A Steady-State Approach to Trend/Cycle Decomposition*

James Morley
Washington University in St. Louis

Jeremy Piger
Federal Reserve Bank of St. Louis

February 29, 2004

ABSTRACT: In this paper, we present a new approach to trend/cycle decomposition. The trend of an integrated time series is estimated using the conditional expectation of the *steady-state level* of the series. Given a nonlinear forecasting model, this steady-state approach can differ in important ways from the related long-horizon forecast decomposition proposed by Beveridge and Nelson (1981). We use generated data from nonlinear regime-switching processes to demonstrate the advantages of the steady-state approach. We then apply the steady-state approach to estimate the trend and cycle of U.S. real GDP implied by a regime-switching forecasting model. Our findings portray a very different picture of the business cycle than implied by standard linear forecasting models.

Keywords: Steady State, Trend/Cycle Decomposition, Nonlinear, Regime Switching
JEL Classification: C15, C22, E32

---

* **Morley:** Department of Economics, Box 1208, Eliot Hall 205, Washington University, One Brookings Drive, St. Louis, MO 63130-4899, (morley@economics.wustl.edu); **Piger:** Research Department, Federal Reserve Bank of St. Louis, 411 Locust St., St. Louis, MO 63166, (piger@stls.frb.org). We would like to thank Gaetano Antinolfi and Charles Nelson for helpful comments and suggestions. Responsibility for any errors is our own. Morley acknowledges support from the Weidenbaum Center on the Economy, Government, and Public Policy. The views expressed in this paper should not be interpreted as those of the Weidenbaum Center, the Federal Reserve Bank of St. Louis, or the Federal Reserve System.

1. Introduction

Trend/cycle decomposition of integrated economic time series is important for both theoretical and statistical reasons. In this paper, we present a new approach and compare it with some existing methods in the literature. Unlike traditional methods such as the Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997) and the Beveridge-Nelson (BN) decomposition (Beveridge and Nelson, 1981), the approach presented in this paper accurately estimates the permanent and transitory components of integrated time series even when the dynamic structure of the data involves nonlinearities such as regime switching.

Our approach to trend/cycle decomposition is based on the premise that the trend of a time series is equivalent to its implicit steady-state level. In a dynamic system, steady state is the hypothetical equilibrium that would occur following the realization of all currently implied transitory dynamics. In practice, the concept of steady state identifies the trend since the implied transitory dynamics for a time series can be calculated from a forecasting model.

By employing a forecasting model to measure transitory dynamics, our steady-state approach is closely related to the long-horizon forecast decomposition proposed by Beveridge and Nelson (1981) and extended to nonlinear processes by Clarida and Taylor (2003). Indeed, for linear forecasting models, the two methods produce equivalent results. However, that equivalence does not generally hold for nonlinear models with regime-switching parameters. We present two examples where the two methods produce different results and the steady-state approach is preferable. In the first example, we generate data from Hamilton's (1989) Markov-switching model. The model assumes that the mean growth rate of a time series undergoes discrete regime shifts according to a Markov process. For data generated from the model, the steady-state approach is preferable because, unlike the long-horizon forecast, it does not include expected future regimes in the estimate of the current level of trend. In the second example, we generate data from Kim and Nelson's (1999a) plucking model. According to this model, the level of output undergoes negative "plucks" that produce transitory dynamics. Again, the steady-state approach is preferable to the long-horizon forecast because it does not include expectations of future transitory "plucks" in the estimate of trend. Importantly, the steady-state approach allows for a non-zero mean cycle and produces equivalent results to the Kalman filter estimates of trend and cycle reported in Kim and Nelson (1999a).

There are two features of the steady-state approach to trend/cycle decomposition that are worth emphasizing. First, it can be applied given a wide variety of forecasting models. Thus, while it produces equivalent results to the Kalman filter for Kim and Nelson's (1999a) plucking model, it is more general since it does not require a known unobserved components (UC) representation for the time series of interest. This generality has the desirable implication that evaluation of different estimates of trend and cycle can ultimately be thought of as a matter of model comparison given a set of possible forecasting models. Second, the concept of steady state is generally sufficient to identify the level of the trend. Thus, there is no need for the arbitrary normalization or identification assumptions that are often employed in unobserved components estimation of trend and cycle. Also, there is no need to impose theory-based identification assumptions that can bias or prejudge subsequent analysis. For example, structural VARs can identify permanent and transitory components through theoretical assumptions about long-run relationships (e.g., Blanchard and Quah, 1989). However, it is often these theoretical relationships that researchers want to investigate empirically, instead of assume *a priori*. The steady-state approach to trend/cycle decomposition allows such investigation without prejudging the results.[1]

The remainder of the paper is organized as follows. Section 2 presents the details of our steady-state approach to trend/cycle decomposition. Section 3 demonstrates the advantages of the steady-state approach when applied to integrated time series generated from known regime-switching processes, namely Hamilton's (1989) autoregressive Markov-switching model of output growth dynamics and Kim and Nelson's (1999a) plucking model. Section 4 presents an application of the approach to estimate the trend and cycle of U.S. real GDP implied by a regime-switching forecasting model developed by Kim, Morley, and Piger (2003). Section 5 concludes.

2. Method

The trend/cycle decomposition method proposed in this paper uses the concept of steady state to identify and estimate permanent and transitory components of integrated time series that follow nonlinear regime-switching processes. In this section, we present the details of the method. First, we introduce the conceptual framework that relates the trend of a time series to its implicit steady-state level. Second, we discuss the principles underlying the calculation of the conditional expectation of steady state given a forecasting model with regime switching parameters. Third, we compare our method to the BN decomposition.

[1] Clarida and Taylor (2003) make a similar point in arguing for the use of the BN decomposition.

2.1 Conceptual Framework

An integrated time series $\{y_t\}_{t=-\infty}^{\infty}$ can always be thought of as the sum of two unobserved components related to permanent and transitory innovations:

$$y_t = t_t + c_t \,. \tag{1}$$

The permanent component t_t is the accumulation of permanent h_t^* innovations:

$$t_t = \sum_{j=0}^{\infty} h_{t-j}^* \,. \tag{2}$$

The transitory component c_t is a weighted-average of transitory w_t^* innovations:

$$c_t = \sum_{j=0}^{\infty} y_{t,j} w_{t-j}^* \,, \tag{3}$$

where the random MA coefficients are normalized by $y_{t,0} = 1$.

We give meaning to the labels “permanent” and “transitory” by imposing the following restriction that links the permanent component to the concept of steady state:

$$\lim_{j \to \infty} y_{t,j} = 0 \,. \tag{4}$$

In words, transitory w_t^* innovations have no long-run impact on the time series $\{y_t\}$. Thus, in the hypothetical absence of any future innovations after time *t*, $\{y_t\}$ would converge to t_t. It is in this sense that t_t is the steady-state level of y_t.

In this paper, we focus our attention on integrated processes with regime switching parameters. In particular, we consider the case where the random MA coefficients in (3) and/or the means of the innovations in (2) and (3) can take on different values according to an unobservable discrete state variable with known distribution, which we denote as S_t:

$$y_{t,j} = y_{S_t,j}, \tag{5}$$

$$h_t^* = m_{S_t} + h_t, \tag{6}$$

$$w_t^* = l_{S_t} + w_t, \tag{7}$$

where $h_t \sim N(0, s_h^2)$, $w_t \sim N(0, s_w^2)$, and $Cov(h_t, w_t) = s_{hw}$.[2] The dependence of $y_{S_t,j}$, m_{S_t}, and l_{S_t} on the state variable S_t need not be linear. However, conditional on S_t, the process for $\{y_t\}$ is linear and assumed to be Gaussian.

The setup in (1)-(7) is quite general. Specifically, the innovations to the permanent and transitory components can be correlated. Also, from (5)-(7), the process can be regime switching for a variety of reasons. In particular, the regime switching can be in terms of the dynamics, the innovation to the permanent component, the innovation

[2] It is also possible to consider the case where the variance-covariance parameters depend on the state vector. We ignore this more general case for simplicity of presentation.

to the transitory component, or any combination of these. Despite this generality, the permanent and transitory components of a regime switching process can often be identified and estimated without prior knowledge of either the correlation between innovations or the source of regime switching. That is, a UC representation need not be known. Instead, as discussed next, the fact that t_t is the steady-state level of y_t provides a possible means for identification and estimation even when only a forecasting model of the first difference Δy_t is available.

2.2 Estimating Steady State

All of the forecasting models for Δy_t considered in this paper can be cast in the following state-space form:

$$\Delta y_t = [h_1 \quad h_2 \quad \cdots \quad h_n] X_t , \tag{8}$$

$$X_t = c_{\tilde{S}_t} + F_{\tilde{S}_t} X_{t-1} + e_t , \tag{9}$$

where $e_t \sim N(0,Q)$. That is, the observation Δy_t can be represented as a linear combination of the elements of the $n \times 1$ vector X_t according to the weights h_i, $i = 1,...,n$. Meanwhile, the vector X_t follows a first-order vector autoregressive process, where the elements of the intercept vector $c_{\tilde{S}_t}$ and the coefficient matrix $F_{\tilde{S}_t}$ may depend on a vector $\tilde{S}_t$ that contains the current and, possibly, lagged values of the unobservable discrete state variable S_t. The vector white noise error e_t is assumed to be Gaussian and uncorrelated with S_t. The first-order form of (9) is more general than it may at first appear since any higher order process can always be recast in its first-order companion form. Also, X_t can include both observables, such as Δy_t and/or its lags, and unobservables, such as moving-average error terms. Thus, the representation in (8) and (9) encompasses a wide variety of forecasting models including all univariate and multivariate Markov-switching ARMA models.

