Trend/cycle decomposition of regime-switching processes

James Morley [a,*], Jeremy Piger [b]

[a] *Washington University in St. Louis, United States*

[b] *University of Oregon, United States*

ABSTRACT

We present a new approach to trend/cycle decomposition of time series that follow regime-switching processes. The proposed approach, which we label the "regime-dependent steady-state" (RDSS) decomposition, is motivated as the appropriate generalization of the Beveridge and Nelson decomposition [Beveridge, S., Nelson, C.R., 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] to the setting where the reduced-form dynamics of a given series can be captured by a regime-switching forecasting model. For processes in which the underlying trend component follows a random walk with possibly regime-switching drift, the RDSS decomposition is optimal in a minimum mean-squared-error sense and is more broadly applicable than directly employing an Unobserved Components model.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Trend/cycle decomposition of integrated economic time series is important for both theoretical and statistical reasons. When an appropriate forecasting model is available, one general approach to trend/cycle decomposition is the method presented in Beveridge and Nelson (1981), or BN hereafter. The BN decomposition extracts a measure of trend from a given series using a long-horizon forecast based on the forecasting model, an approach that has been argued for in Rotemberg and Woodford (1996) and Cogley (2001), among others, and used in countless applications in the literature.

In recent years there has been an explosion of work using nonlinear forecasting models to describe economic time series. Beginning with Hamilton (1989), models with Markov-switching parameters have been particularly popular. While the BN decomposition can be calculated for such models using the techniques detailed in Clarida and Taylor (2003) and Kim (2006), the resulting measures of trend and cycle are often at odds with the underlying regime-switching processes. As one example, the BN decomposition always extracts a trend component with a constant average growth rate, while many regime-switching processes explicitly allow for shifts in the average growth rate across different regimes (e.g., Lam (1990)).

In this paper we present a new approach to trend/cycle decomposition of regime-switching processes, which we refer to as the regime-dependent steady-state (RDSS) decomposition. We motivate the RDSS decomposition as the appropriate generalization of the BN decomposition to the setting where the reduced-form dynamics of a given series can be captured by a regime-switching forecasting model. In particular, our approach is based on the central premise of the BN decomposition, namely that a long-horizon forecast can be used to eliminate predictable momentum implied by the cyclical component of an integrated series and thus extract a measure of its trend. However, we show that the long-horizon forecast should be constructed under the assumption that the series remains inside of a particular regime (hence the label "regime dependent"), rather than averaging across all regimes, as is done with the BN decomposition. Meanwhile, for linear forecasting models in which everything collapses to one regime, the RDSS and BN decompositions are equivalent.

It is useful to compare the RDSS and BN decompositions to direct estimation of trend and cycle using an Unobserved Components (UC) model. As shown in Morley et al. (2003), Kalman filter estimates for linear Gaussian UC models with random walk trends are equivalent to BN measures of trend and cycle based on corresponding reduced-form forecasting models. In this case, because the Kalman filter estimates and BN measures are equal to conditional expectations of the underlying trend and cycle, they provide optimal estimates in a "minimum mean-squared-error" sense. However, the equivalence between the UC approach and the BN decomposition breaks down for UC models with regime-switching parameters, of which there are many examples in the

* Corresponding address: Department of Economics, Box 1208, Eliot Hall 205, Washington University, One Brookings Drive, St. Louis, MO 63130-4899, United States. Tel.: +1 314 935 4437; fax: +1 314 935 4156.

E-mail address: morley@wustl.edu (J. Morley).

literature (e.g., Lam (1990) and Kim and Nelson (1999)). In particular, while filtered estimates based for the UC models are equal to conditional expectations of the underlying trend and cycle, the BN measures of trend and cycle based on the corresponding reduced-form regime-switching forecasting models are biased in the presence of a regime-switching drift for the underlying trend and/or a non-zero mean for the underlying cycle. By contrast, as long as the underlying trend follows a random walk with a possibly regime-switching drift, the RDSS decomposition provides measures of trend and cycle that are equivalent to estimates from the UC approach and are, therefore, optimal. Meanwhile, the RDSS measures are equal to conditional expectations of the underlying trend and cycle for a broader range of regime-switching processes, not just those that correspond to identified or finite-order UC models.

The rest of the paper is organized as follows. The next section begins by describing the regime-switching processes for which our approach is appropriate and discusses the problems with the BN decomposition for such processes. We then present the details of the RDSS decomposition, compare it to the UC approach, and address issues surrounding the identification of the mean of the cycle. Section 3 provides some brief conclusions.

