# Exploiting Non-Linearities in GDP Growth for Forecasting and Anticipating Turning Points

David N. DeJong
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260, USA

Hariharan Dharmarajan
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260, USA

Roman Liesenfeld
Department of Economics
Universität Kiel
24118 Kiel, Germany

Jean-François Richard
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260, USA

September 2008

**Abstract**

We exploit non-linearities in the behavior of GDP growth to produce forecasts of growth itself, and to anticipate regime changes and NBER-dated turning points. The forecasting model we develop features an error-correction mechanism with a drift component that follows regime-specific trajectories. Success in forecasting depends critically on the incorporation of a flexible specification that admits wide heterogeneity in trajectories across regimes, including potentially dramatic departures from linearity.

Keywords: regime switching; error correction mechanism

JEL Codes: C22, C51, C52

# 1 Introduction

Success in forecasting GDP growth depends critically on the ability to anticipate shifts of the economy between episodes of general expansion and contraction. The apparent non-linearity of the process that governs these shifts, coupled with the decrease in the volatility of GDP growth observed since the mid-1980s, renders the task of anticipating regime shifts as challenging. Here we develop a univariate non-linear model designed to detect and anticipate regime shifts. The model exhibits strong in-sample performance in forecasting GDP growth, and is also useful in anticipating NBER turning points.

The model is a refinement of that developed by DeJong, Liesenfeld and Richard (DLR, 2005), and features two key components: regime-switching behavior, and an error-correction mechanism (ECM). The ECM characterizes stochastic fluctuations of GDP growth around a drift component. In turn, regimes are defined in terms of the behavior of the drift component, which shifts stochastically between expansionary and contractionary trajectories (thus GDP growth shifts between phases of general acceleration and deceleration). Trajectories are non-linear and heterogeneous across regimes, with specific characteristics determined by the realization of a set of latent parameters drawn from a fixed distribution. Compared with the specification of DLR, our characterization of drift trajectories is highly flexible, and capable of exhibiting dramatic departures from linearity.

Our use of a regime-change specification builds upon Hamilton (1989) and Tong (1990). Models constructed following Hamilton characterize regime changes as being governed by unobserved regime indicators (extensions of Hamilton's two-regime specification include Boldin, 1996; Clements and Krolzig, 1998; and Kim, Nelson and Piger, 2004). Models constructed following Tong characterize regime changes as being governed by observed indicators constructed as deterministic functions of current and past GDP growth (extensions of Tong's specification include Beaudry and Koop, 1993; Pesaran and Potter, 1997; van Dijk and Franses, 1999; and Ocal and Osborn, 2000).

A third class of models incorporates features in the tradition of both Hamilton and Tong by modelling regime changes as stochastic, with transition probabilities dependent upon observed indicator variables (examples include Durland and McCurdy, 1994; Filardo and Gordon, 1998; and DLR). Like DLR, we link regime-change probabilities to an observed indicator variable via a logistic transformation. We refer to the indicator variable as a tension index, which is constructed as the geometric sum of past deviations of GDP growth from a corresponding sustainable growth rate (interpreted as the growth rate of potential GDP). The tension index tends to increase in growth-acceleration regimes, and decrease in deceleration regimes; in either case, as the index increases in absolute value over the course of a regime, the probability of a regime change heightens.

The tension index, along with NBER-dated recessions (indicated as shaded areas) is illustrated in Figure 1. The index undergoes 16 transitions between periods of general expansion and contraction, with transitions tending to precede NBER-defined business-cycle peaks and troughs by several quarters (transition dates are reported in Table 1). In generating forecasts, we exploit the predictive quality of the index by transforming forecasts of GDP growth into forecasted trajectories for the index, forecasted regime-change probabilities, and forecasted announcements of NBER-defined business-cycle turning points.

As noted, regime changes correspond with changes in the trajectory of the ECM-drift component. Parameters characterizing the nature of trajectories are latent, and assumed to be drawn from a fixed distribution. We also model the volatility of GDP growth as regime-specific, as determined by the realization of an additional latent parameter. As the duration of a given regime increases, the precision of our estimates of the realized parameters sharpens. By modelling drift and volatility parameters as stochastic, we effectively account for the heterogeneity observed across business cycles, along with the decrease in the volatility of GDP growth associated with the Great Moderation (for analyses of this phenomenon, see McConnell and Perez-Quiros, 2000; Kim and Nelson, 1999; and Stock and Watson, 2002).

The enhanced flexibility of the drift component we employ (relative to that of DLR) turns out to be critical in characterizing regime changes realized in the latter portion of the sample period. It also delivers distinct improvements in forecasting performance. For example, using rolling one-step-ahead forecasts over the last regime identified in our sample (2003:IV to 2008:II), the model yields reductions in root-mean-squared (RMSE) forecast errors of 13% and 53% relative to the DLR model and an alternative under which drift trajectories are constrained to be linear. Moving the beginning of the rolling forecast period to 1988:I, and again running through the sample, we obtain reductions of 5% and 30%.

A large and unsettled literature has assessed the importance of non-linearities in accounting successfully for the behavior of GDP growth. The literature consists of three general branches. One branch adopts a hypothesis-testing approach to determine whether linear models can be rejected in favor of non-linear alternatives. While such tests typically fail to reject when the non-linear alternative is Hamilton's (1989) relatively simple two-state Markov-switching model (e.g., see Garcia, 1998; and Hansen, 1992), rejections have been obtained using extensions of Hamilton's model as alternatives (e.g., see Hansen, 1992; and Kim, Morley and Piger, 2005).

A second branch has assessed the ability of alternative specifications to account for various facets of business-cycle behavior (e.g., average lengths of expansions and recessions). In parallel to the hypothesis-testing branch, when non-linear alternatives are relatively simple Markov-switching models, they are found to offer limited value-added relative to linear specifications (e.g., see Hess and Iwata, 1997; Harding and Pagan, 2002; and Clements and Krolzig, 2004). But when relatively complex non-linear alternatives are considered, value-added becomes more pronounced (e.g., see Galvao, 2002; Kim, Morley and Piger, 2005; and Morley and Piger, 2005).

Here we contribute to the third branch, which has analyzed the comparative forecasting performance of linear and non-linear specifications of GDP growth. Notably, Marcellino (2008) assessed the forecasting performance of 55 alternative univariate models, and found that the non-linear specifications included in the comparison set failed to outperform a baseline AR(4) model with a constant term specified for the log level of GDP. However, the comparison set did not include non-linear models in our class. Relative to Marcellino's baseline model, our specification yields reductions in RMSE of 42% and 43% over the (2003:IV to 2008:II) and (1988:I to 2008:II) forecasting horizons referenced above. Thus the non-linearities built into our model deliver substantial gains in forecasting performance.

We conclude this section by noting that we use the model not only to forecast GDP growth, but also to forecast the announcement of NBER turning points. To operationalize turning-point forecasts, we require an algorithm that converts output from the model into an NBER-dating rule.[1] Given the anticipatory relationship noted above between the tension index and NBER turning points, the index features prominently in the algorithm we employ.

We turn now to a description of the model, the methods we use to estimate it, and the algorithm

[1] The need for such an algorithm stems from the fact that NBER's dating method is neither transparent nor reproducible (Chauvet and Hamilton, 2005; Chauvet and Piger, 2008). Many attempts have been made to replicate NBER dates using non-parametric rule-based methods and statistical models (for an overview, see Harding and Pagan, 2002).

used to replicate NBER turning points. We then present model estimates, characterize the real-time evolution of regime-shift probabilities, discuss in-sample forecast performance, and conclude by presenting forecasts obtained using the latest available data.

# 2 The Model

## 2.1 Specification

As noted, the model is a refinement of that developed by DeJong, Liesenfeld and Richard (2005). The mean specification is common to both versions, and features an ECM and regime-switching behavior. Regime switches are manifested in the behavior of the ECM drift component that transitions between periods of general acceleration and deceleration.

Regime changes are triggered by a tension index, constructed as a geometric sum of past deviations of actual GDP growth $g_t$ from a corresponding "sustainable" growth rate $g_t^*$. Denoting the deviations $y_t = g_t - g_t^*$, the tension index $G_t$ is given by

$$G_t = \sum_{i=0}^{\infty} \delta^i y_{t-i} \ , \tag{1}$$

where $0 < \delta < 1$ measures the persistence of past deviations on current $G_t$. We specify the unobservable $g_t^*$ as the sample mean of $g_t$ (alternative specifications that admit slowly-evolving behavior yield similar forecasting characteristics). By implication, $g_t$ tends to pass between phases during which it alternately tends to outstrip and fall short of $g_t^*$. Under the interpretation of our model, neither phase is sustainable: both produce tension buildups (captured by increases in the absolute value of $G_t$) that ultimately lead to regime changes.

