost any false hypothesized value, whereas the bivariate estimation rapidly approaches a rejection probability of one as the hypothesized values move above the true value of the correlation. In particular, the bivariate model would reject with probability one a hypothesized value of zero for the correlation between trend and cycle innovations if the true correlation were -0.9.

We next compute the frequency with which the likelihood ratio test incorrectly rejects a true hypothesized value—that is, the size of the test. We simulate the bivariate model 1,000 times according to the benchmark parameter specification in Table 1, except that we set $\rho_{\eta_y\varepsilon} = \rho^o_{\eta_y\varepsilon}$ for $\rho^o_{\eta_y\varepsilon} \in \{-0.95, -0.9, -0.8, \ldots, -0.1, 0, 0.1, \ldots, 0.8, 0.9, 0.95\}$. We consider

the value of θ_1 to keep the sum of the two the same. We found that the results were very little affected by this change. Thus, we conclude that it is predominatly the low variance of the unit root in the unemployment rate and not the fact that it is a lagging indicator that mostly accounts for its desirability as a cyclical indicator.

[4] A possible explanation for the low power of the univariate test may lie in the finding by Wada (2012) that the likelihood function peaks at a different value than the true correlation.

Figure 9: Power of the Likelihood Ratio Test

(a) $\rho_{\eta_y \varepsilon} = 0$ (b) $\rho_{\eta_y \varepsilon} = -0.9$

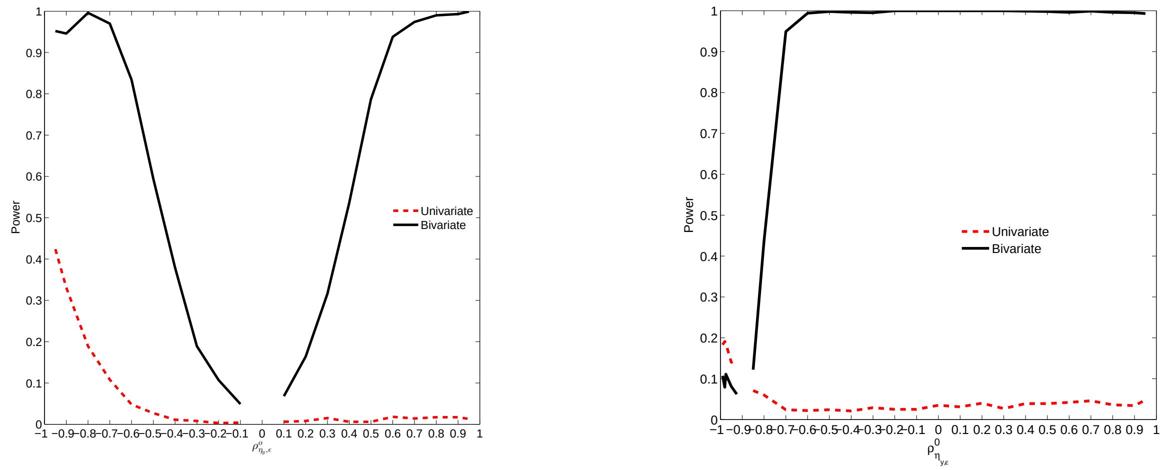

Note: The power corresponds to a significance level $\alpha = 0.05$.

Figure 10: Size of the Likelihood Ratio Test

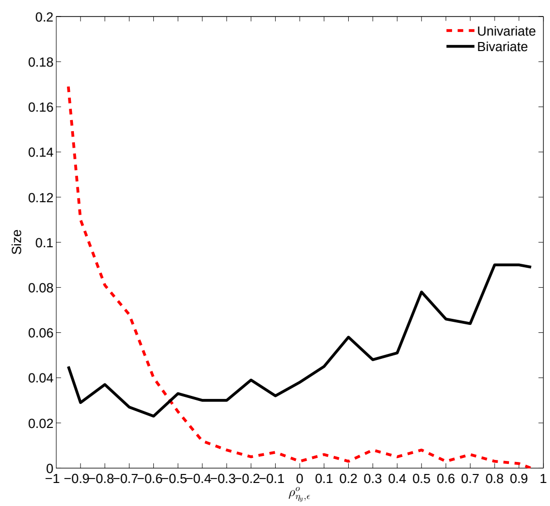

Note: The size corresponds to a significance level $\alpha = 0.05$.

