Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Journal of Monetary Economics 54 (2007) 1670–1701
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Tracking the new economy: Using growth theory to detect changes in trend productivity ☆

James A. Kahn*, Robert W. Rich

Domestic Research Department, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045, USA

Received 19 November 2004; received in revised form 13 March 2006; accepted 5 July 2006
Available online 19 December 2006

Abstract

The acceleration of productivity after 1995 prompted a debate over whether the economy's underlying growth rate would remain high. In this paper, we draw on growth theory to identify variables other than productivity—namely consumption and labor compensation—to help estimate trend productivity growth. We treat that trend as a common factor with two "regimes," high- and low-growth. Our analysis picks up striking evidence of a return in 1997 to the high-growth regime, nearly 25 years after a switch from high- to low-growth. We find that both the common factor and regime-switching aspects of the model are important for identifying changes in trend productivity, and also show that the trend breaks are more difficult to detect with per capita (as opposed to per hour) based data because of persistent labor supply shifts. Finally, we argue that our methodology is effective in detecting changes in trend in real time: In the case of the 1990s, the methodology would have signaled the regime switch by 1999, or within roughly six quarters of when it occurred according to the full sample.

JEL classification: O4; O51; C32

Keywords: Productivity growth; Regime-switching; Neoclassical growth model; Factor model

☆The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. We would like to thank Chang-Jin Kim, Christian Murray, and James Hamilton for assistance with programs used in performing the computations. James Hamilton, Marine Carrasco, Robert Gordon, two anonymous referees, and numerous seminar participants also made helpful comments.
*Corresponding author. Tel.: +1 212 720 1948; fax: +1 212 720 1844.
E-mail address: james.kahn@ny.frb.org (J.A. Kahn).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2006.07.008

1. Introduction

Discerning the underlying trend in productivity growth is a perennial goal of both policymakers and researchers. At least since Solow's (1956) pioneering work on long-term growth, economists have understood that sustained productivity growth is the primary source of long-term growth in living standards. It is also important for short-term policy analysis, as any assessment of "output gaps" or growth "speed limits" ultimately derives from some understanding of trend productivity growth. It is widely believed, for example, that the difficulty of detecting a change in this trend contributed significantly to the economic instability of the 1970s, as policymakers were unaware of the slowdown in productivity growth for many years. Only much later were they able to date the slowdown at approximately 1973.[1] This resulted in overestimating potential GDP (at least so the conventional wisdom goes) and setting interest rates too low, and double-digit inflation followed not long after.

Detecting a change in trend productivity growth is difficult because the series is dominated by transitory fluctuations. Postwar quarterly productivity growth in the nonfarm business (hereafter "nonfarm") sector has a standard deviation of 3.8% (annualized), whereas the difference between the high productivity growth years prior to 1973 and the low growth years that followed through the mid-1990s is only about 1.5%. While some of that volatility is attributable to high frequency movements and can be filtered out relatively easily, productivity growth also has a strong cyclical component. It typically declines at the onset of a recession and rises during a recovery, and also leads cyclical movements in output and employment by about three quarters.[2] The timing of cyclical turning points, however, is usually established only well after the fact. Thus it may take years before a change in the long-term trend becomes apparent.

In the late 1990s, attention turned once again to productivity, this time because of speculation that its trend growth rate had picked up. The growth rate of nonfarm output per hour increased by approximately 1% beginning in 1996 relative to the period 1991–1995, and by about 1.3% relative to 1973–1995. The acceleration of productivity put its growth rate during the period from 1996 to 2000 close to where it had been during the most recent period of strong growth, from roughly 1948 to 1973. This provoked a debate over whether we could have expected an extended period of more rapid productivity growth. Robert Gordon (2000), for example, was pessimistic, attributing about half of the acceleration to a "cyclical" effect, and arguing that much of the remainder was confined to the technology sector. Others (e.g. Stiroh, 2002) were more optimistic, finding evidence that productivity growth had spilled over into other sectors through capital deepening.