Regime-change probabilities are modelled using the logit specification

$$\pi_t = P(s_{t+1} = -s_t | s_t, G_t) = \frac{1}{1 + \exp\{\beta_0 - \beta_1 s_t G_t\}} \ , \tag{2}$$

where $s_t$ indicates the regime prevailing in period $t$, being 1 if $G_t$ is in an expansionary regime and -1 otherwise. Thus as the absolute value of $G_t$ increases, so too does $\pi_t$.

The model for GDP growth, in terms of its deviations from $g_t^*$, is given by

$$y_t = m_t - \nu G_{t-2} + \gamma y_{t-2} + \epsilon_t \ , \quad \epsilon_t | \sigma_t \sim N(0, \sigma_t^2) \ , \tag{3}$$

where $m_t$ represents a stochastic latent regime drift (the rationale behind this lag-2 specification follows the presentation of equations (5) and (6) below). Subtracting $y_{t-2}$ from both sides of (3) casts the model explicitly as an ECM representation, in which the term $\nu G_{t-2}$ reflects an integral correction based upon cumulated past deviations from equilibrium, and $(1-\gamma)y_{t-2}$ represents a proportional correction following the terminology of Phillips (1954, 1957).

The critical departure from DLR embodied in this model is the specification of the stochastic ECM drift component. The specification casts drift as a latent variable that jumps discontinuously at regime-change dates, and that follows piecewise non-linear trajectories that are heterogeneous across regimes. The specification is given by

$$m_t = \overline{m}_j + s_t a_j \left( \frac{e^{\tau b_j} - 1}{b_j} \right), \quad \tau \in \{0, 1, ..., (t(j) - t(j-1) - 1)\} \ , \tag{4}$$

where the index $j$ ($j : 1 \to J$) denotes the regime prevailing in period $t$, and $t(j)$ denotes the date at which regime $j$ gives way to regime $j+1$ (i.e. $t(j)$ is the last period under regime $j$).

The parameters of this specification are taken as random variables drawn from a fixed distribution at the onset of a new regime. The variable $\overline{m}_j$ represents the value of the regime drift in the first period of regime $j$, and the exponential term dictates the curvature of the $m_t$ trajectory during regime $j$. Specifically, $a_j$ represents the absolute value of velocity $\left(\left[\frac{dm_t}{d\tau}\right]_{\tau=0}\right)$ of the drift process at $\tau = 0$, and $a_j b_j$ the absolute value of acceleration $\left(\left[\frac{d^2 m_t}{d\tau^2}\right]_{\tau=0}\right)$. Both $a_j$ and $b_j$ are restricted as non-negative.

Inferred drift trajectories, along with deviations from these trajectories exhibited by the ECM growth term, are depicted in Figure 2 (the method used to infer these trajectories is described below). For comparison, Figure 2 also depicts drift trajectories inferred using the specification employed by DLR. Note that inferred trajectories vary widely across regimes, reflecting the well-known variability of business cycles themselves. In relatively long regimes, the trajectories obtained here and using the DLR specification are barely discernible; but in short regimes, differences are readily apparent: trajectories associated with the refined model tend to exhibit relatively dramatic jumps at regime-change dates, are highly non-linear, and closely track the residual ECM growth term. As we shall see, the payoff of these differences in terms of forecasting performance is substantial.

As noted, in order to allow for variation in the drift process $m_t$ across regimes, we specify $(a_j, b_j, \overline{m}_j)$ as latent random variables. Likewise, to capture heterogeneity in GDP volatility, the conditional variance of growth-rate innovations $\sigma_j^2$ is also latent, random, and regime-specific. An extensive diagnostic analysis led to the specification of a quadrivariate normal distribution for $\Lambda_j \equiv \left[\ln a_j, \ln b_j, \overline{m}_j, \ln \sigma_j^2\right]'$, with mean vector and covariance matrix $(\mu_\Lambda, \Sigma_\Lambda)$ treated as model parameters, and with $\ln \sigma_j^2$ taken as independent.

Returning to the model description, note that by pre-multiplying by $(1 - \delta L)$, (3) can be rewritten in the form of an (overidentified) ARMA(3,1) plus drift process

$$y_t = n_t + \delta y_{t-1} + (\gamma - \nu) y_{t-2} - \gamma \delta y_{t-3} + \epsilon_t - \delta \epsilon_{t-1} \; , \tag{5}$$

where within regime $j$ the variable $n_t$ is given by

$$n_t = (1-\delta)\overline{m}_j + \frac{a_j}{b_j} s_t \left[e^{\tau b_j} - \delta e^{(\tau - 1) b_j} - (1-\delta)\right] \; . \tag{6}$$

Its overidentification provides significant efficiency gains in the estimation stage at virtually no loss of fit. Note that the selection of lag 2 for the ECM representation (3) allows us to capture parsimoniously a non-zero coefficient on $y_{t-3}$ in (5), which turns out to be statistically significant.

## 2.2 Estimation

To characterize estimation, let $\theta$ represent the vector of all model parameters, $X_T$ the data, and $E_T = \{e_t\}_{t=1}^T$ a vector of zeros and ones, where $e_t = 1$ indicates a regime change period (i.e., the next period is the beginning of a new regime). Note that $E_T \in \Xi_T = \langle 0, 1 \rangle^T$, with cardinal $2^T$. Let $B(E_T) = \{t(j)\}_{j=1}^J$ denote the vector of regime change periods associated with $E_T$. It is critical to keep in mind that $J$ and the $t(j)\, s$ are implicit function of $E_T$ (though we do not make this explicit hereafter, for ease of notation).

$E_T$ is not observed, and brute force marginalization of the data density with respect to $E_T$ requires summation over the $2^T$ trajectories in $\Xi_T$ (the vast majority of which are assigned negligible probabilities). Following DLR, we instead apply two maximum likelihood (ML) procedures: ML estimation conditional on a particular trajectory $\widehat{E_T}$ (selected as described below); and ML estimation unconditional on $E_T$, relying on an importance sampling procedure designed to to identify likely trajectories.

For a given $E_T$, the conditional likelihood function is given by

$$L_C(\theta; X_T, E_T) = q\left(E_T; X_T, \theta_q\right) h\left(X_T; E_T, \theta_h\right), \tag{7}$$

where

$$q\left(E_T; X_T, \theta_q\right) = \prod_{t=1}^{T} \pi_t^{e_t}(1-\pi_t)^{(1-e_t)}, \tag{8}$$

$$h\left(X_T; E_T, \theta_h\right) = \prod_{j=1}^{J} f\left(X_j; E_T, \theta_h\right), \tag{9}$$

$$f\left(X_j; E_T, \theta_h\right) = \int \prod_{t=t(j-1)+1}^{t(j)} \frac{1}{\sigma_j^2}\phi\left(\frac{y_t-\mu_t}{\sigma_j}\right) f_N\left(\Lambda_j\right) d\Lambda_j, \tag{10}$$

and $X_j$ denotes the block of the data associated with regime $j$, $\theta = (\theta_q, \theta_h)$, $\theta_q = (\beta_0, \beta_1)$, $\mu_t = m_t - \nu G_{t-1} + \gamma y_{t-2}$, $f_N$ is the Normal density $N\left(\mu_\Lambda, \Sigma_\Lambda\right)$ introduced above, and $\phi()$ denotes the standardized Normal density.

Conditional ML estimation amounts to maximizing $\ln L_C$ with respect to $\theta$ conditionally on $\widehat{E}_T^*$, which is selected iteratively as follows. Let $\widehat{E}_T^0$ denote an initial trajectory (selected, e.g., by visual inspection of Figure 1), let $B\left(\widehat{E}_T^0\right) = \left\{\widehat{t}_0(j)\right\}_{j=1}^{\widehat{J}_0}$, and let $\widehat{\theta}_C^0$ denote the ML estimate of $\theta$ conditional on $\widehat{E}_T^0$. For each $j$ from 1 to $\widehat{J}_0$ we reconsider the location of $\widehat{t}_0(j)$. Specifically, let $t_1^0(j) = \widehat{t}_0(j-1)+1$ and $t_2^0(j) = \widehat{t}_0(j+1)$, so that the interval $\left[t_1^0(j), t_2^0(j)\right]$ represents the complete set of potential dates for $t\left(j\right)$. Then the probability that the $jth$ shift occurred at time $t$ over the interval is given by

$$\widehat{P}_0(t,j) = \frac{L_C(\widehat{\theta}_C^0; X_T, \widehat{E}_T^0\left(t,j\right))}{\sum_{s=t_1^0(j)}^{t_2^0(j)} L_C(\widehat{\theta}_C^0; X_T, \widehat{E}_T^0\left(s,j\right))}, \qquad t \in \left[t_1^0(j), t_2^0(j)\right], \tag{11}$$

where $\widehat{E}_T^0(t,j)$ denotes the modified trajectory obtained by substituting $t$ for $t(j)$ in $\widehat{E}_T^0$ (i.e., by setting $e_{t(j)}^0 = 0$ and $e_t^0 = 1$). For each $j : 1 \to \widehat{J}_0$, a new date $t^*(j)$ is defined as that which maximizes $\widehat{P}_0(t,j)$. Application of this procedure produces a revised trajectory $\widehat{E}_T^*$ and a new conditional ML estimate $\widehat{\theta}_C^0$. The procedure is repeated until convergence to a fixed-point solution $\widehat{\theta}_C^*$, which constitutes the conditional ML estimate of $\theta$. This process typically converges within three rounds (hardly surprising, considering the strong informational content of the index $G_t$, as highlighted in Figure 1). This iterative procedure is akin to the EM-type algorithms typically employed in estimated Markov-switching models (as outlined, e.g., in Hamilton 1994, and Diebold, Lee, and Weinbach 1994).