2. Method

2.1. Underlying regime-switching processes

To explain our proposed approach to trend/cycle decomposition, it is most useful to start with the true data generating processes for which it would be appropriate. We will discuss regime-switching forecasting models that can capture such processes later. In terms of possible underlying processes, our approach is based on the assumption that a time series of interest $\{y_t\}_{t=-\infty}^{\infty}$ is the sum of a trend, τ_t, and a cycle, c_t:

$$y_t = \tau_t + c_t. \tag{1}$$

We are interested in the case where the parameters governing the evolution of both the trend and the cycle can take on different values in N distinct regimes, with the regimes indexed by a discrete state variable, denoted $S_t = i, i = 1, \ldots, N$. For simplicity, we assume that S_t follows an unobserved Markov process with a fixed transition matrix, an assumption that is consistent with much of the applied literature.[1]

We define the trend as the permanent component of the time series process, which is given by the accumulation of permanent innovations:

$$\tau_t = \sum_{j=0}^{\infty} \eta_{t-j}^*, \tag{2}$$

where the mean of the permanent innovations is possibly regime switching, so that

$$\eta_t^* = \bar{\eta}_{S_t} + \eta_t, \quad \eta_t \sim i.i.d.(0, \sigma_\eta^2). \tag{3}$$

These innovations are "permanent" in the sense that their impact on the level of the time series is not expected to be reversed.

Correspondingly, the trend process in (2) and (3) can be thought of as a random walk with regime-switching drift component—i.e., $\tau_t = \bar{\eta}_{S_t} + \tau_{t-1} + \eta_t$.

The cycle is then the transitory component of the time series process, which represents a weighted average of past transitory innovations:

$$c_t = \sum_{j=0}^{\infty} \psi_{j,t} \omega_{t-j}^*, \tag{4}$$

where $\psi_{0,t} = 1$. There are two possible sources of regime switching in the cycle. First, as was the case with permanent innovations, the mean of the transitory innovations can be regime switching:

$$\omega_t^* = \bar{\omega}_{S_t} + \omega_t, \quad \omega_t \sim i.i.d.(0, \sigma_\omega^2). \tag{5}$$

Second, the moving average (MA) coefficients in (4) can depend on the current and past regimes:

$$\psi_{j,t} = \psi_j(S_t, S_{t-1}, \ldots). \tag{6}$$

Meanwhile, to give the notion of "transitory innovations" meaning, we assume that the MA coefficients in (6) are always absolutely summable, $\sum_{j=0}^{\infty} |\psi_{j,t}| < \infty$, and their dependence on past regimes is described by a short-memory process. Thus, conditional on remaining inside of a regime, the cyclical component becomes a covariance-stationary and ergodic process with a regime-dependent mean, $E[c_t | \{S_t = i\}_{-\infty}^{\infty}]$. Finally, in at least one regime, which we label regime i^*, we assume that $\bar{\omega}_{i^*} = 0$. This final assumption identifies the unconditional mean of the cycle, and is discussed in further detail in Section 2.5.

There are two aspects of this regime-switching process that are worth highlighting. First, from (2) and (3), it is apparent that the trend depends only on current and past shocks and regimes and is, therefore, uniquely determined at time t. In other words, future regimes contain no additional information about the current trend above and beyond current and past regimes and information available at time t. Second, given the presence of $\bar{\omega}_{S_t}$ in (5), the cycle is not necessarily unconditionally mean zero.

It should be emphasized that the assumptions in (3), (5) and (6) are quite general. First, the innovations to the trend and cycle can be correlated due to common regime switching and/or correlation between the i.i.d. shocks:

$$\text{Cov}(\eta_t, \omega_t) = \sigma_{\eta\omega}. \tag{7}$$

Second, the overall process can be regime switching for a variety of reasons. In particular, the regime switching can be in terms of the dynamics, the permanent innovation, the transitory innovation, or any combination of these.[2]

Despite this general setup, our proposed approach to identifying and estimating the trend and cycle for such regime-switching processes requires no prior assumptions about the parametric structure of the cyclical component, the correlation between permanent and transitory innovations, or which sources of regime switching apply for a given time series of interest. Instead, our approach requires only the specification of a forecasting model for the first differences, Δy_t, of the integrated time series and is optimal in a minimum mean-squared-error sense given any

---

[1] The assumption of a fixed transition matrix could be relaxed to allow for time-varying transition probabilities based on exogenous variables, as in Diebold et al. (1994) and Filardo (1994). Also, it would be straightforward to consider an observable state process based on exogenous variables, such as in the trivial case of a dummy variable. However, processes in which the regime depends on realized values of endogenous variables, such as so-called "self-exciting threshold autoregressive" models (e.g., Potter (1995)), introduce some complications in terms of the evaluation of regime-dependent long-horizon forecasts. Adapting the basic approach developed in this paper to such processes would provide an interesting extension that we leave for future work.

[2] It is also possible to allow the variances and covariance of the permanent and transitory innovations to be regime switching. Thus, the shocks do not have to be identically distributed over time and their distributions can depend on the regime. However, they must be martingale difference sequences. We ignore this extension for simplicity of presentation.

such model that captures the reduced-form dynamics of the underlying process.[3] For example, suppose the trend component has a two-state regime-switching drift and the cyclical component follows a second-order autoregressive (AR(2)) process. It is straightforward to show that the reduced-form dynamics take the following second-order regime-switching autoregressive moving-average (ARMA(2, 2)) form: $(1 - \phi_1 L - \phi_2 L^2)\left(\Delta y_t - \mu_{S_t}\right) = (1 + \theta_1 L + \theta_2 L^2)e_t$, $e_t \sim i.i.d.(0, \sigma_e^2)$ and $\mu_{S_t} = \mu_1 \cdot I(S_t = 1) + \mu_2 \cdot I(S_t = 2)$, where $I(\cdot)$ is the indicator function that equals one if the argument is true and zero otherwise. Thus, we can directly capture the reduced-form process with a regime-switching ARMA(2, 2) forecasting model. Of course, in this example, we could also have considered an Unobserved Components (UC) model that captures the specified trend and cyclical components and estimated it directly. However, as discussed in more detail in Section 2.4, an identified or finite-order UC model is not always available for the underlying process presented in (1)–(7).