Estimation of the permanent and transitory components of y_t involves using the model given in (8) and (9) to calculate the conditional expectation of t_t. Note that in the absence of regime switching in the underlying time series process and the forecasting model, the BN decomposition calculates the conditional expectation of t_t (see Morley, Nelson, and Zivot, 2003). However, the presence of regime switching complicates identification and estimation of t_t and the BN decomposition, including its extension to nonlinear models by Clarida and Taylor (2003), does not necessarily provide a conditional expectation of t_t.

In calculating the conditional expectation of t_t, we start with the premise that conditional on observing X_t and the sequence $\{S_t\}$, it would be possible to identify and estimate t_t from a forecasting model using the fact that it is a steady-state value. That is, we can generally calculate $E[t_t \mid X_t, \{S_t\}, \Omega_t]$, where Ω_t is information observed at time *t*. First, given (9), solve recursively for the conditional expectation of future values of $\{X_t\}$:[3]

[3] Note that, if we had allowed for more complicated nonlinearities in (9) such as dependence of regimes on the residuals, we would need to use simulation and numerical integration as in Clarida and Taylor (2003) to

$$E[X_{t+j} \mid X_t, \{S_t\}, \Omega_t] = c_{\tilde{S}_{t+j}} + F_{\tilde{S}_{t+j}} E[X_{t+j-1} \mid X_t, \{S_t\}, \Omega_t]. \quad (10)$$

Then, from (8), the conditional expectation of future values of $\{y_t\}$ is

$$E[y_{t+j} \mid X_t, \{S_t\}, \Omega_t] = y_t + \sum_{i=1}^{j} [h_1 \quad h_2 \quad \cdots \quad h_n] E[X_{t+i} \mid X_t, \{S_t\}, \Omega_t]. \quad (11)$$

Suppose that for some $j = j^*$, the following steady-state condition holds:

$$E[\Delta y_{t+j^*+1} \mid X_t, \{S_t\}, \Omega_t] = \cdots = E[\Delta y_{t+j^*+l} \mid X_t, \{S_t\}, \Omega_t]. \quad (12)$$

That is, the expected change in the series remains constant for *l* periods. Given (12), $E[y_{t+j^*} \mid X_t, \{S_t\}, \Omega_t]$ provides a conditional expectation of steady state, and thus of the trend. To see why, note that based on (1), the conditional expectation of the series can always be decomposed as follows:

$$E[y_{t+j} \mid X_t, \{S_t\}, \Omega_t] = E[t_{t+j} \mid X_t, \{S_t\}, \Omega_t] + E[c_{t+j} \mid X_t, \{S_t\}, \Omega_t]. \quad (13)$$

---

calculate this condition expectation. The extension is conceptually straightforward, although computationally burdensome.

When $E[c_{t+j} \mid X_t, \{S_t\}, \Omega_t] \neq 0$, we would expect transitory dynamics to work their way out in future periods, meaning that the condition in (12) would not hold. Conversely, the steady-state condition in (12) implies $E[c_{t+j^*} \mid X_t, \{S_t\}, \Omega_t] = 0$.[4]

The intuition is that, in using (10) to estimate future values of $\{X_t\}$, we are setting future values of $\{e_t\}$ to their expected value of zero. This is equivalent to setting future values of the $\{h_t\}$ and $\{w_t\}$ shocks in (6)-(7) to their expected value of zero. As the horizon j gets large enough, the expectation is that future realizations of $\{y_t\}$ will no longer depend on transitory dynamics due to past realizations of $\{w_t\}$. Thus, at some point, any expected transitory dynamics that prevent the condition in (12) from holding must be caused by the mean of a future transitory innovation taking on non-zero values: i.e., $l_{S_{t+j}} \neq 0$. Meanwhile, if the condition in (12) holds, future transitory innovations due to regime changes are equal to zero for a long enough period for the expected impact of any previous non-zero transitory innovations to die out.

While $E[y_{t+j^*} \mid X_t, \{S_t\}, \Omega_t]$ is a conditional expectation of a steady state, it is not the current steady state t_t. In particular, $E[y_{t+j^*} \mid X_t, \{S_t\}, \Omega_t]$ will reflect future permanent innovations implied by future realizations of $\{S_t\}$. These innovations must be

[4] There is a theoretical possibility that for some set of random MA coefficients in (3) and a particular sequence of transitory shocks, the steady-state condition in (12) may hold by mere coincidence and not because the impact of past transitory innovations has died out. To minimize this possibility, the number of consecutive periods *l* that the steady-state condition is required to hold can be set to an arbitrarily large number. By the long-run restriction in (4), a constant, but non-zero impact of a past transitory innovation must eventually change. Thus, a large enough *l* will prevent an erroneous estimate of steady state, although at the cost of increasing the number of periods before the steady-state condition holds. Depending on the nature of the process, it is possible that the condition in (12) never holds for *l* consecutive periods and we cannot identify trend using the steady-state approach. However, these problems do not arise in the applications considered in this paper, where we set *l*=2 and find that results are robust to higher values of *l*.

removed from $E[y_{t+j^*} \mid X_t, \{S_t\}, \Omega_t]$ in order to measure $E[t_t \mid X_t, \{S_t\}, \Omega_t]$. We measure the long-run impact implied by each specific future realization S_{t+j} as

$$m_{S_{t+j}} \equiv E[\Delta y_t \mid \{S_t\} = S_{t+j}]\,. \tag{14}$$

In particular, the conditional expectation in (14) calculates the long-run growth rate within each regime. Then, the conditional estimate of the current steady state level is

$$E[t_t \mid X_t, \{S_t\}, \Omega_t] = E[y_{t+j^*} \mid X_t, \{S_t\}, \Omega_t] - \sum_{j=1}^{j^*} m_{t+j}\,. \tag{15}$$

In practice, X_t and the sequence $\{S_t\}$ are not observed. We only assumed that they were available conditioning information because it allowed us to use the concept of steady state to identify the trend.[5] Given the conceptual framework above, we would be unable to identify trend without that specific conditioning information. However, once we have identified the trend in theory, we can use Monte Carlo simulation methods to integrate out the unknown conditioning variables in order to calculate the expectation of the trend conditional only on observed data.

In simulating X_t and the sequence $\{S_t\}$, we follow the multi-move Gibbs-sampling procedure employed by Carter and Kohn (1994) and Kim and Nelson (1998,

[5] The choice of the conditioning set is fundamental to identifying steady state. It also turns out to be important to condition only on information up to time *t*, rather than the whole sample. If we condition on future data, our inferences about the future steady state will reflect future permanent shocks, but these shocks are not observed, so we cannot easily identify and remove them to estimate the current steady state.

1999b).[6] The details of the procedure are somewhat involved and are presented in an appendix. However, we sketch out the main features here.

Given an arbitrary initial sequence $\{S_t^{(0)}\}_{t=1}^{t+J}$, where J is sufficiently large to always include j^*, and $m = 1$, iterate through the following steps:

1. Draw $\{X_t^{(m)}\}_{t=1}^{t}$ from the conditional distribution $\{X_t\} | \{S_t^{(m-1)}\}_{t=1}^{t}, \Omega_t$.

2. Draw $\{S_t^{(m)}\}_{t=1}^{t+J}$ from the conditional distribution $\{S_t\} | \{X_t^{(m)}\}_{t=1}^{t}, \Omega_t$.

3. Update $m = m + 1$ and return to step 1 until $m > M$.[7]

For each draw of $X_t^{(m)}$ and the sequence $\{S_t^{(m)}\}$, we can estimate steady state using the condition in (12) to solve for the expectation in (15). Then, denoting the simulated steady-state estimate as $\hat{t}_t^{(m)} = E[t_t \mid X_t^{(m)}, \{S_t^{(m)}\}, \Omega_t]$, our estimate of steady state is

$$\hat{t}_t = \frac{1}{M} \sum_{m=1}^{M} \hat{t}_t^{(m)}, \tag{16}$$

which should converge to the conditional expectation $E[t_t \mid \Omega_t]$ as the number of simulations *M* goes to infinity. Specifically,

$$\lim_{M \to \infty} \left\{ \frac{1}{M} \sum_{m=1}^{M} \hat{t}_t^{(m)} \right\} = E\left[E[t_t \mid X_t^{(m)}, \{S_t^{(m)}\}, \Omega_t] \middle| \Omega_t\right] = E[t_t \mid \Omega_t], \tag{17}$$

[6] For simplicity of presentation, we assume known parameters. In practice, estimation of parameters can be done using either classical or Bayesian methods. Bayesian estimation simply involves the additional step in the Gibbs-sampling procedure of drawing model parameters from their conditional posterior distributions.
[7] In practice, we only keep track of draws after an initial "burn in" period in order to ensure draws are no longer impacted by our initial arbitrary sequence for the Markov-switching state variable.

where the second equality follows from the law of iterated expectations. The estimated transitory component, denoted $\hat{c}_t$, is simply the difference between the level of the series and the estimated trend: $\hat{c}_t = y_t - \hat{t}_t$.