Note that a similar iterative procedure can be applied in order to see whether new regime breaks need to be introduced into the sample. This possibility is particularly relevant as new observations become available, and one wishes to determine whether a new break may have occurred since the last identified break $J$. In such a case the search for a potential new break is conducted on the interval $t_1^0(J+1) = \widehat{t}_0(J) + 1$ to $t_2^0(J+1) = T$. For any $t \in \left[t_1^0(J+1), t_2^0(J+1)\right]$, the modified trajectory $\widehat{E}_T^0(t,-)$ obtains by setting $\widehat{e}_t^0 = 1$ instead of its current zero. Formula (11) still applies, with the additional additive term $L_C(\widehat{\theta}_C^0; X_T, \widehat{E}_T^0)$ appearing in the denominator to account for the possibility of no new break, and also appearing in the numerator when the probability of no new break is calculated. The search for additional breaks in the interior of the sample proceeds analogously.

Unconditional importance sampling estimation employs conditional ML estimates $\widehat{\theta}_q^C$ to produce a sequence $\{\widehat{\pi}_t\}_{t=1}^T$ of probability estimates of regime changes. These probabilities are used as an importance sampler for the unobserved $E_T$. Specifically, these probabilities are used to produce $R = 1,000$ trajectories $\left\{\widetilde{E}_{T,r}\right\}_{r=1}^R$. Given these trajectories, the corresponding IS estimate of the unconditional likelihood function is given by

$$\overline{L}_S(\theta; X_T) = \frac{1}{R}\sum_{r=1}^{R} \frac{q\left(\widetilde{E}_{T,r}; X_T, \theta_q\right)}{q\left(\widetilde{E}_{T,r}; X_T, \widehat{\theta}_q^C\right)} h\left(X_T; \widetilde{E}_{T,r}, \theta_h\right). \tag{12}$$

The unconditional ML estimate of $\theta$ is that which maximizes $\ln \overline{L}_S$.

Both conditionally and unconditionally upon regime-change dates, likelihood evaluation requires integration over the latent parameters $\Lambda_j$, $j = 1, ..., J$. To accomplish this, we utilize the efficient importance sampling (EIS) procedure introduced by Richard and Zhang (2007). Finally, we obtained ML estimates using the simplex algorithm of Nelder and Mead (1965), implemented in IMSL FORTRAN numerical libraries as routine UMPOL.

## 2.3 Dating Recessions

As noted, our forecasting objectives include the anticipation of NBER-dated turning points. Towards this end, we require an algorithm that converts output from the model into an approximated NBER dating rule. A simple and popular example of such an algorithm is the "two consecutive quarters of negative growth" rule for defining the onset of a recession. However (no doubt in part due to the fact that growth-rate data are often revised, while NBER dates are not), this rule is far from infallible in the data-set vintage we analyze.

The algorithm we employ instead features a modification of the two-quarter rule, combined with behavior of the tension index $G_t$. To motivate our use of $G_t$, recall from Figure 1 the proximity of business-cycle peaks and troughs to transitions in $G_t$ from general periods of expansion and contraction. Note also that no NBER-recession fails to coincide with the crossing of $G_t$ over the threshold $-5$. Using this fact, and experimenting with alternative characterizations of changes in $G_t$ trajectories, we developed the following algorithm for dating NBER turning points. Let

$$\Delta\overline{G}_t = \overline{G}_{[t,t+3]} - \overline{G}_{[t-3,t]}, \tag{13}$$

where $\overline{G}_{[t,t+3]}$ represents the arithmetic mean of $G_t$ over the 4-period window $[t, t+3]$. Then:

- **Preselection of potential start date intervals.** A potential interval is defined as a (maximal) sequence of contiguous periods $t$ for which at least two periods of negative growth are observed in the four-period window $[t, t+3]$.

- **Selection of start dates.** Let $[t_1, t_2]$ denote a preselected interval. A recession-start date is defined as the period $t_1^*$ that minimizes $\Delta\overline{G}_t$ subject to the constraint that $\overline{G}_{[t_1^*,t_1^*+3]}$ exceeds $-5$. If no date in $[t_1, t_2]$ exceeds this threshold, the interval is eliminated as containing a possible recession (three such instances arise in our sample).

- **Selection of end dates.** Let $[t_1, t_2]$ denote a preselected interval containing the chosen start date $t_1^*$. The interval of potential end dates is defined as $[t_1^* + 1, t_2]$. The selected end date $t_2^*$ is defined as the last date in the interval such that $G_s < 0$, $s \in [t_1^* + 1, t_2]$, subject to the constraint that $g_{t_2^*} < 2.75$.

Application of this algorithm yields a close approximation to NBER-defined turning points. It succeeds in flagging each NBER-defined recession, and avoids flagging recessions spuriously. Regarding the timing of defined recessions, the algorithm misses only three start dates, each by one quarter: the NBER-defined peaks of (1969:IV, 1980:I, 1981:III) are identified instead as (1969:III, 1979:IV, 1981:IV). And it misses only one end date: 1991:I is identified instead as 1991:III.

# 3 Results

## 3.1 Estimates of Parameters and Latent Regime-Specific Variables

The sequence of regime-shift dates used to obtain conditional ML parameter estimates are reported in Table 1, and conditional and unconditional ML parameter estimates are reported in Table 2. Estimates were obtained using the annualized growth rate of quarterly U.S. GDP measured in chain-weighted 2000 prices, spanning 1950:III to 2008:II.

Conditional and unconditional point estimates are closely comparable, with conditional estimates tending to be relatively precise. Two aspects of these estimates are particularly notable. First, the estimates indicate a strong error-correction effect. For example, the conditional estimate of the autoregressive coefficient $\gamma$ (0.3426, with s.e. 0.0416) indicates non-trivial persistence for shocks; its difference $(1 - \gamma)$ translates into a fairly strong proportional error correction of 0.6574. In turn, the ECM coefficient $\nu$ is estimated as 0.2522 (s.e. 0.051), indicating that past errors in the lagged tension index exert appreciable influence over GDP growth. Second, estimates obtained for the distribution of the latent regime-drift parameters indicate a significant positive interaction between $(\ln a_j, \ln b_j)$ across regimes; recall that these variables jointly determine the initial velocity and acceleration of the regime drift. Additional interactions are insignificant.

Smoothed estimates of the latent regime-specific parameters are illustrated in Figure 3 (plotted estimates were obtained using conditional ML estimates; unconditional estimates are similar). Note that the last five values of $\ln \sigma_j^2$ lie below the estimated sample mean of 2.0709, illustrating the impact of the Great Moderation. In addition, estimates of the regime-drift parameters exhibit mild downward trends across regimes. It is possible to exploit these patterns by introducing additional complexity into the distributions specified for these latent parameters (e.g., means of the latent parameters could be modelled as decreasing functions of identified regimes). However, with only 17 regimes identified in the sample, the precision with which tendencies across regimes could be identified is limited, thus the payoff of this added complexity is low. Moreover, as highlighted in Section 3.3 below, the in-sample forecasting performance of the model is very strong, and does not appear to suffer from the maintained assumption of a stable distribution for regime characteristics.

### 3.2 The Evolution of Regime-Shift Probabilities

The regime-shift dates indicated in Table 1 were obtained using the full data set. Figure 4 illustrates how inferences regarding the last four selected break dates evolved in real time. To see how, consider the upper-left panel of the figure. Given the regime-break date identified in 1984:I, and using data observed through 1989:IV, conditional model estimates were obtained and used to calculate the probability that a subsequent break date had not occurred. As the figure indicates, the probability assigned to this scenario is essentially 1. Adding an additional data point, re-estimating the model, and re-calculating break probabilities, we continue to roll through the sample. Beginning in 1990:III, a second probability is plotted: that associated with a possible break in 1990:III; four additional plots are added over the next four quarters. Jointly, the probability plots illustrate the ability of the model to track the evolution of regime changes.