2.2. The Beveridge–Nelson decomposition

When we have an appropriate reduced-form forecasting model for Δy_t, a useful approach to identifying and estimating the trend and cycle of y_t is the decomposition suggested by Beveridge and Nelson (1981). The idea behind the BN decomposition is that, because the cyclical component is ergodic, a forecast of the time series into the infinite future will no longer be influenced by the expected cyclical momentum that exists at time t. Thus, such a forecast should reveal the influence of the permanent component on the time series, and can be used to estimate the trend.[4]

In precise terms, the BN measure of trend is the long-horizon forecast of a time series, adjusted to account for any future deterministic drift:

$$\hat{\tau}_t^{\text{BN}} \equiv \lim_{j \to \infty} \left\{ E^F\left[y_{t+j}|\Omega_t\right] - j \cdot E^F\left[\Delta y_t\right] \right\}, \tag{8}$$

where $E^F[\cdot]$ is the expectations operator with respect to the forecasting model and Ω_t is the set of relevant and available information observed up to time t. To see how the BN decomposition works, we can start with a simplified version of the process in (1)–(7) without regime-switching parameters, so that $\bar{\eta}_i = \bar{\eta}$ and $\bar{\omega}_i = 0$, $\forall i$. Letting $E[\cdot]$ specifically denote the expectations operator with respect to the process, the BN measure of trend, $\hat{\tau}_t^{\text{BN}}$, will be equal to $E[\tau_t|\Omega_t]$ as long as the forecasting model captures the reduced-form dynamics of the process such that $E^F\left[y_{t+j}|\Omega_t\right] = E\left[y_{t+j}|\Omega_t\right]$. To see this, substitute $E[\cdot]$ for $E^F[\cdot]$ in (8) and note from (1) that $y_{t+j} = \tau_{t+j} + c_{t+j}$ in order to re-write the expression for the BN measure of trend as

$$\hat{\tau}_t^{\text{BN}} = \lim_{j \to \infty} \left\{ E\left[\tau_{t+j}|\Omega_t\right] - j \cdot E\left[\Delta y_t\right] \right\} + \lim_{j \to \infty} \left\{ E\left[c_{t+j}|\Omega_t\right] \right\}$$

$$= E\left[\tau_t|\Omega_t\right] + \lim_{j \to \infty} \left\{ \sum_{k=1}^{j} E\left[\Delta \tau_{t+k}|\Omega_t\right] - j \cdot E\left[\Delta y_t\right] \right\}$$

$$+ \lim_{j \to \infty} \left\{ E\left[c_{t+j}|\Omega_t\right] \right\}. \tag{9}$$

In the absence of regime switching, the trend component defined in (2) and (3) is simply a random walk with drift, so that, for $k > 0$, the conditional expectation $E[\Delta \tau_{t+k}|\Omega_t] = \bar{\eta}$. Also,

given the ergodicity of the cyclical component, it is straightforward to show that the unconditional expectation $E[\Delta y_t] = \bar{\eta}$. Finally, the ergodicity of the cyclical component along with the fact that $\bar{\omega}_i = 0$, $\forall i$, guarantees that the limiting conditional expectation $\lim_{j \to \infty}\left\{E\left[c_{t+j}|\Omega_t\right]\right\} = 0$. Taken together, these three expectations imply that, in the absence of regime-switching parameters, the last two terms in (9) drop out and the BN measure of trend is equal to the conditional expectation of the underlying trend component. See Watson (1986) and Morley et al. (2003) on this point.

By contrast, when the process for y_t has regime-switching parameters, the last two terms in (9) do not necessarily drop out and $\hat{\tau}_t^{\text{BN}}$ does not, in general, equal $E[\tau_t|\Omega_t]$, even if $E^F\left[y_{t+j}|\Omega_t\right] = E\left[y_{t+j}|\Omega_t\right]$. First, in the presence of regime switching in the drift parameter of the trend component, the conditional expectation $E[\Delta \tau_{t+k}|\Omega_t]$ no longer equals the unconditional expectation $E[\Delta y_t]$ for all $k > 0$. This is because information contained in Ω_t about current and past regimes is useful for predicting future regimes, and thus future changes in the trend. Thus, $\hat{\tau}_t^{\text{BN}}$ will be biased by expected *future* changes in the trend. Second, in the presence of regime switching, the ergodicity of c_t implies only that the limiting conditional expectation $\lim_{j \to \infty} E\left[c_{t+j}|\Omega_t\right]$ converges to the unconditional mean of the cyclical component, which may or may not be zero. Thus, $\hat{\tau}_t^{\text{BN}}$ will also be biased by the expected effects of future regime shifts on the level of the cyclical component.[5]

2.3. Measuring the trend with a regime-dependent steady state

While the BN decomposition generally fails to provide $E[\tau_t|\Omega_t]$ for the regime-switching processes described in Section 2.1, the principle of using a long-horizon forecast to eliminate the influence of the cyclical component can still be used to construct an alternative decomposition that will yield $E[\tau_t|\Omega_t]$. In this subsection we lay out the details of this alternative decomposition, which we refer to as the "regime-dependent steady-state" (RDSS) decomposition.