2.3 Comparison with BN Decomposition

Given a linear forecasting model, the steady-state estimate of trend is always equivalent to the long-horizon forecast (minus any deterministic growth) estimate developed by Beveridge and Nelson (1981):

$$BN_t = \lim_{j \to \infty} \{E_t[y_{t+j} - j \cdot m]\}. \tag{18}$$

A linear model is a model that can be cast as a special case of (8) and (9) in which the intercept vector c and the coefficient matrix F do not depend on S_t. For such a model, we can calculate $\hat{t}_t = E[t_t \mid \Omega_t]$ using a simplified version of the steady-state approach outlined in (10)-(16).

A simple example may help illustrate the equivalence of the steady-state approach and the BN decomposition for linear models. Consider an integrated time series that can be forecast by a stationary AR(1) model in first differences:

$$\Delta y_t - m = f(\Delta y_{t-1} - m) + e_t, \tag{19}$$

where $|f| < 1$ and $e_t \sim N(0, s_e^2)$. For this model, the state-space representation is simply the intercept form of the AR(1) model, where the intercept parameter $c = m(1-f)$. Then, the expectation in (11) simplifies to $E[y_{t+j} \mid \Omega_t] = E[y_{t+j} \mid X_t, \{S_t\}, \Omega_t]$ since X_t is observed ($X_t = \Delta y_t$) and the process does not depend on $\{S_t\}$. For this model, it is straightforward to solve for this expectation using (11) and the expected change in the series implied by (19):

$$E[\Delta y_{t+j} \mid \Omega_t] = m + f^j(\Delta y_t - m), \tag{20}$$

Since $\lim_{j\to\infty} f^j = 0$, there will be a j^* that satisfies the steady-state condition in (12) as $j \to \infty$:

$$E[\Delta y_{t+j^*+1} \mid \Omega_t] = \cdots = E[\Delta y_{t+j^*+l} \mid \Omega_t] = m\,. \tag{21}$$

The deterministic growth m reflects the expectation of future permanent innovations. Thus, in order to calculate the conditional expectation of the current steady state, we must remove the impact of expected future permanent innovations on $E[y_{t+j^*} \mid \Omega_t]$, as is done in (15):

$$E[t_t \mid \Omega_t] = E[y_{t+j^*} \mid \Omega_t] - j^* \cdot m\,. \tag{22}$$

Given the forecasting equation in (20) and $j \to \infty$, it is straightforward to show that (22) is equivalent to

$$\hat{t}_t = y_t + \frac{f}{1-f}(\Delta y_t - m), \tag{23}$$

which is the same as the BN trend for an AR(1) forecasting model (see Morley, 2002).

By contrast, the steady-state approach and the BN decomposition are generally not the same for nonlinear forecasting models with regime-switching parameters as in (8) and (9). The reason is that the BN estimate of trend in (18) implicitly includes the expectation of future innovations that are not part of the actual current trend. In the next section, we demonstrate this difference for two well-known regime-switching processes.

3. Examples

In this section, we present two examples of the steady-state approach to trend/cycle decomposition applied to integrated time series generated from regime-switching processes. In both cases, we generate data from unobserved components models, allowing us to compare estimates of the trend and cycle with their true values. We also use the generated data to illustrate how the steady-state approach identifies the level of trend and the mean of the cycle.

3.1 Hamilton's (1989) Markov-Switching Model

The first data generating process we consider is based on Hamilton's (1989) regime-switching model of U.S. real GNP. The original specification is a forecasting model of the first differences, Δy_t. However, in order to observe the actual permanent and transitory components for the generated data, we consider the following unobserved-components representation for y_t:

$$
\begin{aligned}
y_t &= t_t + c_t \\
t_t &= m_0 + m_1 S_t + t_{t-1} + h_t \\
f(L)c_t &= w_t
\end{aligned}
\qquad (24)
$$

where $h_t \sim N(0, s_h^2)$, $w_t \sim N(0, s_w^2)$, $\mathrm{cov}(h_t, w_t) = 0$, $f(L) = 1 - \sum_{i=1}^{k} f_i L^i$ is a $k-th$ order lag polynomial will all roots outside the unit circle, and $S_t = \{0,1\}$ is a Markov-switching state variable with transition probabilities $\Pr[S_t = 0 \mid S_{t-1} = 0] = q$ and $\Pr[S_t = 1 \mid S_{t-1} = 1] = p$. The permanent component of real output, t_t, follows a random walk with a regime-switching drift component. The transitory component, c_t, follows a linear stationary autoregressive process. In the example, we set lag length $k = 2$, so that c_t is an AR(2) process.

We generate a sample of 200 observations from the data generating process in (24) using the following parameter calibration: $m_0 = 0.8$, $m_1 = -1.3$, $s_h = 0.25$, $s_w = 0.25$, $f_1 = 0.2$, $f_2 = 0.1$, $q = 0.96$ and $p = 0.8$. In practice, given this UC model and parameters, the trend and cycle can be estimated using the filtering techniques in Lam (1990) and Kim (1994). However, we consider the steady-state approach here in order to evaluate its performance. Of course, it should be emphasized that, unlike the filtering techniques, the steady-state approach does not require knowledge of the UC representation or any arbitrary normalization or identification assumptions that often accompany UC models. It can be applied given a forecasting model for the first differences only.

For the steady-state approach to trend/cycle decomposition, we cast a forecasting model into the state-space form given by (8) and (9). One way to do so for the UC model in (24) is to solve for corresponding reduced-form representation for Δy_t,[8] which in this case is a regime-switching ARMA(2,2) process:

$$(1 - f_1 L - f_2 L^2)(\Delta y_t - m_0 - m_1 S_t) = (1 + q_1 L + q_2 L^2) e_t, \qquad (25)$$

where $e_t \sim N(0, s_e^2)$ and q_1 and q_2 are complicated nonlinear functions of the parameters in (24).[9] In terms of the state-space form given by (8) and (9), the variable vectors are $\Delta y_t = [0 \quad 0 \quad 1 \quad 0] X_t$, $X_t = [e_t \quad e_{t-1} \quad \Delta y_t \quad \Delta y_{t-1}]'$, and $e_t = [e_t \quad 0 \quad e_t \quad 0]'$. The intercept vector is $c_{\tilde{S}_t} = [0 \quad 0 \quad a_{\tilde{S}_t} \quad 0]$, where $a_{\tilde{S}_t} \equiv (1 - f_1 L - f_2 L^2)(m_0 + m_1 S_t)$, which means that $\tilde{S}_t = [S_t \quad S_{t-1} \quad S_{t-2}]'$. Finally, the companion matrix, which does not depend on $\tilde{S}_t$ for the model in (25), is

[8] It should be noted that we could have directly cast the UC model in (24) into the state-space form given by (8) and (9). However, the point of the example is to show that trend and cycle can be estimated even if only a reduced-form representation were known.

[9] We calculate values for the moving-average parameters by solving the system of nonlinear equations that relates the reduced-form parameters to the autocovariances implied by (24). Of the multiple solutions available, we use the one that corresponds to real numbers for the parameters and an invertible representation for the moving-average component.

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ q_1 & q_2 & f_1 & f_2 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Given the model in this state-space form, $\hat{t}_t = E[t_t \mid \Omega_t]$ can be calculated using the steady-state approach outlined in (10)-(16). We perform M = 500 simulations to integrate X_t and the sequence $\{S_t\}$ out of the conditional distribution $E[t_t \mid X_t, \{S_t\}, \Omega_t]$ after an initial 1000 simulations to ensure convergence of the Gibbs sampler.

Figure 1 displays the generated time series, its true trend, and some estimates of the trend. In addition to the steady-state estimate of trend, we consider the HP filter estimate, with the smoothing parameter set to 1600, and an estimate based on the long-horizon forecast (i.e., the BN decomposition) calculated using the method presented in Clarida and Taylor (2003).[10] The first panel of Figure 1 shows that most of the variation in the generated series reflects the regime switching in the trend. The second panel shows that the steady-state approach is able to capture this form of nonlinearity. Indeed, the steady-state estimate of trend is virtually indistinguishable from the true trend. Meanwhile, the third panel shows that the HP filter estimate of trend misses the nonlinearity in the data. Instead, it essentially traces out a smooth line through the series. Finally, the fourth panel shows that the long-run forecast estimate of trend is much more variable than the true trend. In particular, whenever there is a change in regime, the

[10] Briefly, Clarida and Taylor's (2003) method involves generating a large number of simulated future realizations of the time series from the conditional distribution implied by a forecasting model and the observed data at each point of time. Then, the conditional expectation of the series at any future horizon is calculated by averaging the simulated future realizations at that horizon. As with Beveridge and Nelson (1981), the estimate of trend is the long-horizon conditional expectation of the series minus any deterministic growth, which is defined as the unconditional expectation of the change in the time series.

estimated trend adjusts by more than the series itself due to the persistence of the regimes. For example, the persistence of recessionary regimes means that a transition from an expansionary regime to a recessionary regime will greatly increase the probability that future regimes are also recessionary. Thus, the long-run conditional expectation of the series will be dramatically lowered by the regime shift.