In the upper-left figure, note that 1990:IV initially appears as a strong potential break date: the initial probability of a break at this date is 60%. However, in the next quarter this probability drops below 20%, and the ultimate break date of 1991:I is assigned a break-point probability of 75%. By 1991:II the probability of no break is essentially zero, with 1990:IV and 1991:I both appearing as likely candidates for break dates. By 1992:III, 1991:I emerges as the clear choice for a break.

In contrast to the rapid identification of 1991:I as a clear candidate for a break, Figure 4 indicates that probabilities assigned to the subsequent break dates of 1999:IV, 2001:III and 2003:III evolved relatively slowly and gradually. In turn, note from Table 1 that the full-sample probabilities assigned to these specific dates are relatively low: 62%, 87%, and 68%, relative to the 94.5% probability assigned to 1991:I.

The message from Figure 4 is that the identification of regime changes often takes time, and in general we face uncertainty regarding the current state of the economy as we seek to generate forecasts. We account for this uncertainty explicitly by generating sets of conditional forecasts, where conditioning is with respect to the full range of break-date scenarios evident given the latest available data. We then compute unconditional forecasts as the weighted average of the conditional forecasts, with weights given by the probabilities assigned to the alternative conditional scenarios. Characteristics of these forecasts are detailed below.

### 3.3 In-Sample Forecasting Performance

As noted, a large and unsettled literature has assessed the importance of non-linearities in accounting for the behavior of GDP growth. Marcellino (2008) contributed to this literature by comparing the in-sample forecasting performance of 55 models relative to a baseline AR(4) model with constant specified for the logged-levels of GDP (from which growth-rate forecasts are straightforward to generate via a difference transformation of levels forecasts). While the non-linear models he considered failed to outperform his baseline specification, he did not consider non-linear models in our class. Thus we present a modification of his analysis to assess the contribution of the non-linearities built into our model.

Following Marcellino, for a given model we obtained a series of forecasts by truncating the sample period, generating a one-step-ahead forecast, augmenting the sample with an additional observation, re-estimating, and generating another forecast. We then computed root-mean- and root-variance-squared-error (RMSE and RVSE) statistics over various forecasting windows, including the period spanning 1988:I through 2008:II (the end of our sample), and the sub-periods 1988:I-1999:I, 1991:II-1999:IV, 2000:I-2001:III, 2001:IV-2003:III, 2003:IV-2008:II. The sub-periods correspond with the last five regimes identified in the sample.

In addition to our model, we generated forecasts using Marcellino's AR specification, a random walk with drift specified for GDP growth, the DLR model, and a version of our model under which regime-drift trajectories were restricted to be linear (hereafter, the linear-$m_t$ model). Comparisons with these latter two models are particularly revealing, because identified regime changes are identical across specifications: departures from linearity in regime-shift specifications are the sole source of differences in forecasting performance across these specifications.

Table 3 reports RMSE and RVSE statistics; and Figure 5 illustrates the time series of squared forecast errors obtained using our model and the AR(4) specification (vertical lines indicate regime-change dates and NBER-defined turning points). In the brief sub-period 2000:I-200I:III, the AR(4) model yields optimal performance on the basis of RMSE comparisons; in all other cases the model we have presented dominates, often substantially. Over the entire in-sample horizon, our model yields a 43% reduction in RMSE relative to the AR(4) model, and 30% relative to the linear-$m_t$ model. From the figure, note that squared forecast errors associated with the AR(4) model spike around the last two regime-change dates we identify, but additional spikes are also observed during the relatively tranquil period of the late 1990s.

### 3.4 Forecasting Future Growth Rates and Recessions

Beyond forecasting GDP growth, the model is also useful for anticipating future regime changes, along with NBER-dated turning points. We illustrate this in Figures 6 and 7, which present the results of model simulations computed over a twelve-quarter horizon. The simulations are of forecasted trajectories for growth, conditional on the model estimated using data through 2008:II, and on the set of identified regime-break dates reported in Table 1.

Recall from Table 1 that the economy was inferred to have entered a growth-deceleration regime in 2003:III. Using subsequent observations of GDP growth, we compute smoothed values of the parameters that determine the trajectory followed by the ECM drift $m_t$, and thus smoothed values of the trajectory itself. Extending the trajectory beyond the end-of-sample period $T$, and simulating subsequent ECM innovations using the smoothed estimate of the latent innovation-variance parameter, we obtain a simulated trajectory for growth using (3).

For a given simulated trajectory, so long as a regime break does not occur over the forecast horizon, the trajectory follows the uninterrupted path for $m_{T+i}$ described above. However, for each period in the forecast horizon, the probability of a regime change triggered by the continuation of the trajectory is calculated using (2). We then simulate a coin toss that triggers a regime change with the calculated probability. When a regime change is realized, say, at date $T+i$, a new trajectory is inferred for $\{m_{T+i+1}, m_{T+i+2}, ...\}$.

The top panel of Figure 6 reports a histogram of regime-break probabilities (including the probability of no break) over the forecast horizon obtained using 500,000 simulated trajectories. The forecasted probability of no break over the horizon is 1.6%, and 2010:II and 2010:III appear as the dates most likely to feature a break (break probabilities in both cases are 14.8%).

The middle panel of Figure 6 plots mean growth trajectories simulated for each possible break-date scenario over the forecast horizon. Note that mean forecasts jump discontinuously whenever a break is inferred to occur. In turn, the bottom panel of the figure integrates over all simulated trajectories to produce unconditional point forecasts and 95% confidence intervals. The outlook based on the figure is bleak: growth is predicted to remain sluggish and flat (below 2% with near-zero slope) over the next two years.

Figure 7 illustrates implications of this outlook for the likely occurrence of an NBER-dated recession by mapping model simulations into the dating algorithm described above. The top, middle and bottom panels plot histograms of recession start-date, end-date, and duration probabilities. The probability of no recession is 11.4%. Start dates are fairly evenly distributed, with 2009:I appearing most likely for a start (15.2%); end dates are also evenly distributed, with 2010:II appearing most likely (11.7%). Finally, the histogram for recession length is sharply peaked at 4 quarters, which is assigned a probability of 47.3%.

## 4 Conclusion

We have presented a regime-switching model of GDP growth that can be used to anticipate NBER-dated turning points. Parameters charaterizing growth trajectories and innovation volatilities are modelled as latent and regime-specific, thus the model captures temporal changes in growth behavior in the absence of changes in its underlying structure. There is considerable heterogeneity in the model's characterization of growth trajectories across regimes, with distinct departures from non-linearity appearing frequently. Based on an in-sample forecasting exercise, allowances for non-linear behavior generate distinct improvements in forecasting performance.

## References


[1] Beaudry, P., and C. Koop, 1993. "Do Recessions Permanently Change Output?" *Journal of Monetary Economics* 31:149-163.

[2] Boldin, M.D., 1996. "A Check on the Robustness of Hamilton's Markov Switching Model Approach to the Economic Analysis of the Business Cycle," *Studies in Nonlinear Dynamics and Econometrics* 1: 35-46.

[3] Chauvet, M. and J. D. Hamilton, 2005. "Dating Business Cycle Turning Points," Working Paper 11422, National Bureau of Economic Research.

[4] Chauvet, M. and J. Piger, 2008. "A Comparison of the Real-Time Performance of Business Cycle Dating Methods," *Journal of Business and Economic Statistics,* 26:42-49.

[5] Clements, M.P. and H.-M. Krolzig, 1998. "A Comparison of the Forecast Performance of Markov-Switching and Threshold Autoregressive Models of US GNP," *Econometrics Journal* 1:47-75.

[6] Clements, M.P. and H.-M. Krolzig, 2004. "Can Regime-Switching Models Reproduce the Business Cycle Features of US Aggregate Consumption, Investment and Output?" *International Journal of Finance and Economics* 9:1-14.

[7] DeJong, D. N., R. Liesenfeld, and J.F. Richard, 2005. "A Non-Linear Forecasting Model of GDP Growth," *Review of Economics and Statistics* 87(4): 97-708.

[8] Diebold, F.X., J.-H. Lee, and G.C. Weinbach, 1994. "Regime Switching with Time-Varying Transition Probabilities." In *Nonstationary Time Series Analysis and Cointegration*, edited by Hargraves C., Oxford: Oxford University Press, 283:302.

[9] Durland, J.M. and T.H. McCurdy, 1994. "Duration-Dependent Transitions in a Markov Model of U.S. GDP Growth," *Journal of Business and Economic Statistics* 12:279-288.