the null hypotheses in the univariate and bivariate estimations as $H_0 : \rho_{\eta_y \varepsilon} = \rho^o_{\eta_y \varepsilon}$. Figure 10 plots the size of the likelihood ratio test for the univariate and bivariate models using a 5 percent significance level. As can be seen, the univariate estimation delivers very low size for hypothesized values of the correlation coefficient above -0.5, meaning that the test will not reject the hypothesized true value often enough. For correlations below -0.5, the size is too large, meaning that the test of the null hypothesis will be rejected too often.[5][6] The bivariate model, on the other hand, keeps the size between approximately 3 percent and 9 percent, close to the chosen value of 5 percent.

5 Estimation Results of the GDP-Unemployment Bivariate Model

We use quarterly GDP and unemployment rate data from the St. Louis FRED databases as observable variables for the period 1947:1-2013:4 to estimate the unrestricted unobserved components (UC-UR) model in Equations (1)-(7) to extract the trend and the cycle of GDP. Results appear in Table 4.

Table 4: GDP and Unemployment Bivariate Model Estimates

	Estimate	Standard Error	Z-statistic
μ_y	0.83	0.04	21.30
ϕ_1	1.58	0.07	21.98
ϕ_2	-0.66	0.06	-10.47
σ_ε	0.67	0.11	5.89
σ_{η_y}	0.84	0.07	12.60
$\rho_{\eta_y \varepsilon}$	-0.52	0.14	-3.62
θ_1	-0.35	0.07	-5.40
θ_2	-0.21	0.05	-4.05
σ_{η_u}	0.24	0.04	5.42
$\rho_{\eta_u \varepsilon}$	1.00	0.02	66.20
$\rho_{\eta_y \eta_u}$	-0.67	0.06	-11.11

LogL = -320.03, BIC = 701.35

The correlation between the innovations of the cycle and trend components of GDP is estimated to be -0.52, a lower value than the correlation obtained by MNZ, although still

[5]Based on their simulations, MNZ argue that the size of the likelihood ratio test of the hypothesis $H_0 : \rho_{\eta_y \varepsilon} = 0$ is approximately correct. This interpretation is broadly consistent with our findings because we find that at $H_0 : \rho_{\eta_y \varepsilon} = 0$, the null hypothesis will not be rejected often enough, making MNZ's rejection of the hypothesis all the more convincing.

[6]Again, the reason for the incorrect size in the univariate case may lie in the fact that the likelihood function peaks at a value different from the true correlation coefficient.

Table 5: GDP and Unemployment Bivariate Model Estimates Imposing $\rho_{\eta_y \varepsilon} = \rho_{\eta_u \varepsilon} = 0$

	Estimate	Standard Error	Z-statistic
μ_y	0.82	0.04	21.33
ϕ_1	1.60	0.05	13.40
ϕ_2	-0.66	0.06	-10.79
σ_ε	0.53	0.06	8.94
σ_{η_y}	0.66	0.05	13.40
θ_1	-0.26	0.06	-4.02
θ_2	-0.29	0.08	-3.84
σ_{η_u}	0.18	0.01	13.12
$\rho_{\eta_y \eta_u}$	-0.35	0.15	-2.32

LogL = -320.54, BIC = 691.23

negative and, according to the z-ratio, statistically significant. The coefficients that relate the unemployment rate to the cycle have the expected signs, implying that the unemployment rate is countercyclical and a lagging variable. As expected, the correlation between the GDP and the unemployment trend innovations is negative.[7] However, the correlation between the innovations of the unemployment trend and of the cycle is estimated to be 1, indicating a corner solution from the optimizer.