3.2 Kim and Nelson's (1999a) "Plucking" Model

The second data generating process we consider is based on the plucking model of U.S. real GDP proposed by Kim and Nelson (1999a).[11] Instead of allowing for regime switching in the permanent component, Kim and Nelson assume that the regime switching is in the transitory component. In particular, the model is

$$\begin{aligned} y_t &= t_t + c_t \\ t_t &= m + t_{t-1} + h_t \\ f(L)c_t &= lS_t + w_t \end{aligned} \tag{26}$$

where, h_t, w_t, $f(L)$ and S_t are the same as in (24). The permanent component of real output, t_t, follows a random walk with constant drift. The transitory component, c_t, follows a stationary autoregressive process that depends on the value of S_t. For example, if $l < 0$, c_t is "plucked" downward by the amount l when $S_t = 1$. However, the effects of this pluck are transitory, as they are worked off through the dynamics of c_t. Importantly, note that this model implies that the transitory component, c_t, has a non-

[11] The "plucking" terminology is due to Milton Friedman (1964, 1993).

zero mean given by $E(c_t) = \dfrac{l \cdot \Pr[S_t = 1]}{f(1)}$, where $\Pr[S_t = 1] = \dfrac{1-q}{2-p-q}$. As before, we set lag length $k = 2$.

We generate a sample of size 200 from the data generating process in (26) using the following calibration: $m = 0.8$, $s_h = 0.25$, $s_w = 0.02$, $l = -2$, $f_1 = 1.2$, $f_2 = -0.4$ $p = 0.96$ and $q = 0.80$. Consistent with the parameter estimates of Kim and Nelson (1999a), s_w is very small relative to l, meaning most variation in c_t is due to "plucks". The reduced-form representation for Δy_t implied by the UC model in (26) is again a regime-switching ARMA (2,2) process:

$$(1 - f_1 L - f_2 L^2)(\Delta y_t - m) = l \cdot \Delta S_t + (1 + q_1 L + q_2 L^2) e_t, \qquad (27)$$

where e_t, q_1, and q_2 are the same as in (25). The model in (27) can be cast in the state-space form given by (8) and (9) in much the same way as the Hamilton model above. The only difference is in the $a_{\tilde{S}_t}$ element of the intercept vector $c_{\tilde{S}_t}$. Specifically, for the model in (27), $a_{\tilde{S}_t} \equiv (1 - f_1 - f_2)m + l\Delta S_t$ which means that $\tilde{S}_t = \begin{bmatrix} S_t & S_{t-1} \end{bmatrix}'$.

Again, $\hat{t}_t = E[t_t \mid \Omega_t]$ can be calculated using the steady-state approach outlined in (10)-(16). As before, we perform $M = 500$ simulations to integrate X_t and the sequence $\{S_t\}$ out of the conditional distribution $E[t_t \mid X_t, \{S_t\}, \Omega_t]$ after an initial 1000 simulations to ensure convergence of the Gibbs Sampler.

Figure 2 displays the series, the true trend, and some estimates of the trend. Again, in addition to the steady-state estimate, we display the HP filter estimate and the long-horizon forecast estimate calculated using the method presented in Clarida and Taylor (2003). The first panel of Figure 2 shows the "plucking" nature of the generated series. It remains very close to trend except when it is plucked in the downward direction. As with the previous example, the second panel shows that the steady-state estimate of trend is able to capture the true trend. The third panel shows that the HP filter estimate of trend misses the nonlinear aspects of the data and essentially traces out a smooth line through the series, labeling some of the transitory variation in the series as variation in the trend. Finally, the fourth panel shows that the long-run forecast estimate of trend moves with the true trend, but is shifted downward. That is, the long-run forecast method fails to identify the level of the trend or, equivalently, the mean of the cycle.

This can be seen more clearly in Figure 3, which displays the true generated cycle for the plucking model along with the steady-state, H-P filter and long-run forecast estimates of the cycle. The ability of the steady-state estimate to capture the negative mean of the true cycle is apparent. Meanwhile, the H-P filter and long-run forecast technique produce estimates of the cycle that are shifted upward from the true cycle. Indeed, the ability of the steady-state approach to identify the level of the trend or, equivalently, the mean of the cycle, is a key advantage over other methods. We turn to this issue in more detail next.

3.3 Identifying the Level of Trend

It may appear at first that the idea of cycle with a non-zero mean is merely a matter of normalization. That is, the initial level of the trend could always be set to make the resulting cycle have a mean of zero. However, the resulting re-normalized trend may not correspond to steady state. Put another way, the concept of steady state uniquely identifies the level of the trend or, equivalently, the mean of the cycle. Kim and Nelson's (1999a) plucking model allows for a straightforward demonstration of this point.

Figure 4 displays the level and the first-difference of the cycle generated from the plucking model. While it would be possible to shift the level of the cycle upwards to have a mean of zero, it should be noted that doing so would have no impact on the first-difference of the cycle. Meanwhile, it is the first-difference of the cycle that identifies when the series is in steady state. In particular, the series can only be thought to be in steady state when the first-difference of the cycle is constant and equal to zero. Otherwise, the cycle must be changing, which from (3) means that past transitory innovations are still affecting the cycle and the series is not in steady state. Thus, we can use the first-difference of the cycle and the concept of steady state to identify the mean of the cycle or, equivalently, the level of the trend. Specifically, steady state implies that the cycle is equal to zero and the trend is equal to the actual level of the series only when the first-difference of the cycle is equal to zero. This relationship holds for the cycle displayed in Figure 4, but would not hold for a re-normalized cycle with a mean of zero.

4. Application

Separating transitory or cyclical variation in time series of macroeconomic activity from permanent or trend variation has a rich history in macroeconomics. A primary reason for the attention this enterprise has drawn is that the "business cycle" is typically measured using the cyclical component of an output series such as real Gross Domestic Product (GDP), while the trend component is used to measure long-run growth.[12] In order to test macroeconomic theories of the business cycle or long-run growth, we need accurate estimates of these components.

There are many alternative techniques that have been used to extract trend and cycle from real GDP. A popular early technique was to assume that the trend is a deterministic polynomial. More recently, the focus has shifted to stochastic trends, in which the permanent component is not merely a deterministic function of time, but contains random elements. Popular techniques to extract stochastic trends from real GDP include the Hodrick-Prescott (1997) filter, the linear unobserved components models of Clark (1987) and Watson (1986), and the Beveridge-Nelson (1981) decomposition.

Much of this work is based on the assumption that the data generating process for real GDP is linear. Indeed, the UC models referenced above and the BN decomposition are based on ARIMA forecasting models for real GDP and are, therefore, explicitly linear. At the same time, there is growing evidence that the time-series properties of real GDP are well described by models containing departures from linear ARIMA models, including threshold models and regime-switching models. However, the implications of these nonlinear dynamics for measuring trend and cycle have been largely ignored. Exceptions are generally within the UC framework, where nonlinearities are explicitly entered into the processes for the trend and cycle (see, for example, Kim and Nelson, 1999a).

In this section, we apply the steady-state approach to extract the trend and cycle from U.S. real GDP under the assumption of a nonlinear data generating process. We

[12] This statistical definition of trend and cycle is not without problems. For example, Blanchard and Quah (1989) point out that the structural trend of output may also experience transitory fluctuations and thus trend need not correspond to the permanent component of output.

note that the steady-state approach requires a forecasting model. One advantage of the steady state approach is that it can be applied given any forecasting model, which reduces the problem of evaluating a particular decomposition to a matter of model selection. However, performing model selection across an appropriately large set of forecasting models is beyond the scope of this paper. Thus, we focus here on a particular nonlinear forecasting model of U.S. real GDP developed by Kim, Morley, and Piger (KMP) (2003) and leave model selection to future research.

The KMP model is a regime-switching ARIMA model for U.S. real GDP:

$$f(L)\left(\Delta y_t - m_0 - m_1 S_t - l\sum_{j=1}^{\overline{m}} S_{t-j}\right) = e_t\,, \tag{28}$$

where $e_t \sim N(0, s_e^2)$, the lag operator $f(L)$ is k-th order with roots outside the unit circle, Δy_t is the first difference of the logarithm of real GDP, and S_t is an unobserved Markov-switching state variable that takes on discrete values of 0 or 1 according to transition probabilities $\Pr[S_t = 0 \mid S_{t-1} = 0] = q$ and $\Pr[S_t = 1 \mid S_{t-1} = 1] = p$. The states are normalized by restricting $m_1 < 0$. That is, $S_t = 1$ corresponds to a "lower growth" regime or, if $m_0 + m_1 < 0$, a "contractionary" regime.

The model in (28) admits a three-phase representation of GDP growth dynamics. When $S_t = 0$, output grows at the rate m_0, with deviations from this growth rate caused by e_t which are propagated by the autoregressive dynamics $f(L)$. When $S_t = 1$, output enters a low growth or contractionary regime and grows at the average rate $m_0 + m_1$.

However, as the recession progresses and ultimately ends, the term $l\sum_{j=1}^{\overline{m}} S_{t-j}$ augments the growth rate of the forecasting equation. This additional term is best described as a pressure variable, obtaining larger values as the recession is longer and more severe. It then implies a "bounce-back" effect from the recession if $l > 0$. That is, output growth will be above-average for the first $\overline{m}$ periods of an "expansionary" regime. Using this model, KMP are able to reject both linearity and the absence of a "bounce-back" effect in U.S. real GDP. Also, in a formal model comparison exercise, KMP find that this three-phase model outperforms linear ARIMA models and alternative regime-switching models in mimicking a number of key features in the data.