[10] Filardo, A.J. and S.F. Gordon, 1998. "Business Cycle Durations," *Journal of Econometrics* 85:99-123.

[11] Galvao, A.B., 2002. "Can Non-Linear Time Series Models Generate US Business Cycle Asymmetric Shape?" *Economics Letters* 77:187-194.

[12] Garcia, R., 1998. "Asymptotic Null Distribution of the Likelihood Ratio Test in Markov Switching Models," *International Economic Review* 39:763-788.

[13] Hamilton, J.D. 1994. *Time Series Analysis.* Princeton: Princeton University Press.

[14] Hamilton, J.D., 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica* 57: 357-384.

[15] Hansen, B.E., 1992. "The Likelihood Ratio Test Under Non-Standard Conditions: Testing the Markov Switching Model of GNP," *Journal of Applied Econometrics* 7:S61-S82.

[16] Harding, D. and A. Pagan, 2002. "A Comparison of Two Business Cycle Dating Methods," *Journal of Economic Dynamics and Control,* 27:1681-1690.

[17] Hess, G.D., and S. Iwata, 1997. "Measuring and Comparing Business-Cycle Features," *Journal of Business and Economic Statistics,* 15:432-444.

[18] Kim, C.J., J. Morley and J. Piger, 2005. "Nonlinearity and the Permanent Effects of Recessions," *Journal of Applied Econometrics* 20:291-309.

[19] Kim, C.J., and C. R. Nelson, 1999. "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of Business Cycle," *Review of Economics and Statistics* 81:1–10.

[20] Kim, C.J., C.R. Nelson and J. Piger, 2004. "The Less-Volatile U.S. Economy: A Bayesian Investigation of Timing, Breadth, and Potential Explanations," *Journal of Business and Economics Statistics* 22:80-93.

[21] Marcellino, M., 2008. "A Linear Benchmark for Forecasting GDP Growth and Inflation?" *Journal of Forecasting,* 27:305-340.

[22] McConnell, M.M. and G. Perez-Quiros, 2000. "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review* 90:1464-1476.

[23] Morley, J. and J. Piger, 2005. "The Importance of Nonlinearity in Reproducing Business Cycle Features," Washington University in St. Louis Working Paper.

[24] Nelder, J.A. and R. Mead, 1965. "A Simplex Method for Function Minimization," *Computer Journal* 7:308-313.

[25] Ocal, N. and D.R. Osborn, 2000. "Business Cycle Non-Linearities in UK Consumption and Production," *Journal of Applied Econometrics* 15:27-44.

[26] Pesaran, M.H. and S.M. Potter, 1997. "A Floor and Ceiling Model of US Output," *Journal of Economic Dynamics and Control* 21:661-695.

[27] Phillips, A.W., 1954. "Stabilization Policy in a Closed Economy," *Economic Journal,* 64:290–333.

[28] Phillips, A.W., 1957. "Stabilization Policy and the Time Form of Lagged Response," *Economic Journal,* 67:265-277.

[29] Press, W.H., S.A. Teukolsky, W.T. Vetterling and B.P. Flannery, 1996. *Numerical Recipes in Fortran 90*, Cambridge: Cambridge University Press.

[30] Richard, J.-F. and W. Zhang, 2007. “Efficient High-Dimensional Monte Carlo Importance Sampling,” *Journal of Econometrics,* 141:1385-1411.

[31] Stock, J.H. and M.W. Watson, 2002. “Has the Business Cycle Changed and Why?” *NBER Macroeconomics Annual,* 17:159–218.

[32] Tong, H. 1990. *Non-Linear Time Series: A Dynamical Systems Approach.* Oxford: Oxford University Press.

[33] van Dijk, D. and P.H. Franses, 1999. “Modeling Multiple Regimes in the Business Cycle,” *Macroeconomic Dynamics* 3:311-340.

# 5 Tables and Figures

**Table 1. Regime-Shift Dates and Associated Probabilities**

| No. | Break Date | Estimated Probability | No. | Break Date | Estimated Probability |
|---|---|---|---|---|---|
| **1** | **1953:IV** | **0.977** | **10** | **1978:II** | **0.955** |
| **2** | **1955:I** | **0.956** | **11** | **1980:II** | **0.965** |
| **3** | **1958:I** | **0.961** | | 1983:IV | 0.341 |
| | | | **12** | **1984:I** | **0.454** |
| | 1959:I | 0.102 | | 1984:II | 0.114 |
| **4** | **1959:II** | **0.882** | | | |
| | | | **13** | **1991:I** | **0.945** |
| **5** | **1960:IV** | **0.949** | | | |
| | | | | 1997:III | 0.097 |
| | 1965.III | 0.105 | | 1998:IV | 0.075 |
| **6** | **1965:IV** | **0.705** | **14** | **1999:IV** | **0.620** |
| **7** | **1970:IV** | **0.901** | | 2001:II | 0.062 |
| | | | **15** | **2001:III** | **0.872** |
| **8** | **1973:I** | **0.915** | | 2001:IV | 0.050 |
| **9** | **1974.III** | **0.891** | **16** | **2003:III** | **0.678** |
| | | | | 2003:IV | 0.102 |
| | | | | 2004:III | 0.074 |

Note: Identified regime-shift dates are in boldface; local alternatives receiving non-negligible probabilities are also listed.

**Table 2. Parameter Estimates**

| | Conditional ML | | Unconditional ML | |
|---|---|---|---|---|
| Parameter | Estimate | Asym. Std. Error | Estimate | Asym. Std. Error |
| $\nu$ | -0.2522 | 0.0510 | -0.2427 | 0.1183 |
| $\gamma$ | 0.3426 | 0.0416 | 0.3328 | 0.0664 |
| $\beta_0$ | 10.877 | 0.9028 | 10.161 | 1.2494 |
| $\beta_1$ | 1.2124 | 0.0784 | 1.3132 | 0.2136 |
| $\delta$ | 0.5875 | 0.1910 | 0.5733 | 0.2745 |
| $\ln a_j$ | -2.7623 | 0.2320 | -2.8740 | 0.7324 |
| $\ln b_j$ | -1.1255 | 0.2214 | -1.2918 | 0.4133 |
| $\overline{m}_j$ | 0.4322 | 0.1322 | 0.3923 | 0.1325 |
| $\ln \sigma_j^2$ | 2.0709 | 0.2880 | 2.0173 | 0.5317 |
| $Cov(\ln a_j, \ln b_j)$ | 0.6335 | 0.0423 | 0.5662 | 0.2139 |
| $Cov(\ln a_j, \overline{m}_j)$ | -0.0039 | 0.0358 | 0.0069 | 0.0867 |
| $Cov(\ln b_j, \overline{m}_j)$ | -0.0087 | 0.0399 | -0.0174 | 0.1851 |
| $Var(\ln a_j)$ | 0.9127 | 0.2830 | 0.9828 | 0.3934 |
| $Var(\ln b_j)$ | 0.4935 | 0.1719 | 0.4240 | 0.1654 |
| $Var(\overline{m}_j)$ | 3.5381 | 0.4933 | 3.6405 | 0.5695 |
| $Var(\ln \sigma_j^2)$ | 1.0292 | 0.5940 | 1.0921 | 0.6038 |
| Log-Likelihood | -611.130 | | -604.617 | |

Note: The sequence of regime-shift dates used for Conditional Maximum Likelihood are reported in Table 1.

## Table 3. In-Sample Forecast Performance

| Model: | DHLR | DLR (2005) | $m_t$-Linear | Random Walk | AR(4) |
|---|---|---|---|---|---|
| | **1988:I to 1991:I** | | | | |
| Root Mean Squared Error | 100 | 100 | 122 | 114 | 128 |
| Root Variance Squared Error | 100 | 127 | 151 | 109 | 135 |
| | **1991:II to 1999:IV** | | | | |
| Root Mean Squared Error | 100 | 105 | 126 | 134 | 193 |
| Root Variance Squared Error | 100 | 103 | 199 | 188 | 340 |
| | **2000:I to 2001:III** | | | | |
| Root Mean Squared Error | 100 | 112 | 134 | 148 | 90 |
| Root Variance Squared Error | 100 | 138 | 116 | 223 | 177 |
| | **2001:IV to 2003:III** | | | | |
| Root Mean Squared Error | 100 | 104 | 134 | 102 | 100 |
| Root Variance Squared Error | 100 | 107 | 132 | 67 | 117 |
| | **2003:IV to 2008:II** | | | | |
| Root Mean Squared Error | 100 | 113 | 153 | 153 | 142 |
| Root Variance Squared Error | 100 | 109 | 228 | 283 | 156 |
| | **1988:I to 2008:II** | | | | |
| Root Mean Squared Error | 100 | 105 | 130 | 131 | 143 |
| Root Variance Squared Error | 100 | 116 | 156 | 174 | 183 |

Reported statistics are relative to DHLR, normalized to 100 in all forecast horizons. Forecast errors were computed using rolling one-period-ahead forecasts of GDP growth (with each model re-estimated following updates in the observed sample). The random walk model included a drift term; the AR(4) model included a constant, and was specified for the log-level of GDP (log-level forecasts were then converted into forecasted growth rates).