To begin, given an underlying process that corresponds to (1)–(7), the appropriate reduced-form forecasting model for Δy_t will involve current and, possibly, lagged values of the regime indicator variable. To develop our approach, we initially proceed as if the relevant regimes, given by the vector $\tilde{S}_t \equiv (S_t, S_{t-1}, \ldots, S_{t-m})'$, were observed in period t. Later, we address the fact that $\tilde{S}_t$ is unobserved by marginalizing inferences with respect to its distribution. For now, we condition on $\tilde{S}_t$ and on a particular future sequence of regimes in order to construct a hypothetical "regime-dependent" j-step forecast of y_{t+j}:

$$E^F\left[y_{t+j}| \left\{S_{t+k} = i^*\right\}_{k=1}^{j}, \tilde{S}_t, \Omega_t\right]. \tag{10}$$

In words, (10) is the period t forecast of y_{t+j} given the hypothetical knowledge that the state process will enter regime i^* in period $t+1$ and remain there through period $t + j$.

Now, define a regime-dependent long-horizon forecast as follows:

$$\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) \equiv \lim_{j \to \infty} \left\{ E^F\left[y_{t+j}| \left\{S_{t+k} = i^*\right\}_{k=1}^{j}, \tilde{S}_t, \Omega_t\right] \right.$$

$$\left. - j \cdot E^F\left[\Delta y_t| \left\{S_t = i^*\right\}_{-\infty}^{\infty}\right] \right\}. \tag{11}$$

---

[3] Of course, it should always be acknowledged that if the forecasting model provides a poor approximation to the reduced-form dynamics of the underlying process, the inferences about trend and cycle will also be poor.

[4] The long-horizon forecast can also be used to define the trend, although there are some issues with this interpretation of the BN decomposition (see Morley (2007)).

[5] In contemporaneous work, Chen and Tsay (2006) present a modification of the BN decomposition designed to incorporate Markov-switching in the average growth rate of the trend component. However, their approach does not address the possibility of regime switching in the cyclical component.

We can show that $\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = E\left[\tau_t | \tilde{S}_t, \Omega_t\right]$ as long as the forecasting model captures the reduced-form dynamics of the process such that $E^F\left[y_{t+j} | \Omega_t\right] = E\left[y_{t+j} | \Omega_t\right]$. Again, substituting $E\left[\cdot\right]$ for $E^F\left[\cdot\right]$ and noting $y_{t+j} = \tau_{t+j} + c_{t+j}$, the expression in (11) can be re-written as

$$\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = \lim_{j \to \infty}\left\{E\left[\tau_{t+j} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]\right.$$
$$\left. - j \cdot E\left[\Delta y_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right]\right\}$$
$$+ \lim_{j \to \infty} E\left[c_{t+j} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]. \tag{12}$$

From Section 2.1, the cyclical component is ergodic inside of each regime. Thus, the second limit on the right hand side of (12) converges to the unconditional mean of the cyclical component inside of regime i^*, which is zero given our assumption that $\bar{\omega}_{i^*} = 0$.[6] Thus, (12) simplifies to

$$\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = \lim_{j \to \infty}\left\{E\left[\tau_{t+j} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]\right.$$
$$\left. - j \cdot E\left[\Delta y_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right]\right\}. \tag{13}$$

The first term on the right-hand-side of (13) can be decomposed as follows:

$$E\left[\tau_{t+j} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right] = E\left[\tau_t \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]$$
$$+ \sum_{k=1}^j E\left[\Delta\tau_{t+k} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right], \tag{14}$$

where from the definition of the trend component in (2) and (3), we can solve for the last term in (14):

$$\sum_{k=1}^j E\left[\Delta\tau_{t+k} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]$$
$$= \sum_{k=1}^j E\left[\bar{\eta}_{S_{t+k}} + \eta_{t+k} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right] = j \cdot \bar{\eta}_{i^*}. \tag{15}$$

Then, from (1) to (6), the regime-dependent expectation in the second term on the right-hand-side of (13) can be decomposed as follows:

$$E\left[\Delta y_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right] = E\left[\Delta\tau_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right]$$
$$+ E\left[\Delta c_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right]$$
$$= E\left[\bar{\eta}_{S_t} + \eta_t \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right] + E\left[\sum_{j=0}^\infty \psi_j\left(S_t, S_{t-1}, \ldots\right)\omega_{t-j}^*\right.$$
$$\left. - \sum_{j=0}^\infty \psi_j\left(S_{t-1}, S_{t-2}, \ldots\right)\omega_{t-1-j}^* \left| \{S_t = i^*\}_{-\infty}^\infty\right.\right]$$
$$= \bar{\eta}_{i^*} + E\left[\sum_{j=0}^\infty \psi_j\left(i^*, i^*, \ldots\right)\bar{\omega}_{i^*} - \sum_{j=0}^\infty \psi_j\left(i^*, i^*, \ldots\right)\bar{\omega}_{i^*}\right]$$
$$= \bar{\eta}_{i^*}. \tag{16}$$