We estimate the KMP model using the first difference of log U.S. real GDP from 1949:Q1 to 2003:Q1. The series is multiplied by 400 in order to correspond to annualized growth rates. We set the lag order $k = 2$ and, following the findings in KMP, $\overline{m} = 6$ quarters. The model is estimated via maximum likelihood using the filter given in Hamilton (1989). Table 1 reports the estimated parameters, while Figure 5 plots the smoothed probability $\Pr[S_t = 1 | \Omega_T]$, along with the log real GDP series. As in KMP, $m_0 + m_1 < 0$, which suggests that $S_t = 1$ corresponds to a "contractionary" regime. Indeed, the smoothed probability $\Pr[S_t = 1 | \Omega_T]$ suggests that the contractionary regime corresponds fairly closely to most NBER recession dates. Meanwhile, $l > 0$ suggests that the quarters following the end of a contractionary regime correspond to above-average economic growth.

To obtain the steady-state estimate of trend, we cast the model in (28) into state-space form. In terms of the state-space form given by (8) and (9), the variable vectors are

$\Delta y_t = [1 \;\; 0] X_t$, $X_t = [\Delta y_t \;\; \Delta y_{t-1}]'$, and $e_t = [e_t \;\; 0]'$. The intercept vector is $c_{\tilde{S}_t} = [a_{\tilde{S}_t} \;\; 0]'$, where $a_{\tilde{S}_t} \equiv (1 - f_1 L - f_2 L^2)(m_0 + m_1 S_t + l \sum_{j=1}^{\bar{m}} S_{t-j})$, which means that $\tilde{S}_t = [S_t \;\; S_{t-1} \;\; S_{t-2} \;\; S_{t-3} \;\; S_{t-4} \;\; S_{t-5} \;\; S_{t-6} \;\; S_{t-7} \;\; S_{t-8}]'$. Finally, the companion matrix is:

$$F = \begin{bmatrix} f_1 & f_2 \\ 1 & 0 \end{bmatrix}$$

We set each parameter to its maximum likelihood point estimate listed in Table 1. Finally, we calculate $\hat{t}_t = E[t_t \mid \Omega_t]$ using the steady-state approach outlined in (10)-(16). We perform M = 500 simulations to integrate X_t and the sequence $\{S_t\}$ out of the conditional distribution $E[t_t \mid X_t, \{S_t\}, \Omega_t]$ after an initial 1000 simulations to ensure convergence of the Gibbs sampler.

Figure 6 displays the steady-state estimates of the trend and cycle for log U.S. real GDP. The estimated cycle has several noticeable features. First, it clearly has a negative mean, suggesting that it is more typical for the economy to operate below trend than above it. Second, it turns strongly negative during most NBER dated recessions, while it is near zero during expansions. This finding is consistent with Milton Friedman's (1964, 1993) plucking model, in which output is generally close to trend, and is occasionally "plucked" below trend. We note, however, that for the three recessions of 1970, 1990-91, and 2001 most of the movement in real GDP is in the trend component. That is, according to the model, these recessions had largely permanent effects on real GDP.

5. Conclusion

The steady-state approach to trend/cycle decomposition has many advantages over competing methods. It is able to deal with nonlinearities in the data and it identifies the level of the trend or, equivalently, the mean of the cycle. The application to U.S. real GDP reveals these advantages and portrays a picture of the business cycle that is very different than what is implied by standard linear models.

We conclude by noting that our approach, by allowing for nonlinear dynamics, provides a possible reconciliation of the NBER notion of the business cycle with the concept of an output gap. In particular, according to the forecasting model employed in our application, the transitory component of real GDP is very small in expansions and large and negative only during NBER recessions. Meanwhile, the fact that some NBER recessions correspond only to movements in the permanent component suggests that these episodes may reflect fundamentally different economic conditions than the other recessions. Of course, even within the trend/cycle decomposition framework developed in this paper, a more complete investigation of the business cycle must involve some consideration of alternative forecasting models. We leave the model selection issue and a full resolution of what the business cycle actually looks like to future research.

Appendix: Sampling the Unobserved $\{X_t\}$ and $\{S_t\}$

The steady-state approach to trend-cycle decomposition outlined in Section 2 requires a step to draw simulated values of $\{X_t\}_{t=1}^{t}$ and $\{S_t\}_{t=1}^{t+J}$ from their conditional distributions $\{X_t\}|\{S_t\}_{t=1}^{t},\Omega_t$ and $\{S_t\}|\{X_t\}_{t=1}^{t},\Omega_t$. In the following we describe algorithms for this data generation. The presentation will closely follow that in Kim and Nelson (1999b).

A.1 Sampling $\{X_t\}$

To draw from $\{X_t\}_{t=1}^{t}|\{S_t\}_{t=1}^{t},\Omega_t$ we use the multimove algorithm detailed in Carter and Kohn (1994). To begin, the state space form of the forecasting model is written so that the first $R \ x \ R$ block of Q, denoted Q^*, is positive definite, whereas the remaining elements of Q are zero. Denote the first R elements of X_t as X_t^* and the first R rows of $F_{\tilde{S}_t}$ as $F_{\tilde{S}_t}^*$. As an example, consider the ARMA(2,2) model in (25):

$$(1-f_1L-f_2L^2)(\Delta y_t - m_0 - m_1S_t) = (1+q_1L+q_2L^2)e_t$$

This forecasting model has a state-space form that meets the requirements for the Q matrix given above:

$$(1-f_1L-f_2L^2)(\Delta y_t - m_0 - m_1S_t) = \begin{bmatrix}1 & q_1 & q_2\end{bmatrix}X_t$$

$$X_t = c_{\tilde{S}_t} + F_{\tilde{S}_t} X_{t-1} + e_t$$

where $X_t = [e_t \quad e_{t-1} \quad e_{t-2}]'$, $e_t = [e_t \quad 0 \quad 0]'$, $c_{\tilde{S}_t} = [0 \quad 0 \quad 0]'$, $F_{\tilde{S}_t} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$ and

$Q = \begin{bmatrix} s_e^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$. Here, $R = 1$, so $Q^* = s_e^2$, X_t^* is e_t and $F_{\tilde{S}_t}^* = [0 \quad 0 \quad 0]$.

The distribution of interest, $\{X_t\}_{t=1}^t \mid \{S_t\}_{t=1}^t, \Omega_t$, is then factored as:

$$\{X_t\}_{t=1}^t \mid \{S_t\}_{t=1}^t, \Omega_t = \left(X_t \mid \{S_t\}_{t=1}^t, \Omega_t\right) \cdot \prod_{i=1}^{t-1} \left(X_i \mid X_{i+1}^*, \{S_t\}_{t=1}^t, \Omega_i\right)$$

This factorization is established by the Markov structure of X_i. In particular, conditional on X_{i+1} and Ω_i, there is no additional information regarding X_i in X_{i+2} and Ω_{i+1}. Also, note that X_i is conditioned on X_{i+1}^* rather than on X_{i+1}. This is because the remaining terms in X_{i+1} beyond those in X_{i+1}^* determine elements of X_i with certainty, which creates a singularity that makes generating X_i impossible.

Given this factorization, a draw from $\{X_t\}_{t=1}^t \mid \{S_t\}_{t=1}^t, \Omega_t$ can proceed by first drawing a realization of X_t from $\left(X_t \mid \{S_t\}_{t=1}^t, \Omega_t\right)$, denoted $X_t^{(g)}$ and then drawing recursively from $\left(X_i \mid X_{i+1}^{*(g)}, \{S_t\}_{t=1}^t, \Omega_i\right)$. Operationally, this is performed by first running the Kalman Filter on the state-space model to compute and save $\{X_{t|t}\}_{t=1}^t$ and $\{P_{t|t}\}_{t=1}^t$, where $X_{t|t}$ and $P_{t|t}$ are the filtered estimate of X_t and its variance-covariance matrix respectively. The Gaussian structure of the state-space model implies $\left(X_t \mid \{S_t\}_{t=1}^{t}, \Omega_t\right) \sim N\left(X_{t|t}, P_{t|t}\right)$, which can be used to generate $X_t^{(g)}$. We then simply need $\left(X_i \mid X_{i+1}^{*(m)}, \{S_t\}_{t=1}^{t}, \Omega_i\right)$ to complete the algorithm. Carter and Kohn (1994) show that this is given by:

$$\left(X_i \mid X_{i+1}^{*(g)}, \{S_t\}_{t=1}^{t}, \Omega_i\right) \sim N\left(X_{i|i, X_{i+1}^*}, P_{i|i, X_{i+1}^*}\right),$$

where:

$$X_{i|i, X_{i+1}^*} = E(X_i \mid X_{i+1}^{(g)*}, \{S_t\}_{t=1}^{t}, \Omega_i) = X_{i|i} + P_{i|i} F_{\tilde{S}_t}^{*'} (F_{\tilde{S}_t}^{*} P_{i|i} F_{\tilde{S}_t}^{*'} + Q^*)^{-1} (X_{i+1}^{*(g)} - c_{\tilde{S}_t}^{*} - F_{\tilde{S}_t}^{*} X_{i|i})$$

$$P_{i|i, X_{i+1}^*} = Cov(X_i \mid X_{i+1}^{*}, \{S_t\}_{t=1}^{t}, \Omega_i) = P_{i|i} - P_{i|i} F_{\tilde{S}_t}^{*'} (F_{\tilde{S}_t}^{*} P_{i|i} F_{\tilde{S}_t}^{*'} + Q^*)^{-1} F_{\tilde{S}_t}^{*} P_{i|i}$$

Note that for many forecasting models, such as the regime-switching autoregressive model in (28), $\{X_t\}$ is observed. In this case, we simply condition on the observed $\{X_t\}$, rather than sample values of $\{X_t\}$ from its conditional distribution.