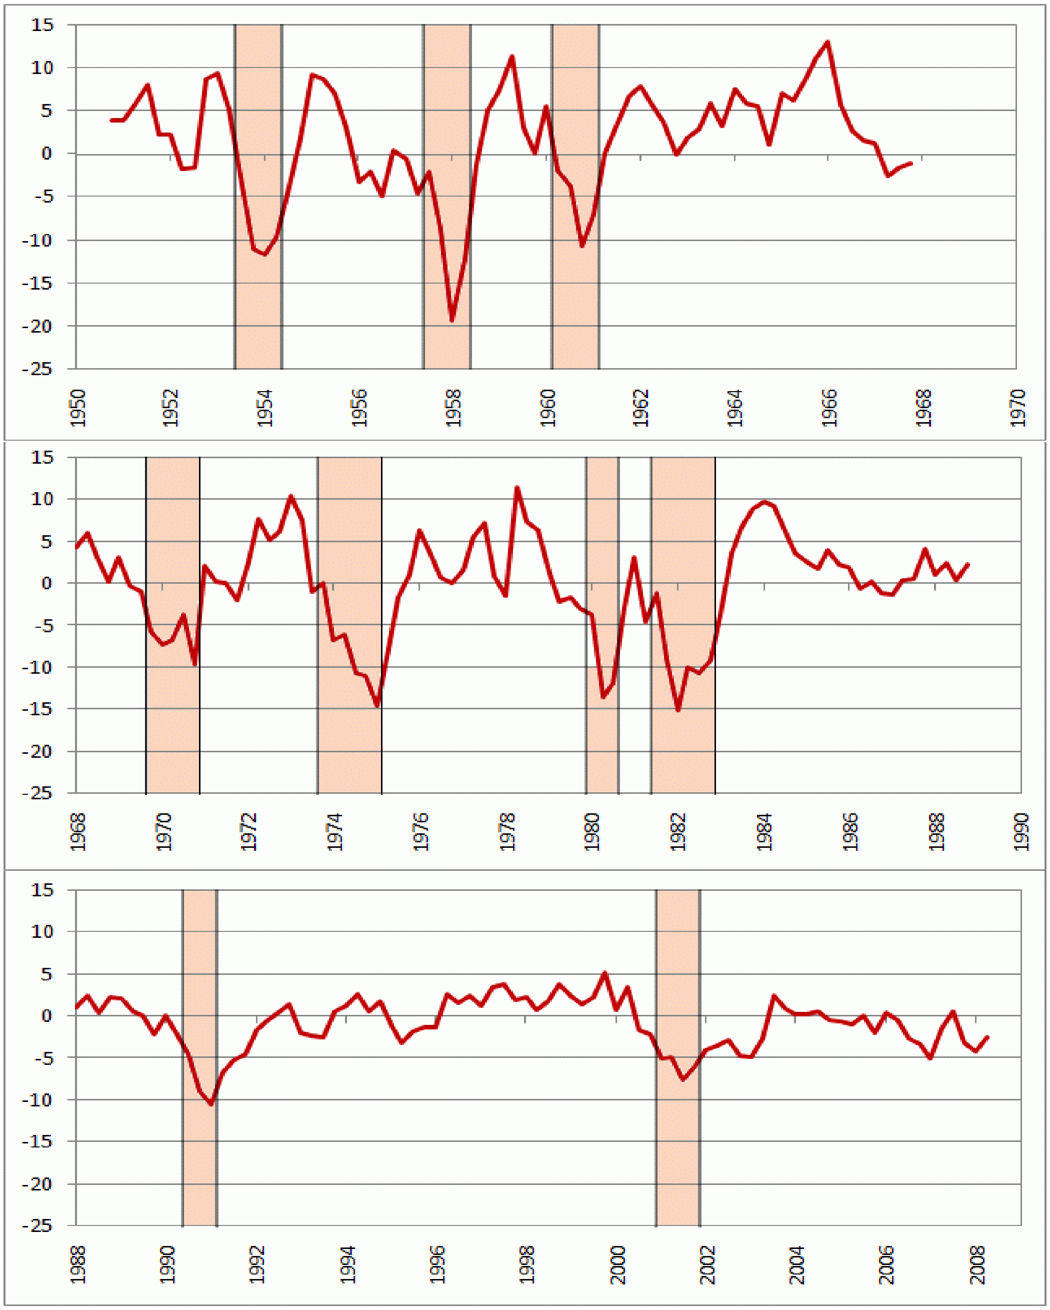


**Figure 1. Tension Index and NBER-Dated Recessions**

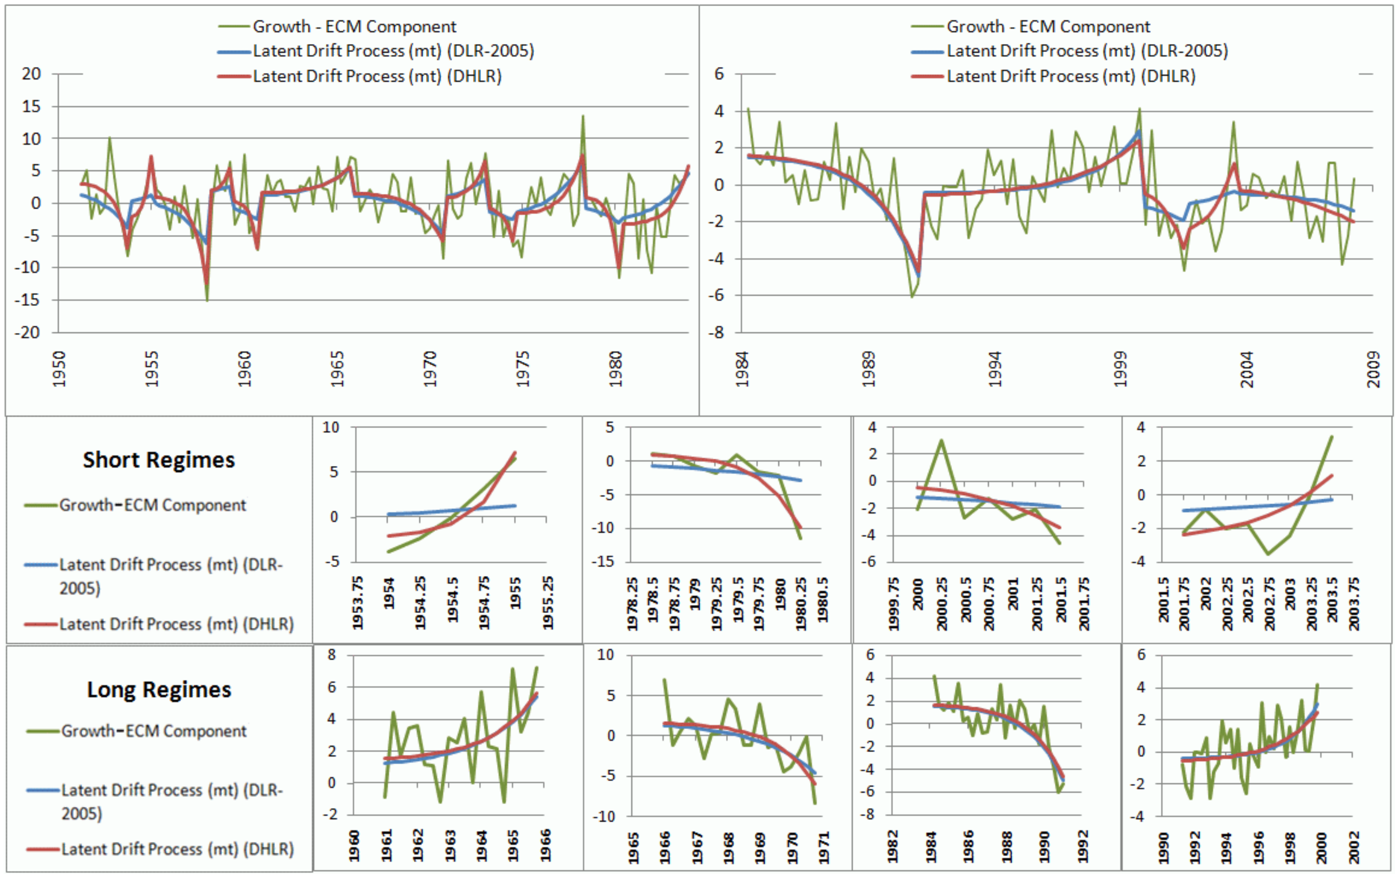


**Figure 2. Comparison of the Latent Drift Process ($m_t$).**

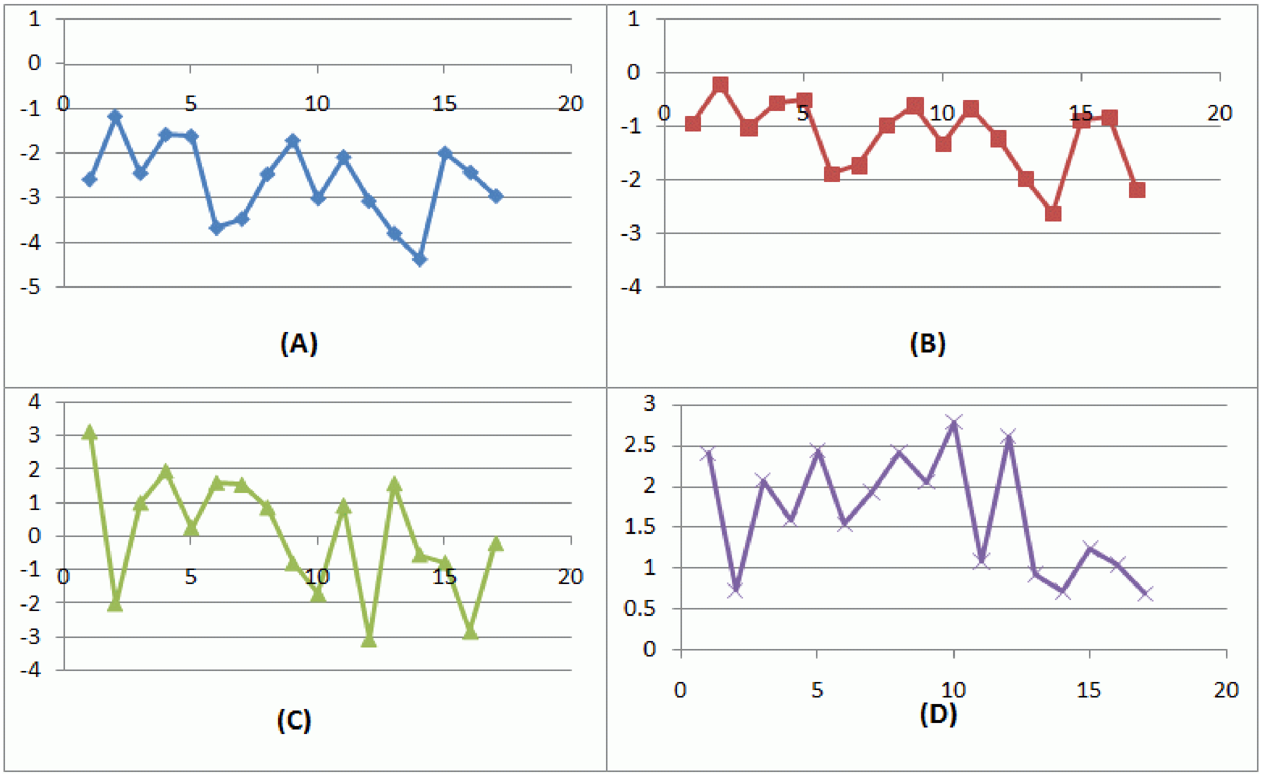


**Figure 3. Smoothed Values of the Latent Variables.**

**Panel (A)** $\rightarrow a_j$, **(B)** $\rightarrow b_j$, **(C)** $\rightarrow \overline{m}_j$, **(D)** $\rightarrow \ln \sigma_j^2$.