Given the counterintuitive and extreme estimate of the correlation between the cycle and trend innovations of the unemployment rate, we conduct a likelihood ratio test of the restriction $H_0 : \rho_{\eta_u \varepsilon} = 0$. The log-likelihood value of the restricted model (shown in Table A1 of Appendix A) is -320.52, and the Schwarz criterion is 696.76. The LR test does not reject the null hypothesis, while the Schwarz criterion favors the restricted model. With the restriction $\rho_{\eta_u \varepsilon} = 0$ imposed, the estimated correlation between the innovations of the trend and cycle of GDP is very small, and the z-statistic does not reject the null hypothesis $\rho_{\eta_y \varepsilon} = 0$. We therefore estimate the proposed bivariate model imposing the constraints $\rho_{\eta_y \varepsilon} = \rho_{\eta_u \varepsilon} = 0$. Table 5 shows results with these restrictions imposed. Again, the log-likelihood value is again not very different from that in Table 4. The joint restriction is not rejected at the 5 percent significance level, and the BIC very strongly favors the restricted model. We explored the size of the likelihood ratio test of the joint hypothesis $H_0 : \rho_{\eta_y \varepsilon} = \rho_{\eta_u \varepsilon} = 0$ for this model via Monte Carlo experiments. The size under a significance level of 5 percent is 3.6 percent. We also explored the power of the likelihood ratio test at the 5 percent level of significance for the joint hypothesis $H_0 : \rho_{\eta_y \varepsilon} = -0.9, \rho_{\eta_u \varepsilon} = 1$ and obtained a 99 percent rejection probability. As a consequence, the estimates in Table 5 are our preferred results.

The estimates of the autoregressive coefficients of the GDP cycle imply a strong hump-shaped pattern of the responses to a business-cycle shock, similar to the results of Fleischman

[7]One reason we would expect a negative correlation between the GDP and unemployment trend innovations is that an increase in the unemployment trend likely corresponds to a reduction in long-term labor supply, which should, in turn, reduce trend output.

Figure 11: Bivariate UC-UR Model - Smoothed GDP Cycle Imposing $\rho_{\eta_y \varepsilon} = \rho_{\eta_u \varepsilon} = 0$