A.2 Sampling $\{S_t\}$

To draw from $\{S_t\}_{t=1}^{t+J} \mid \{X_t\}_{t=1}^{t}, \Omega_t$ we use the multimove algorithm detailed in Kim and Nelson (1998). To begin, note that $\{S_t\}_{t=1}^{t} \mid \{X_t\}_{t=1}^{t}, \Omega_t$ can be factored as:

$$\{S_t\}_{t=1}^{t} \mid \{X_t\}_{t=1}^{t}, \Omega_t = \left(S_t \mid \{X_t\}_{t=1}^{t}, \Omega_t\right) \cdot \prod_{i=1}^{t-1} \left(S_i \mid S_{i+1}, \{X_t\}_{t=1}^{t}, \Omega_i\right).$$

Again, this factorization is established by the Markov nature of S_i. In particular, conditional on S_{i+1} and Ω_i, there is no additional information regarding S_i in S_{i+2} and Ω_{i+1}. Given this, a draw from $\{S_t\}_{t=1}^{t} \mid \{X_t\}_{t=1}^{t}, \Omega_t$ can proceed by first drawing a realization of S_t from $\left(S_t \mid \{X_t\}_{t=1}^{t}, \Omega_t\right)$, denoted $S_t^{(g)}$ and then drawing recursively from $\left(S_i \mid S_{i+1}^{(g)}, \{X_t\}_{t=1}^{t}, \Omega_i\right)$.

Operationally, this is performed by first running the Hamilton (1989) filter on the forecasting model, conditional on $\{X_t\}$. For example, conditional on $\{X_t\}$, and thus on $\{e_t\}$, the forecasting model in (25) is an ARMA (2,2) for which the moving average component, $q_1 e_{t-1} + q_2 e_{t-2}$ is observed. The Hamilton (1989) filter produces the filtered probabilities:

$$\left\{\Pr\left[S_t = w \mid \{X_t\}_{t=1}^{t}, \Omega_t\right]\right\}_{t=1}^{t}$$

The final filtered probability, $\Pr\left[S_t = w \mid \{X_t\}_{t=1}^{t}, \Omega_t\right]$, gives us $\left(S_t \mid \{X_t\}_{t=1}^{t}, \Omega_t\right)$, from which we can generate $S_t^{(g)}$. We then draw recursively from $\left(S_i \mid S_{i+1}^{(g)}, \{X_t\}_{t=1}^{t}, \Omega_i\right)$, which is given by:

$$\Pr\left[S_i = w \mid S_{i+1}^{(g)}, \{X_t\}_{t=1}^{t}, \Omega_i\right] = \frac{\Pr\left[S_{i+1}^{(g)} \mid S_i = w\right] \cdot \Pr\left[S_i = w \mid \{X_t\}_{t=1}^{t}, \Omega_i\right]}{\sum_{w=0}^{1} \Pr\left[S_{i+1}^{(g)} \mid S_i = w\right] \cdot \Pr\left[S_i = w \mid \{X_t\}_{t=1}^{t}, \Omega_i\right]}$$

where $\Pr\left[S_{i+1}^{(g)} \mid S_i = w\right]$ are simple functions of the transition probabilities, p and q.

Finally, given $S_t^{(g)}$, $\Pr\left[S_{t+1} = w \mid S_t^{(g)}, \{X_t\}_{t=1}^{t}, \Omega_t\right] = \Pr\left[S_{t+1} = w \mid S_t^{(g)}\right]$ is again a simple function of the transition probabilities, p and q, and can be used to generate $S_{t+1}^{(g)}$. In turn, $\Pr\left[S_{t+2} = w \mid S_{t+1}^{(g)}\right]$ is then used to generate $S_{t+2}^{(g)}$. This is repeated to generated $\{S_t^{(g)}\}_{t=t+1}^{t+J}$.

References

Beveridge, S. and C.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, 151-174

Blanchard, O. and D. Quah, 1989, The Dynamic Effects of Aggregate Demand and Supply Disturbances, *American Economic Review* 79, 655-673.

Carter, C. K. and P. Kohn, 1994, On Gibbs Sampling for State Space Models, *Biometrica* 81, 541-553.

Clarida, R. H. and M.P. Taylor, 2003, Nonlinear Permanent-Temporary Decompositions in Macroeconomics and Finance, *The Economic Journal* 113, C125-C139.

Clark, P.K., 1987, The Cyclical Component of U.S. Economic Activity, *Quarterly Journal of Economics* 102, 797-814.

Friedman, M., 1964, Monetary studies of the National Bureau, the National Bureau enters its 45th Year, 44th Annual Report, 7-25 (NBER, New York); Reprinted in Friedman, M., 1969, The optimum quantity of money and other essays (Aldine, Chicago).

Friedman, M. 1993, The "plucking model" of business fluctuations revisited, *Economic Inquiry* 31, 171-177.

Hamilton, J.D., 1989, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57, 357-384.

Hodrick, R. J. and E. C. Prescott, 1997, Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit and Banking* 29, 1-16.

Kim, C.-J., 1994, Dynamic linear models with Markov-switching, *Journal of Econometrics* 60, 1-22.

Kim, C.-J., Morley, J. and J. Piger, 2003, Nonlinearity and the Permanent Effects of Recessions, working paper, Federal Reserve Bank of St. Louis.

Kim, C.-J. and C.R. Nelson, 1998, Business Cycle Turning Points, a New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model with Regime Switching, *The Review of Economics and Statistics* 80, 188-201.

Kim, C.-J. and C.R. Nelson, 1999a, Friedman's Plucking Model of Business Fluctuations: Tests and Estimates of Permanent and Transitory Components, *Journal of Money, Credit and Banking* 31, 317-34.

Kim, C.-J. and C.R. Nelson, 1999b, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, MIT Press, Cambridge and London.

Lam, P.-S., 1990, The Hamilton Model with a General Autoregressive Component: Estimation and Comparison with Other Models of Economic Time Series, *Journal of Monetary Economics* 26, 409-432.

Morley, J., 2002, A State-Space Approach to Calculating the Beveridge-Nelson Decomposition, *Economics Letters* 75, 123-127.

Morley, J.C., Nelson, C.R. and E. Zivot, 2003, Why are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP so Different? *The Review of Economics and Statistics* 85, 235-243.

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

Table 1
Maximum Likelihood Estimates for KMP Model of U.S. Real GDP

Parameter	*Estimate*	*Standard Error*
m_0	0.831	0.080
m_1	-2.005	0.242
l	0.319	0.059
q	0.956	0.018
p	0.679	0.111
f_1	0.138	0.082
f_2	0.076	0.082
s_e	0.764	0.044

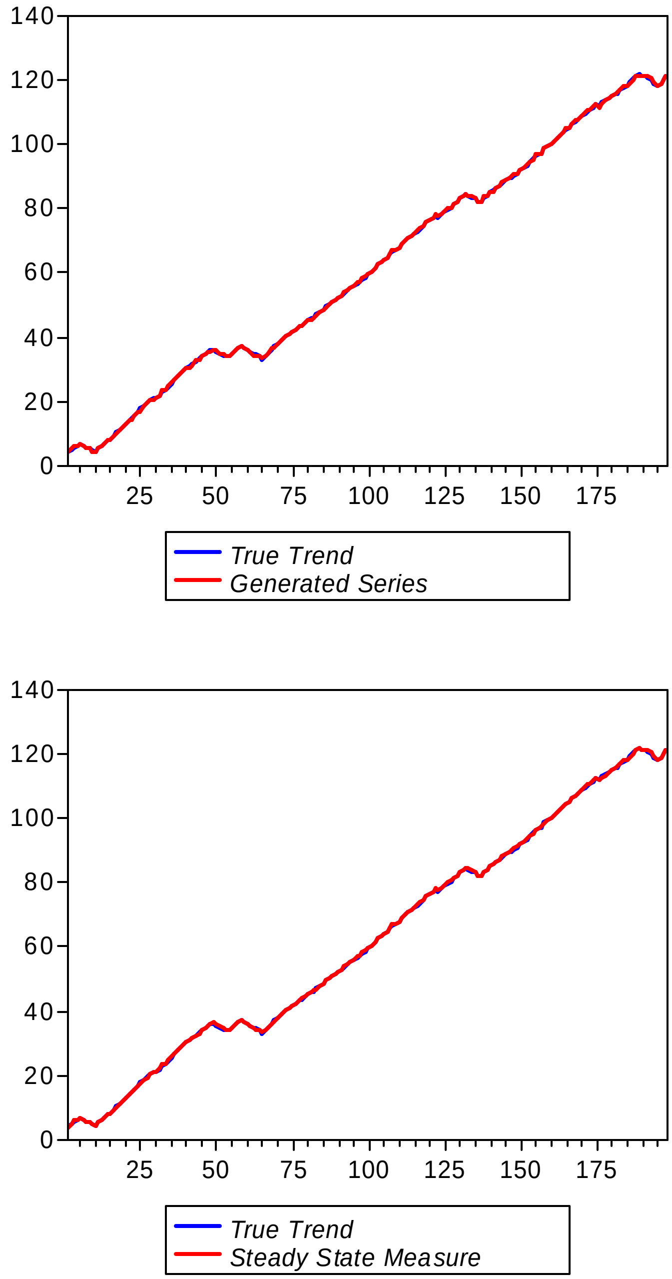

Fig. 1
Generated series, true trend, and estimates of trend for the Hamilton model