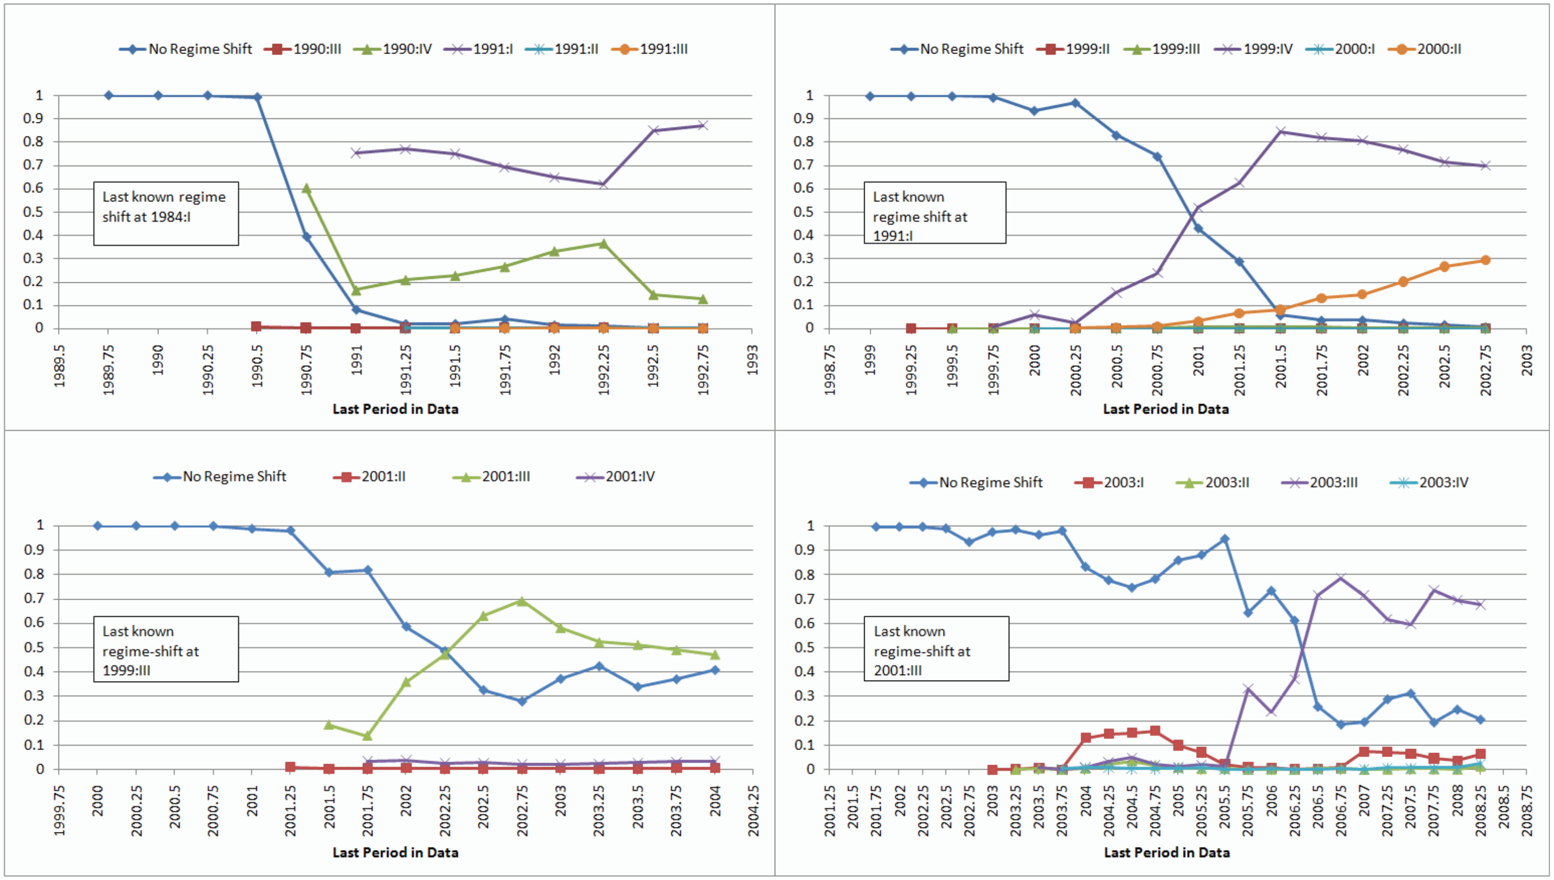


**Figure 4. Identifying Regime Shifts.**

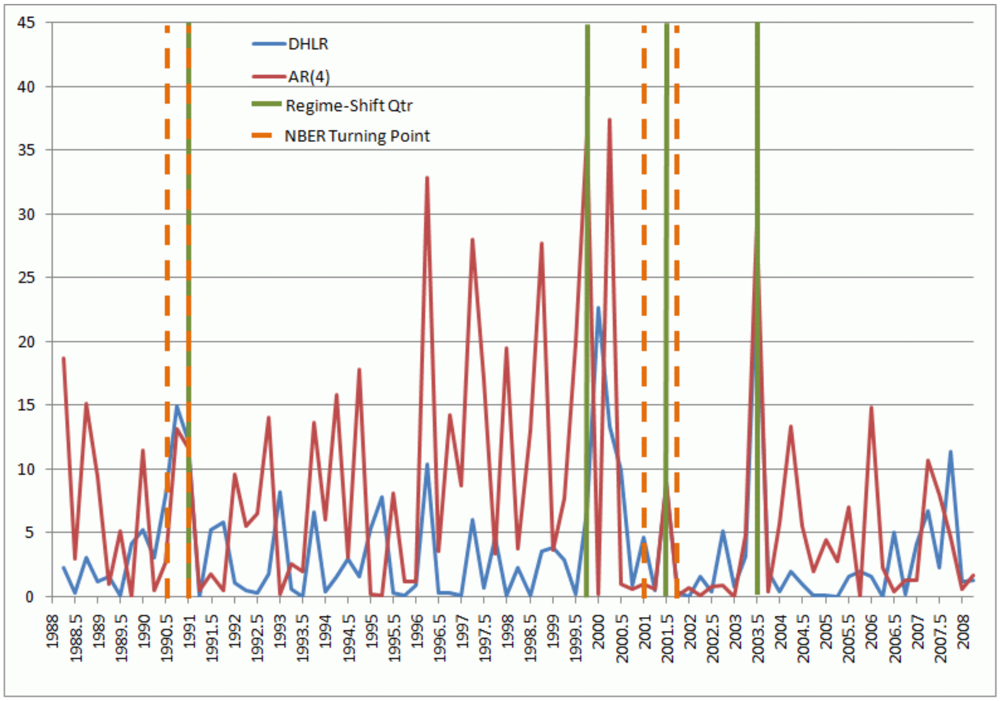


**Figure 5. Squared Forecast Error Comparisons.**

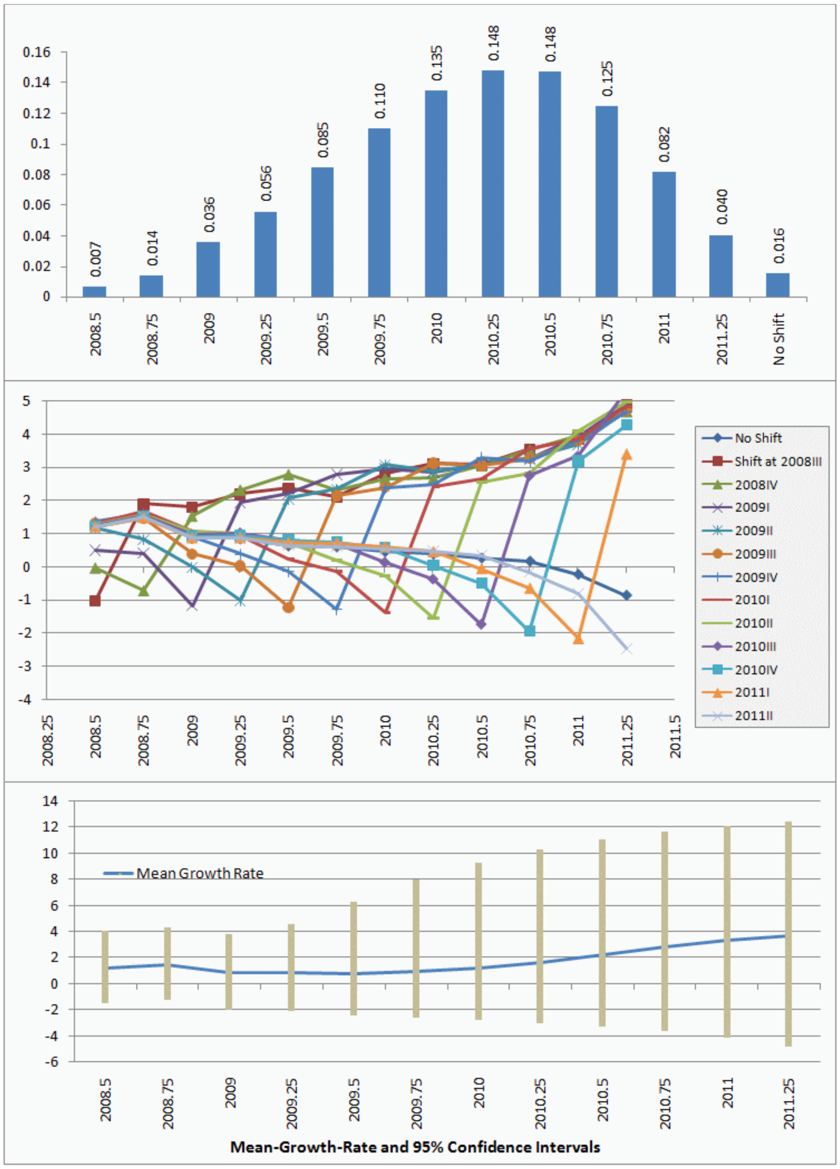


**Figure 6. Forecasted Regime-Shift Probabilities and Growth Rates.**

**Top Panel: Regime-Shift Probabilities.