Thus, substituting (14)–(16) into (13) yields

$$\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = E\left[\tau_t \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t\right.\right]. \tag{17}$$

Finally, the Markov property of the state variable ensures that the sequence of future regimes, $\{S_{t+k} = i^*\}_{k=1}^j$, provides no additional information about the current trend beyond the other conditioning information in (17). Specifically, from (2) and (3), τ_t only depends on current and lagged regimes. Therefore, the only relevance of knowing the future regimes in (17) would be to provide information about current and past regimes. However, because the state variable follows a Markov process, the probability distribution $p\left(S_{t-1} | S_t, S_{t+1}\right)$ simplifies to $p\left(S_{t-1} | S_t\right)$. Thus, given the knowledge of S_t, which is included in the vector $\tilde{S}_t$, the future regimes in $\{S_{t+k} = i^*\}_{k=1}^j$ provide no additional information about current and past regimes. As a result, (17) simplifies to

$$\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = E\left[\tau_t | \tilde{S}_t, \Omega_t\right]. \tag{18}$$

In practice, the relevant current and past regimes, $\tilde{S}_t$, are not observed, but need to be integrated out of the right hand side of (18) to arrive at $E\left[\tau_t | \Omega_t\right]$. This can be done using the following weighted sum:

$$E\left[\tau_t | \Omega_t\right] = \sum_{\tilde{S}_t} E\left[\tau_t | \tilde{S}_t, \Omega_t\right] \cdot p\left(\tilde{S}_t | \Omega_t\right), \tag{19}$$

where $p\left(\tilde{S}_t | \Omega_t\right)$ represents the probability distribution of each possible sequence of relevant current and past regimes. Thus, in general, the RDSS measure of trend is

$$\hat{\tau}_t^{\text{RDSS}} \equiv \sum_{\tilde{S}_t}\left\{\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) \cdot p^F\left(\tilde{S}_t | \Omega_t\right)\right\}, \tag{20}$$

where $p^F(\cdot)$ is the probability distribution with respect to the forecasting model and the above analysis, demonstrating that $\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right) = E\left[\tau_t | \tilde{S}_t, \Omega_t\right]$ directly implies that $\hat{\tau}_t^{\text{RDSS}} = E\left[\tau_t | \Omega_t\right]$, as long as the forecasting model captures the state variable dynamics such that $p^F\left(\tilde{S}_t | \Omega_t\right) = p\left(\tilde{S}_t | \Omega_t\right)$. In this case, it can also be easily shown that $\hat{c}_t^{\text{RDSS}} = E\left[c_t | \Omega_t\right]$, where $\hat{c}_t^{\text{RDSS}} \equiv y_t - \hat{\tau}_t^{\text{RDSS}}$ is the RDSS estimate of the cycle.

It is worth emphasizing the computational simplicity of the RDSS decomposition. In particular, for each possible $\tilde{S}_t$, one need only compute the quantity in (11), $\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right)$, which is simply a long-horizon forecast of the time series conditional on $\tilde{S}_t$, Ω_t, and $\{S_{t+k} = i^*\}_{k=1}^\infty$, adjusted to subtract the expected average growth of the time series in regime i^* that accumulates over the forecast horizon. Given the assumption that $\bar{\omega}_{i^*} = 0$, the calculation of this adjusted long-horizon forecast is mechanically identical to the BN decomposition for the conditionally-linear forecasting model that would be implied by a sequence of known regimes.[7] Then, to arrive at a measure that is equivalent to $E\left[\tau_t | \Omega_t\right]$, the values of $\hat{\tau}_t^{\text{RDSS}}\left(\tilde{S}_t\right)$ can be averaged, as in (20), using the estimated state probabilities, $p^F\left(\tilde{S}_t | \Omega_t\right)$. In most cases, these probabilities are a direct by-product of the filter used to estimate the parameters for a regime-switching forecasting model of Δy_t. For convenience, Box I summarizes the mechanics of the RDSS decomposition.

---

[6] It is this convergence that gives the RDSS decomposition its name. In particular, the ergodicity of the cyclical component inside of each regime ensures that the regime-dependent long-horizon forecast of the series will no longer be influenced by the expected cyclical dynamics, and is thus in a "steady-state" in which expected growth in the series will be determined entirely by expected growth in the trend.