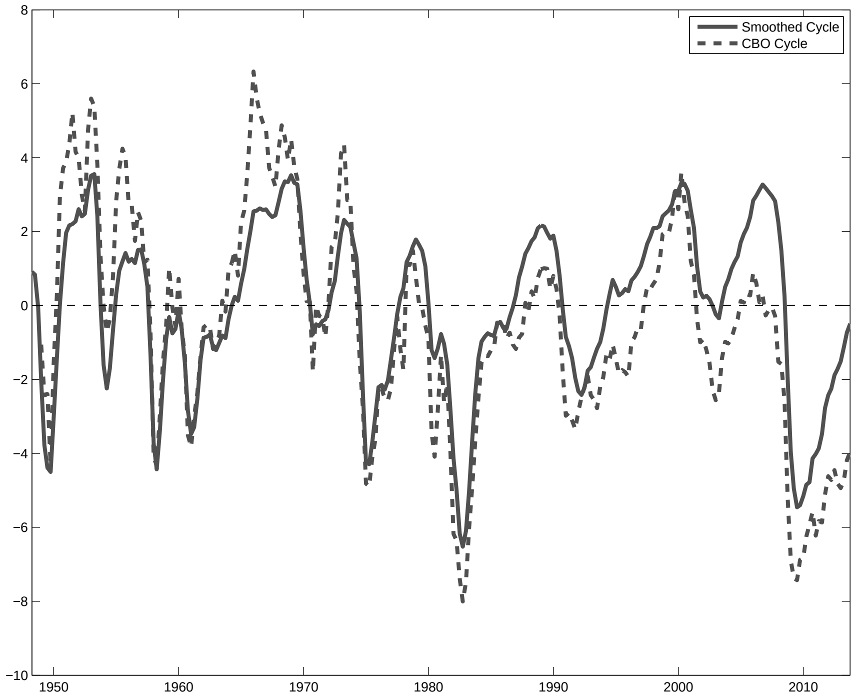

and Roberts (2011) and of MNZ in the estimation of the constrained univariate unobserved components model (labeled UC-0 in their paper). The estimated parameters imply complex roots, with a period of 7.8 years. The variance decomposition indicates that about 94 percent of the variations in GDP growth are explained by variations in the cycle. The estimates of the coefficients that relate the cycle to the unemployment rate, θ_1 and θ_2, suggest a conventional Okun's law relationship, with the unemployment rate reacting to cyclical shocks with a lag relative to GDP, and a total effect after two quarters of -0.55, not far from conventional estimates of -0.5 (see Abel et al., 2013). Again, trend innovations of GDP and unemployment have the expected negative correlation.

The smoothed cycle obtained from this estimation appears in Figure 11 along with the CBO-implied cycle. The two cycles behave similarly, although the cycle obtained from our model shows somewhat less variability than the CBO cycle. At the end of the sample (2013:Q4), our bivariate model predicts that the output gap is less than 1 percent, while CBO estimates that output is still about 4 percent below its trend. Mechanically, this difference occurs because our model predicts smaller increases in the trend component of GDP in the latter years of the sample. This slowdown in trend growth can be seen in Figure 12, which shows the observed GDP series and its corresponding estimated trend. The GDP trend rises only 1 percent per year from 2010 to 2013, compared with 2.2 percent per year from 2007 to 2009 (?) and 1.9 percent per year in the three years before that.

Figure 12 also shows the unemployment rate and its trend. The unemployment rate trend moves up fairly steadily from the beginning of the sample to the mid-1960s and then again through the 1970s, reaching as high as 7 percent in the early 1980s. The trend then moves down over the rest of the 1980s and through the 1990s, reaching as low as 5-1/2 percent.

Figure 12: Observed GDP and Unemployment Rate and Corresponding Trends

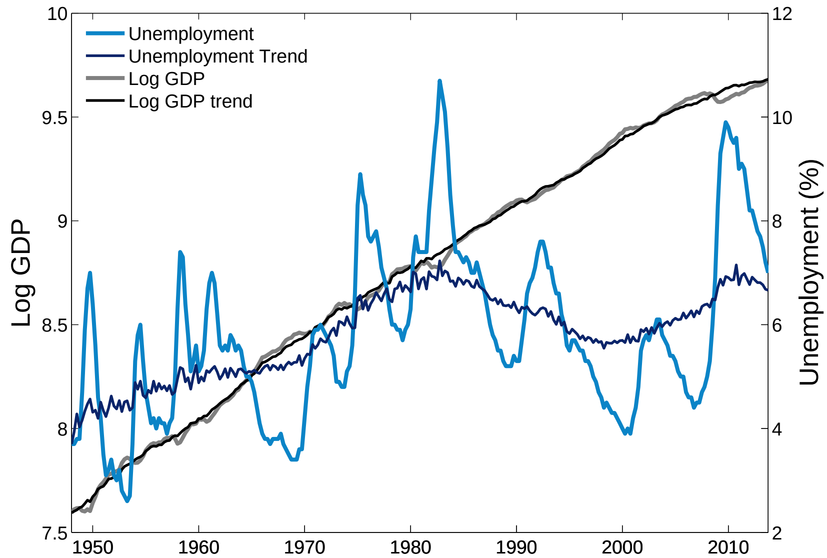

There is an upward sloping trend from the early 2000s and a faster increase that starts at the beginning of the financial crisis and lasts for about two quarters. After that, the unemployment rate trend edges down a bit, to around 6-1/2 percent at the end of 2013.

6 Conclusions

In this paper, we investigated the performance of different bivariate unobserved components models to estimate the trend and cycle of GDP. We found that the best variable to accompany GDP in the bivariate specification is the unemployment rate, which is superior in performance to two alternatives, namely inflation and gross domestic income. Our results suggest that the main reason the unemployment rate is especially helpful is that its unit root component (trend) has a relatively small variance relative to the cycle component. We estimated the cycle using GDP and unemployment data and found that there is not evidence against the assumption of orthogonal trend-cycle innovations of GDP and the unemployment rate. Moreover, our Monte Carlo experiments suggested that the results of our statistical tests could be trusted.