**
**Middle Panel: Conditional forecasted mean growth trajectories.**

**Bottom Panel: Unconditional forecasted mean growth trajectories.**

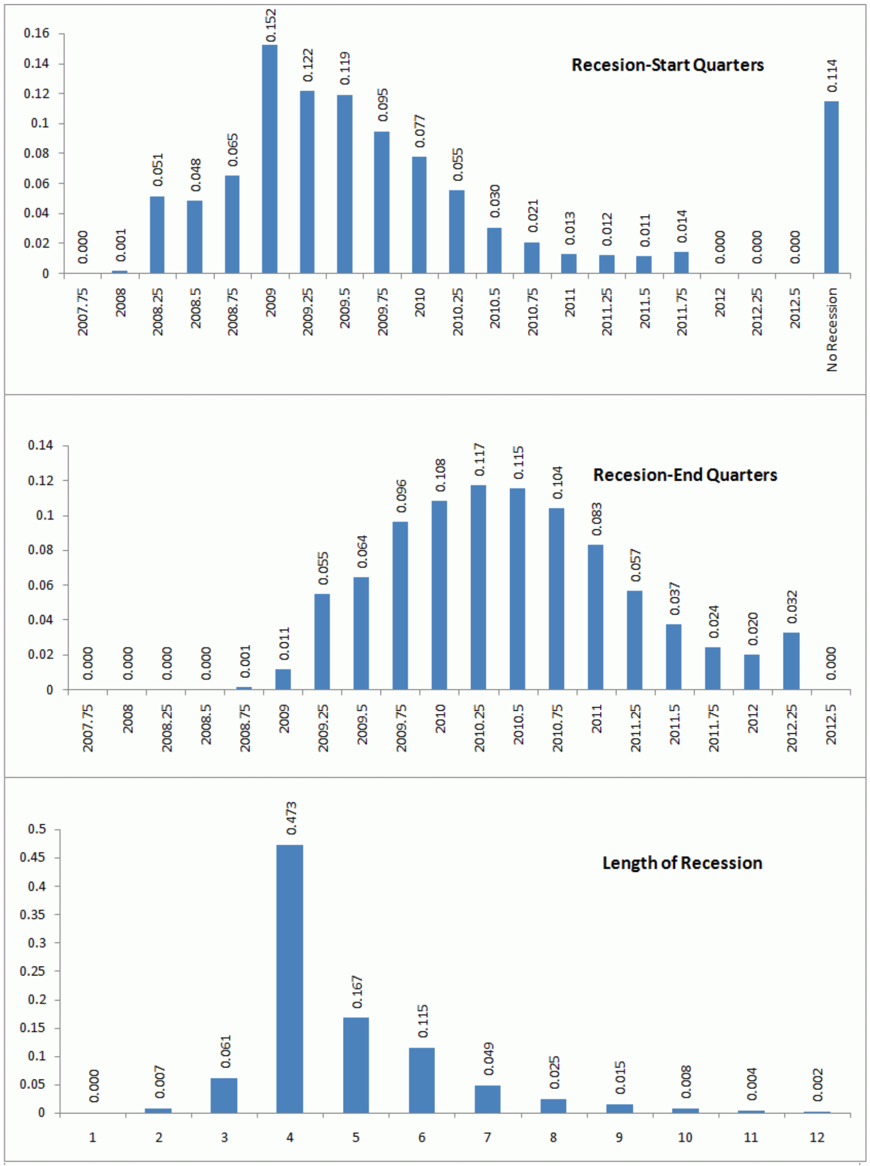


**Figure 7. Forecasting NBER-Dated Recessions.**

**Top Panel: Recession-Start Probabilities.**
**Middle Panel: Recession-End Probabilities.**

**Bottom Panel: Duration Probabilities.**