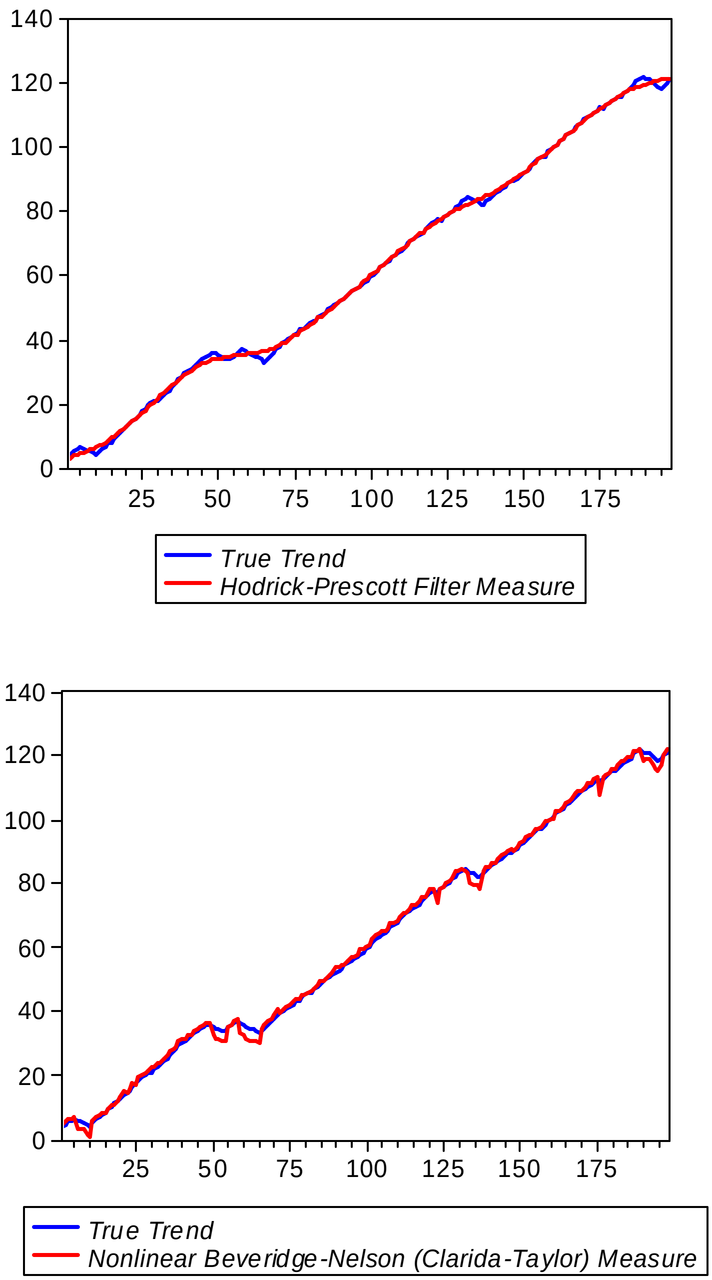

Fig. 1 (Continued)
Generated series, true trend, and estimates of trend for the Hamilton model

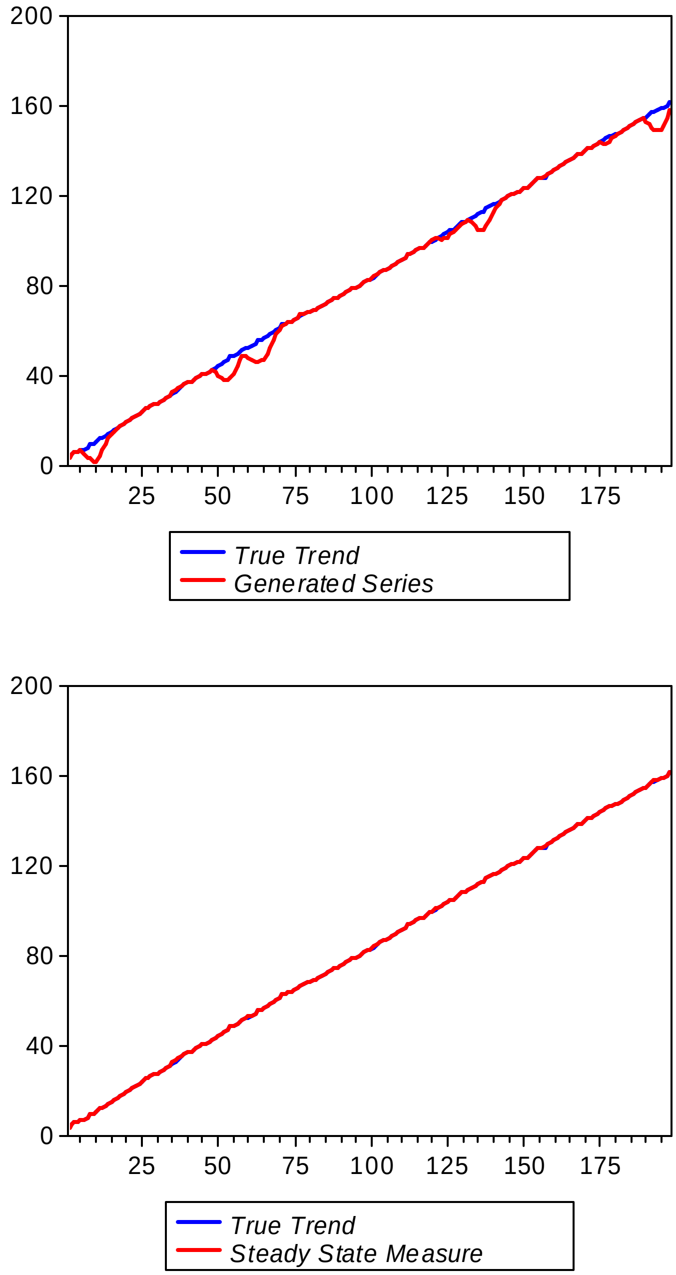

Fig. 2
Generated series, true trend, and estimates of trend for the plucking model

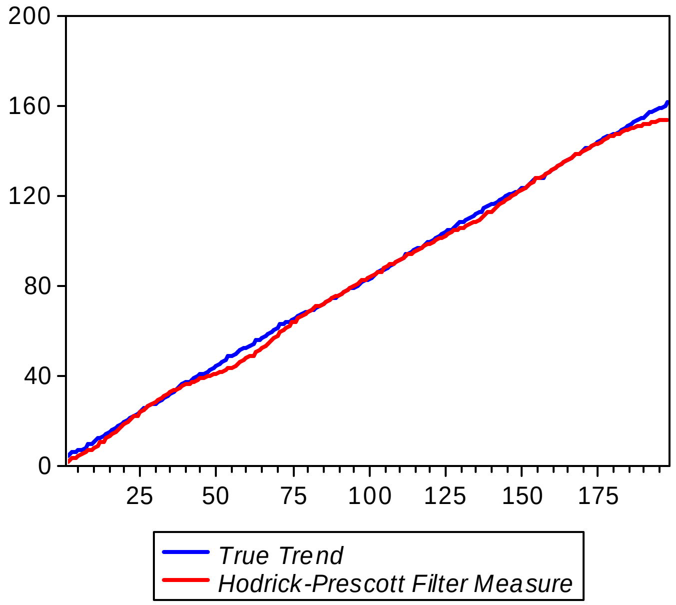

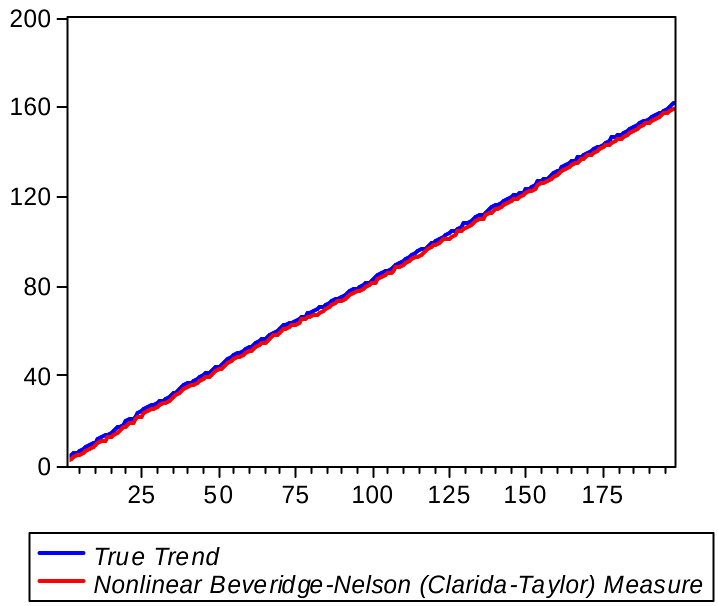

Fig. 2 (Continued)
Generated series, true trend, and estimates of trend for the plucking model