[7] In the case of a Gaussian disturbance term for the forecasting model, this BN decomposition can be easily computed using the analytical approach presented in Morley (2002).

$$\hat{\tau}_t^{\text{RDSS}} \equiv \sum_{\tilde{S}_t} \left\{ \hat{\tau}_t^{\text{RDSS}} \left(\tilde{S}_t \right) \cdot p^F \left(\tilde{S}_t | \Omega_t \right) \right\},$$

$$\text{where } \hat{\tau}_t^{\text{RDSS}} \left(\tilde{S}_t \right) \equiv \lim_{j \to \infty} \left\{ E^F \left[y_{t+j} | \left\{ S_{t+k} = i^* \right\}_{k=1}^{j}, \tilde{S}_t, \Omega_t \right] - j \cdot E^F \left[\Delta y_t | \left\{ S_t = i^* \right\}_{-\infty}^{\infty} \right] \right\}$$

Box I. The RDSS Measure of Trend.

We conclude this subsection by noting that the RDSS decomposition is equivalent to the BN decomposition when the forecasting model for Δy_t does not involve regime switching. In particular, in the absence of regime switching, the regime indicator variable S_t becomes redundant information and is no longer needed. In this case, the regime-dependent long-horizon forecast in (11) collapses to the expression in (8) for the BN measure of trend.

2.4. Comparison to the unobserved components approach

A popular approach to extracting estimates of the trend and cycle from an integrated time series y_t is to specify functional forms for these components directly in a UC model (e.g. Harvey (1985), Watson (1986) and Clark (1987)). For linear Gaussian UC models, the Kalman filter can then be used to calculate $E^{UC}[\tau_t | \Omega_t]$, where $E^{UC}[\cdot]$ is the expectations operator with respect to the UC model. Morley et al. (2003) show that if the trend component follows a random walk with constant drift, the cyclical component follows a finite-order ARMA process, and the UC model is identified and appropriate such that $E^{UC}[\tau_t | \Omega_t] = E[\tau_t | \Omega_t]$, then the BN decomposition for the corresponding reduced-form ARMA forecasting model of Δy_t will also yield $E^F[\tau_t | \Omega_t] = E[\tau_t | \Omega_t]$, implying that the UC approach and BN decomposition are equivalent and optimal in this case.

A similar equivalence exists for the UC approach and the RDSS decomposition when the UC model has regime-switching parameters. In particular, several authors, notably Lam (1990) and Kim and Nelson (1999), have specified UC models for y_t in which the trend and/or cycle are regime switching. For such models, posterior densities for the trend and cycle will depend on the entire history of regimes. Thus, even given Gaussian shocks, exact analytical inference based on the Kalman filter is computationally infeasible due to the need to keep track of up to N^T possible sequences of regimes, where T is the sample size under consideration. However, it is possible to use the Gibbs sampler to compute $E^{UC}[\tau_t | \Omega_t]$ numerically (Carter and Kohn, 1994). If the trend component in the UC model is specified as a random walk with potentially regime-switching drift, as in (2)–(3), the cyclical component follows a potentially regime-switching finite-order ARMA process, and the UC model is identified and appropriate such that $E^{UC}[\tau_t | \Omega_t] = E[\tau_t | \Omega_t]$, then the RDSS decomposition for the corresponding reduced-form regime-switching ARMA forecasting model of Δy_t will also yield $E^F[\tau_t | \Omega_t] = E[\tau_t | \Omega_t]$, and thus the UC approach and RDSS decomposition are equivalent and optimal. By contrast, because the BN decomposition does not generally produce $E[\tau_t | \Omega_t]$ for a regime-switching process, as discussed in Section 2.2, it is not equivalent to the UC approach, even given the appropriate reduced-form forecasting model.[8]

Given that the RDSS decomposition produces identical results to the UC approach when an identified and appropriate UC model is available, a reasonable question is whether there is any value added from using the RDSS decomposition. Beyond computational

simplicity, the primary benefit of the RDSS decomposition is that it is appropriate for a broader range of regime-switching processes than the UC approach. First, there are some relevant processes that correspond to (1)–(7) for which UC models are not identified. In this case, the RDSS decomposition will be robust in the sense that the RDSS decomposition will provide optimal estimates of trend and cycle when applied to forecasting models that capture the reduced-form dynamics of the set of underlying unidentified processes. Second, there are also relevant processes that correspond to (1)–(7) for which the reduced-form dynamics can be captured by flexible, yet tightly parameterized regime-switching forecasting models, while the dynamics cannot be captured by finite-order UC models.

To illustrate the robustness of the RDSS decomposition, consider simple "random-walk-plus-noise" UC models for y_t, which correspond to MA(1) reduced-form dynamics for Δy_t. These UC models are unidentified in the sense that the correlation between the shocks to trend and cycle cannot be estimated. In such a setting, a standard approach is to consider only one of the possible UC models by making an assumption about the correlation parameter (usually that it is zero). However, such an assumption can place false restrictions on the parameters related to the autocovariance structure of the process (see Nelson and Plosser (1982) and Morley et al. (2003), on this point). Instead, by directly estimating the reduced-form MA(1) forecasting model and applying the RDSS decomposition (which would just be the BN decomposition given the linear MA(1) model in this simple example), the parameter estimates will be consistent and the inferences about trend and cycle will be optimal for the range of unidentified random-walk-plus-noise processes. In this sense, the RDSS decomposition provides a robust approach to estimation of the trend and cycle when UC parameters are not identified.