Appendix

A Additional Results

Table A1: Bivariate UC-UR Model Estimates $\rho_{\eta_u \varepsilon} = 0$

	Estimate	Standard Error	Z-statistic
μ_y	0.82	0.04	21.03
ϕ_1	1.59	0.06	25.23
ϕ_2	-0.66	0.06	-10.63
σ_ε	0.54	0.11	4.90
σ_{η_y}	0.67	0.09	7.50
$\rho_{\eta_y \varepsilon}$	-0.06	0.36	-0.17
θ_1	-0.23	0.15	-1.53
θ_2	-0.30	0.09	-3.23
σ_{η_u}	0.18	0.02	10.68
$\rho_{\eta_y \eta_u}$	-0.37	0.17	-2.19

LogL = -320.52, BIC = 696.76

References

Abel, Andrew B., Ben Bernanke, and Dean Croushore (2013) *Macroeconomics (8th Edition)* (Prentice Hall)

Basistha, Arabinda (2007) 'Trend-cycle correlation, drift break and the estimation of trend and cycle in canadian gdp (corrlation tendance-cycle, discontinuit, et estimation de la tendance et du cycle dans le pib canadien).' *The Canadian Journal of Economics / Revue canadienne d'Economique* 40(2), pp. 584–606

_ (2009) 'Hours per capita and productivity: evidence from correlated unobserved components models.' *Journal of Applied Econometrics* 24(1), 187–206

Basistha, Arabinda, and Charles R. Nelson (2007) 'New measures of the output gap based on the forward-looking new Keynesian Phillips curve.' *Journal of Monetary Economics, Elsevier* 54(2), 498–511

Clark, Peter K. (1987) 'The cyclical component of u. s. economic activity.' *The Quarterly Journal of Economics* 102(4), pp. 797–814

Clark, Peter K. (1989) 'Trend reversion in real output and unemployment.' *Journal of Econometrics* 40(1), 15 – 32

Fixler, Dennis J., and Jeremy J. Nalewaik (2007) 'News, noise, and estimates of the "true" unobserved state of the economy.' Finance and Economics Discussion Series 2007-34, Board of Governors of the Federal Reserve System (U.S.)

Fleischman, Charles, and John M. Roberts (2011) 'From many series, one cycle: improved estimates of the business cycle from a multivariate unobserved components model.' Finance and Economics Discussion Series 2011-46, Board of Governors of the Federal Reserve System (U.S.)

Kuttner, Kenneth N. (1994) 'Estimating potential output as a latent variable.' *Journal of Business & Economic Statistics* 12(3), pp. 361–368

Morley, James C, Charles R Nelson, and Eric Zivot (2003) 'Why are the beveridge-nelson and unobserved-components decompositions of gdp so different?' *Review of Economics and Statistics* 85(2), 235–243

Nalewaik, Jeremy J. (2010) 'The income- and expenditure-side estimates of u.s. output growth.' *Brookings Papers on Economic Activity* 41(1 (Spring)), 71–127

Nelson, Charles R, and Charles R Plosser (1982) 'Trends and random walks in macroeconmic time series: some evidence and implications.' *Journal of monetary economics* 10(2), 139–162

Perron, Pierre, and Tatsuma Wada (2009) 'Let's take a break: Trends and cycles in us real gdp.' *Journal of Monetary Economics* 56(6), 749–765

Roberts, John M. (2001) 'Estimates of the productivity trend using time-varying parameter techniques.' *The B.E. Journal of Macroeconomics* 1(1), 1–32

Stock, James H., and Mark W. Watson (1998) 'Business cycle fluctuations in u.s. macroeconomic time series.' NBER Working Papers 6528, National Bureau of Economic Research, Inc

Wada, Tatsuma (2012) 'On the correlations of trendcycle errors.' *Economics Letters* 116(3), 396 – 400

Watson, Mark W. (1986) 'Univariate detrending methods with stochastic trends.' *Journal of Monetary Economics* 18(1), 49 – 75