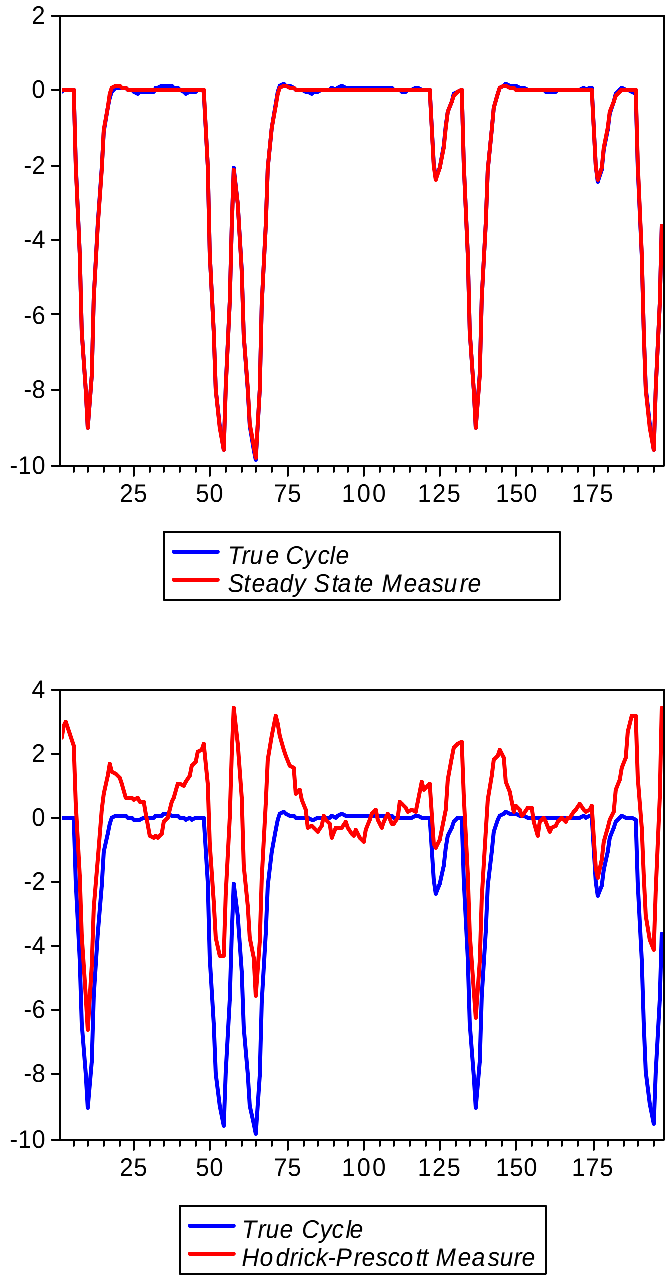

Fig. 3
True cycle and estimates of cycle for the plucking model

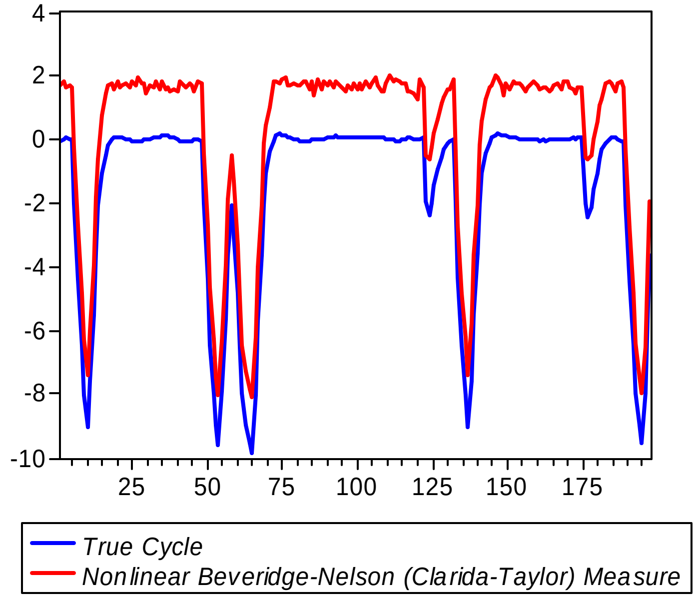

Fig. 3 (Continued)
True cycle and estimates of cycle for the plucking model

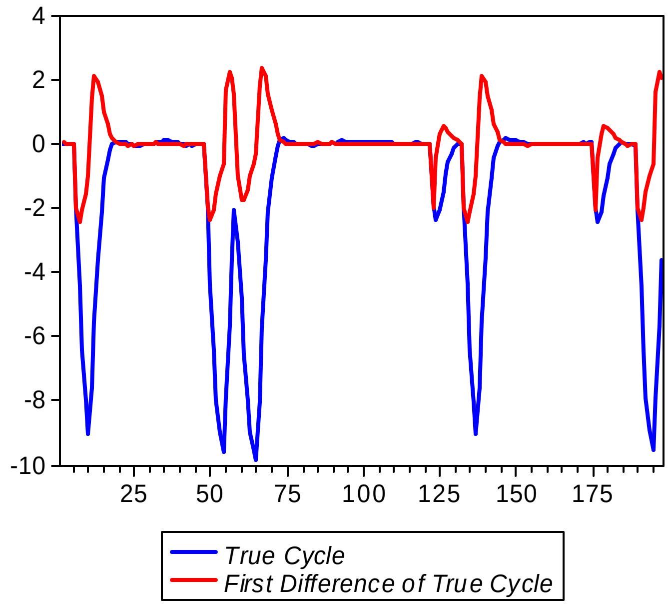

Fig. 4
The level and the first difference of the true cycle for the plucking model

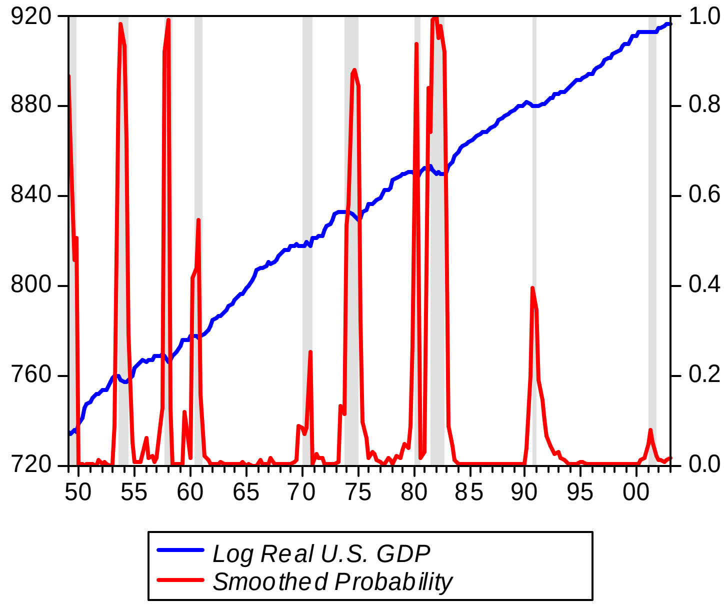

Fig. 5
U.S. real GDP and smoothed probability of a “contractionary” regime

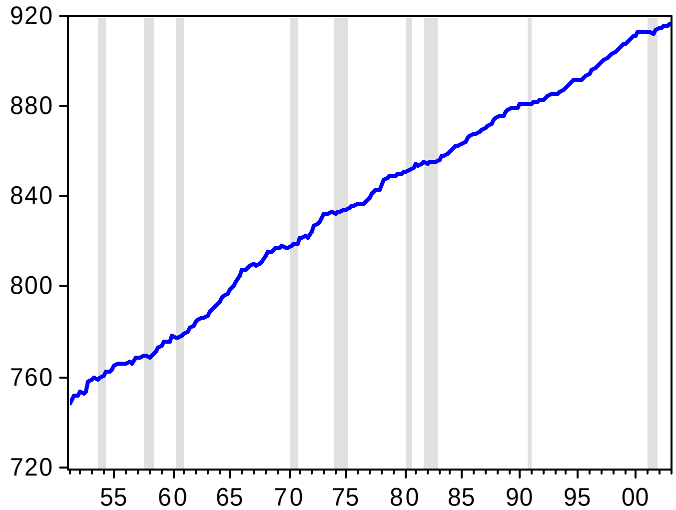

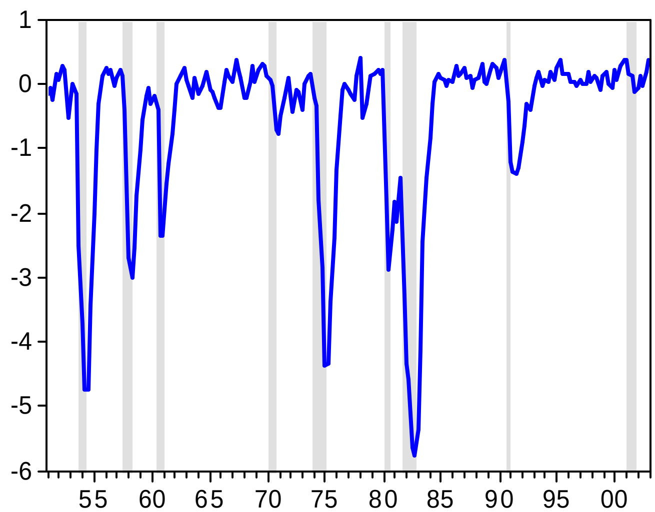

Fig. 6
Steady-state estimates of trend and cycle for U.S. real GDP implied by the KMP model (NBER recessions shaded)