To illustrate the relative flexibility of the RDSS decomposition compared to the UC approach, note that identified UC models correspond to reduced-form models for Δy_t that include MA dynamics, rather than just AR dynamics only. Thus, unlike the RDSS decomposition, the UC approach implicitly limits the reduced-form dynamics that can be considered to the case of more complicated ARMA models.[9] Restricting the class of reduced-form models in this way is undesirable for at least two reasons. First, many popular regime-switching models in the literature include only AR dynamics (e.g. Hamilton (1989), Hansen (1992) and Garcia and Perron (1996)). Ruling out such models from consideration is problematic if their popularity reflects, at least in part, their empirical relevance. Second, UC models and their corresponding reduced-form ARMA models are often difficult to estimate and may not be good forecasting models in practice. Estimation difficulties arise due to weak identification in the presence of near cancellation

---

[8] As implied by the analysis in Section 2.2, the BN decomposition for a regime-switching forecasting model yields the conditional expectation of the trend only in the special case that, for the corresponding UC model, the trend is a random walk with constant drift and the cycle is unconditionally mean zero.

[9] It might seem counterintuitive that an ARMA structure does not nest the simpler AR structure. However, given an identified UC model, the MA parameter for the corresponding reduced-form model is always nonzero, except in the limiting case of no cyclical component. Of course, in this limiting case, making inference about the trend is as simple as observing the series. Meanwhile, a reduced-form AR model provides a parsimonious way to capture the infinite-order MA structure for the first-differences of an integrated time series that would be implied by a cyclical component with a general infinite-order MA structure.

of roots in the AR and MA polynomials.[10] Forecasting failures occur because, as pointed out by Campbell and Mankiw (1987), estimated ARMA models (and, implicitly, UC models) sometimes massively overstate the long-horizon predictability inherent in an integrated process due to a pile-up problem for the likelihood at a unit MA root. This is a serious drawback for such models, as the main criterion for a model to be appropriate for use in a model-based decomposition is whether it credibly captures the actual predictability inherent in a given integrated time series process.

Another advantage of the RDSS decomposition over the UC approach that is worth mentioning is the ease with which non-Gaussian shocks can be accommodated. In particular, the parameters of regime-switching AR forecasting models with non-Gaussian errors are relatively straightforward to estimate, as the likelihood function can be computed with minor modifications to existing recursive filters. Long-horizon forecasts are also relatively straightforward to calculate for such models. By contrast, both the construction of the likelihood function and the calculation of $E^{UC}[\tau_t | \Omega_t]$ for UC models with non-Gaussian shocks (regardless of the presence of regime-switching) present substantial computational challenges, requiring the use of nonlinear filtering techniques (e.g. Kitagawa (1987) and Gordon et al. (1993)). This is relevant, as there is evidence for non-Gaussian shocks in many macroeconomic and financial time series (e.g. Hamilton (2005)).

2.5. Identifying the mean of the cycle

The RDSS decomposition relies on the assumption that $\bar{\omega}_{i^*} = 0$, which implies that the regime-dependent mean of the cycle is zero in regime i^*, where the regime is specified by the researcher. This assumption was necessary to identify the unconditional mean of the cycle. To see this, note that, while the long-horizon forecast in (11) guarantees that the cyclical component has converged to its regime-dependent mean, it does not identify the level of this mean. Also, note that the *differences* in regime-dependent means of the cycle are identified by (11). In particular, suppose we consider an alternative version of (11), where instead of assuming that the process entered and remained in regime i^* beginning in period $t + 1$, we assume that the process entered and remained in a different regime, labeled j^*. It is straightforward to show from (12) that the difference between (11) with regime i^* and its alternative version with regime j^* simplifies to

$$\lim_{j \to \infty} \left\{ E\left[c_{t+j} \left| \{S_{t+k} = i^*\}_{k=1}^j, \tilde{S}_t, \Omega_t \right. \right] \right\}$$
$$- \lim_{j \to \infty} \left\{ E\left[c_{t+j} \left| \{S_{t+k} = j^*\}_{k=1}^j, \tilde{S}_t, \Omega_t \right. \right] \right\}, \tag{21}$$

which is simply the difference in the mean of the cycle in regime i^* from the mean of the cycle in regime j^*. Thus, this difference is readily calculable using (11) for different regimes, meaning that, if the mean of the cycle in regime i^* is zero, it is possible to identify the regime-dependent mean of the cycle in all other regimes, and thus the unconditional mean of the cycle.

Of course, this identification requires the choice by the researcher of which regime to label i^*, and thus in which regime the cyclical component has a zero mean. It is worth noting that this choice is also required in the UC approach to trend/cycle decomposition for regime-switching processes. In particular, UC models require explicit functional forms for the cyclical component in each regime, with the overall mean of the cycle only identified by fixing the mean of the cyclical component in one of the regimes.