While the return to faster trend productivity growth is now better established, the difficulty of detecting such changes in trend (i.e. distinguishing permanent and transitory movements in data) is an ongoing challenge. The challenge is particularly problematic in real time. In this paper we attack this problem by drawing on neoclassical growth theory to

[1]See, for example, Sims (2001), who writes that during the 1970s, "unemployment rose and inflation rose because of real disturbances that lowered growth …. Since such 'stagflation' had not occurred before on such a scale, they faced a difficult inference problem, which it took them some years to unravel."

[2]See Estrella (2004).

help identify variables that should be cointegrated with productivity—namely consumption and labor compensation—and consequently can help estimate its trend. We treat that trend as a stochastic process whose mean growth rate has two "regimes," high and low, with some probability of switching between the two at any point in time. We model the business cycle as a second process common to all of the variables in the analysis, also with two regimes of its own, based on the so-called "plucking" model of Friedman (1969, 1993).

There are several advantages to this approach. First, we show that aggregate productivity data alone do not provide as clear or as timely a signal of changes in trend growth as does the joint signal from the series we examine. Second, we do not have to choose break dates *a priori*, as we let the data speak for themselves. Third, the model in principle not only provides information about when regime switches occurred, it also provides estimates of how long the regimes are likely to last. This last property contrasts with even the most sophisticated structural break tests, such as those described by Bai et al. (1998) and Hansen (2001).

Also worth emphasizing is that the reliance on theory allows us to confine our analysis to a low dimensional system of variables and to impose parameter restrictions in the estimation procedure. Thus our approach contrasts with atheoretical applications of factor models that involve a large number of variables or that do not place theory-based restrictions on estimated coefficients.[3] Here there are both advantages and disadvantages. While we may lose information contained in some omitted variables, we also avoid the potential noise from variables for which theory provides less guidance about how they relate to the productivity trend.

In spite of the parsimonious approach, our analysis picks up striking evidence of a switch in the mid-1990s to a higher long-term growth regime, some 25 years after a switch from higher to lower growth in the early 1970s. While these findings themselves may not be surprising in light of subsequent data, our results point to further conclusions as well. First, one could not decisively conclude that there was a return to a higher growth regime on the basis of productivity data alone, even with the inclusion of a second variable to control for the business cycle. Only the corroborating evidence from other cointegrated series can swing the balance strongly in favor of a regime switch. Second, our approach appears effective in detecting changes in trend in real time: In the case of the 1990s, the methodology would have provided a strong signal of a regime switch within roughly six quarters of its actual occurrence according to subsequent data.

The remainder of the paper proceeds as follows. Section 2 motivates the theoretical restrictions based on a variant of the neoclassical growth model. Section 3 describes the statistical model and discusses other related work. Section 4 describes the data and presents parameter estimates for the benchmark model and several alternative specifications. Section 5 discusses the estimated common trends in output per hour and output per capita and their interpretation in light of evidence of changing trends in labor supply. Section 6 examines the benchmark model more closely to gauge the importance of the regime-switching and common factor aspects of the model. Section 7 looks at the performance of the model in real time during the late 1990s. Section 8 concludes.

[3]See for example, Stock and Watson (1989, 2002), Kim and Piger (2002).

2. Implications of the neoclassical growth model

2.1. Background

Over forty years ago, Nicholas Kaldor (1961) established a set of stylized facts about economic growth that have guided empirical researchers ever since. His facts are: (1) labor and capital's income shares are relatively constant; (2) growth rates and real interest rates are relatively constant; (3) the ratio of capital to labor grows over time, and at roughly the same rate as output per hour, so that the capital-output ratio is roughly constant. To these facts, more recent research has added another: (4) measures of work effort show no clear tendency to grow or shrink over time on a per capita basis. The important implication of this additional fact is that wealth and substitution effects roughly offset each other. This means, for example, that a permanent change in the level of labor productivity has no permanent impact on employment.

Of course, closer inspection suggests that none of the above "stylized facts" is literally true. Indeed the premise of much work on U.S. productivity is that productivity growth was systematically higher from 1948 to 1973 than it was over the subsequent 20-plus years. As we shall see, work effort per capita has been anything but stationary since World War II, and there have been large shifts in capital-output ratios. But Kaldor's facts still provide a starting point for modeling economic growth, particularly since there may be reasonable explanations for departures from those facts that do not require discarding the framework that they inspired. We begin in this section with a neoclassical growth model consistent with the Kaldor facts, but then relax all but the first one. We then examine the implications of the generalized model for empirical efforts to assess growth trends.