It should be stressed that the choice of i^* need not be completely arbitrary (i.e., it is not merely a matter of normalization). While the decomposition method itself provides no guide as to how this choice should be made and the estimated variation in trend and cycle are robust to the choice, there may be compelling reasons in any particular application to choose one regime over another. For example, for UC models of the business cycle, the usual practice is to assume that the mean of the cycle in the "normal" regime, defined as the most frequently occurring, is zero (e.g. Kim and Nelson (1999)). This choice is driven by the argument that the large, asymmetric shocks that yield non-zero mean cycles are abnormal events. In certain cases, it might also be possible to use subsequent analysis with the different measures of the cycle obtained from different choices of i^* in order to discriminate amongst them.[11]

Finally, and in contrast to the RDSS decomposition and the UC approach, the BN decomposition identifies the mean of the cycle by assuming that it is zero. Thus, the BN decomposition ties the hands of the researcher wanting to extract the trend and cycle from a regime-switching process. In particular, even if the researcher has *a priori* reasons for allowing a non-zero mean cycle, the BN decomposition does not permit it.

3. Conclusion

We have developed a new approach to trend/cycle decomposition of time series that follow regime-switching processes. Because of its mechanics, we refer to this new approach as the "regime-dependent steady-state" (RDSS) decomposition. The RDSS decomposition is useful because it provides optimal estimates of trend and cycle for a broader class of regime-switching processes than either the Beveridge–Nelson decomposition or the Unobserved Components approach. In future research, we plan to apply the RDSS decomposition to study the US business cycle using a range of linear and nonlinear forecasting models.

References

Beveridge, S., Nelson, C.R., 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.

Campbell, J., Mankiw, G., 1987. Are output fluctuations transitory? Quarterly Journal of Economics 102, 857–880.

---

[10] See Nelson and Startz (2007) on weak identification of ARMA models.

[11] For example, Keynesian macroeconomic theory suggests that inflation will have a tendency to increase or decrease depending on whether the economy is above or below trend. In this case, the plausibility of alternative measures of the cyclical component of economic activity could be evaluated based on their ability to explain the direction of change in inflation.

Carter, C.K., Kohn, P., 1994. On Gibbs sampling for state space models. Biometrika 81, 541–553.

Chen, C.-C., Tsay, W.-J., 2006. The Beveridge–Nelson decomposition of Markov-switching processes. Economics Letters 91, 83–89.

Clarida, R.H., Taylor, M.P., 2003. Nonlinear permanent–temporary decompositions in macroeconomics and finance. The Economic Journal 113, C125–C139.

Clark, P.K., 1987. The cyclical component of US economic activity. Quarterly Journal of Economics 102, 797–814.

Cogley, T., 2001. Alternative definitions of the business cycle and their implications for business cycle models: A reply to Torben Mark Pederson. Journal of Economic Dynamics and Control 25, 1103–1107.

Diebold, F., Lee, J.-H., Weinbach, G., 1994. Regime switching with time-varying transition probabilities. In: Hargreaves, C. (Ed.), Non-Stationary Time Series Analysis and Cointegration. Oxford University Press, Oxford, UK.

Filardo, A.J., 1994. Business-cycle phases and their transitional dynamics. Journal of Business and Economic Statistics 12, 299–308.

Garcia, R., Perron, P., 1996. An analysis of the real interest rate under regime shifts. Review of Economics and Statistics 78, 111–125.

Gordon, N.J., Salmond, D.J., Smith, A.F.M., 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings-F 140, 107–113.

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

Hamilton, J.D., 2005. What's real about the business cycle? Federal Reserve Bank of St. Louis Review 87, 435–452.

Hansen, B.E., 1992. The likelihood ratio test under nonstandard conditions: Testing the Markov switching model of GNP. Journal of Applied Econometrics 7, S61–S82.

Harvey, A.C., 1985. Trends and cycles in macroeconomic time series. Journal of Business and Economic Statistics 3, 216–227.

Kim, C.-J., 2006. The Beveridge–Nelson decomposition and impulse-response analysis in the presence of Markov-switching: Has the persistence of real GDP changed since the mid-1980s? Working Paper.

Kim, C.-J., Nelson, C.R., 1999. Friedman's plucking model of business fluctuations: Tests and estimates of permanent and transitory components. Journal of Money, Credit and Banking 31, 317–334.

Kitagawa, G., 1987. Non-Gaussian state-space modeling of nonstationary time series. Journal of the American Statistical Association 82, 1032–1063.

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., 2007. The two interpretations of the Beveridge–Nelson decomposition. Working Paper.

Morley, J.C., Nelson, C.R., Zivot, E., 2003. Why are the Beveridge–Nelson and unobserved-components decompositions of GDP so different? The Review of Economics and Statistics 85, 235–243.

Nelson, C.R., Plosser, C.I., 1982. Trends and random walks in macroeconomic time series: Some evidence and implications. Journal of Monetary Economics 10, 139–162.

Nelson, C.R., Startz, R., 2007. The zero-information-limit condition and spurious inference in weakly identified models. Journal of Econometrics 138, 47–62.

Potter, S., 1995. A nonlinear approach to US GNP. Journal of Applied Econometrics 10, 109–125.

Rotemberg, J.J., Woodford, M., 1996. Real-business-cycle models and the forecastable movements in output, hours, and consumption. American Economic Review 86, 71–89.

Watson, M.W., 1986. Univariate detrending methods with stochastic trends. Journal of Monetary Economics 18, 49–75.