2.2. A growth model with nonstationary labor supply

In our analysis we allow for exogenous changes in preferences between consumption and leisure to account for long-term movements in work effort (as measured by hours) that show up in the data. Specifically, let C denote aggregate consumption, Y aggregate output, N population (measured in person-hours and growing at rate n), K capital, X effective labor per unit of labor input, and L aggregate labor input (in hours). We also assume that there is a production function

$$Y_t = K_{t-1}^{\alpha}(L_t X_t)^{1-\alpha}. \tag{1}$$

Here X_t, which measures the level of technical progress, is stochastic and exogenous, with an average growth rate of g.

Preferences are defined in terms of a present discounted value of single-period utility

$$U(C_t/N_t, \ell_t) = \Lambda_t \ln\left(C_t/N_t\right) + v(1-\ell_t), \tag{2}$$

where $\ell \equiv L/N$ represents the proportion of available hours devoted to work. The marginal rate of substitution between consumption and leisure is $\Lambda^{-1}(C/N)v'(1-\ell)$, where v is a concave differentiable function, v' is strictly decreasing, and Λ is a taste parameter that can shift over time. Note that while Λ is modeled as a preference shock, it could reflect taxation or other labor market distortions (see Mulligan, 2002), as well as demographic shifts.

We will also allow Λ to be non-stationary. For the sake of exposition we will specify it as a unit root process with zero drift, though it could also be a deterministic function of time,

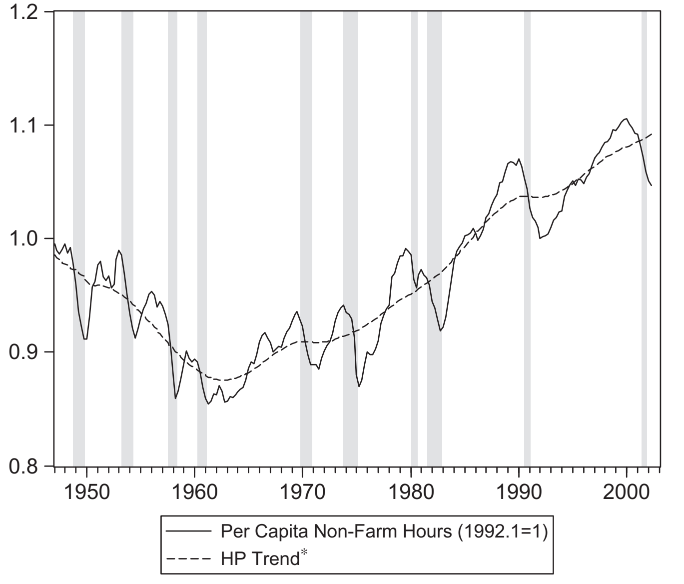

*Obtained from log(hours) using a parameter of 9600, then subtracting log(population).

Fig. 1. Hours of work per capita in the postwar U.S.

or a combination of the two. This is in recognition of the fact that there is significant low-frequency variation of work effort in postwar U.S. data, as seen in the behavior of per capita hours in the nonfarm sector since 1947 (Fig. 1).[4] Apart from the large middle frequency fluctuations associated with business cycles, there are clear secular changes. There was a decline of roughly 15% between the end of World War II and the early 1960s, followed by an increase of about 20% from the mid-1960s to the present. Studies that have assumed that aggregate per capita output, along with consumption and investment, have the same permanent component as (i.e. are cointegrated with) labor productivity implicitly assume that hours per capita is a stationary time series.[5] We argue below that the stochastic trend in per capita output is better described as two separate trend components, one demographic (i.e. labor supply), the other technological (i.e. labor demand), and that by doing so we are able to identify regime shifts in the latter that otherwise would be obscured by movements in the former.

We assume that the economy evolves as if a planner solves the following problem:

$$\max E_t\left\{\sum_{t=0}^{\infty}\beta^t U(C_t/N_t,\ell_t)\right\}, \tag{3}$$

[4]Per capita variables are obtained by dividing by the total resident population, averaging the monthly data to obtain a quarterly series, and extrapolating to extend the series beyond 2001. Note that the share of nonfarm to total employment has varied only slightly over the sample period and is not responsible for the low frequency movements visible in Fig. 1.

[5]For example, Bai et al. (1998).

subject to

$$C_t + I_t \leqslant K_{t-1}^{\alpha}(L_t X_t)^{1-\alpha}, \tag{4}$$

$$K_t \leqslant I_t + (1-\delta)K_{t-1}, \tag{5}$$

where I_t is investment and β is a discount factor.

Let $Z_t = L_t X_t$. Also, let W denote labor compensation, and note that $W_t = (1-\alpha)K_{t-1}^{\alpha} Z_t^{1-\alpha}$. It is straightforward to show that the variables $c_t \equiv C_t/Z_t$, $y_t \equiv Y_t/Z_t$, $w_t \equiv W_t/Z_t$, $k_{t-1} \equiv K_{t-1}/Z_t$, $i_t \equiv I_t/Z_t$ and $h_t \equiv h(\ell_t, \Lambda_t) = \Lambda_t^{-1} v'(1-\ell_t)\ell_t$, are stationary along a balanced growth path. Thus the economy will grow in per hour terms on average at the rate g, and Y/L, C/L, W/L, I/L and K/L will have a common stochastic trend X. L (or L/N), on the other hand, will have its own stochastic trend that is a function of Λ_t. Aggregate or per capita quantities such as Y or Y/N will have a common stochastic trend made up of two components: the trend in L (arising from the nonstationarity of the labor supply trend Λ) and the trend in X.

These implications regarding common trends are robust to generalizations of the model provided they are consistent with the same balanced growth path, i.e. so long as they result in only transitory deviations from that path. For example, the observables may be measured with (stationary) errors, or may have transitory dynamics that reflect imperfect information, adjustment costs, or other rigidities. So long as such deviations (which we will allow for in the estimation) are transitory, the ratios noted above should still be cointegrated. Our approach is essentially to remain agnostic about short-term adjustment dynamics, and rely only on the more robust implications regarding cointegration for the purpose of estimating changes in the trend.

On the other hand, there are other growth models that have different cointegration implications. One popular generalization of the neoclassical growth model allows for vintage capital or embodied technical progress (e.g. Greenwood et al., 1997). In the empirical section we will examine a number of alternative specifications to check the robustness of our findings.

We should note that our focus on labor productivity rather than on total factor productivity (TFP) is intentional. Anything that *permanently* raises output per hour will enter our estimated "technology" component, whether it be capital deepening, growth in human capital, or TFP. Of course growth theory suggests that capital deepening is unlikely to be an *independent* contributor to sustained growth. Rather, it is a symptom of underlying technological progress and/or growth in human capital. Thus, for example, the capital deepening of the late 1990s, much of which can be attributed to computer and related high-tech investment, ultimately reflects TFP in the sectors that produce that equipment.[6]

One final issue relates to our treatment of Λ as non-stationary. There has been considerable debate in the economics literature on whether per capita hours (ℓ) is better characterized as having a unit root or as stationary (perhaps around a trend).[7] We are not taking a stand on that question, as the non-stationarity of Λ could merely reflect, for example, permanent changes due to demographics, or to anything that could be handled by a deterministic time trend. As we shall see, controlling for the apparent trend in $\log \Lambda$ turns out to be important in enabling us to detect changes in trend productivity growth. At the same time, we should point out that although a bounded variable cannot literally

[6]See, for example, Gordon (2000), Stiroh (2002).

[7]See, for example, Francis and Ramey (2005), Christiano et al. (2004).

follow a simple linear unit root process without violating its bounds, there is nothing preventing $\log(\Lambda)$ from following such a process. For example, if $v(\ell) = \log(\ell)$, then $\log(\ell/(1-\ell))$ is cointegrated with $\log \Lambda$ and ranges over the entire real line, while ℓ itself remains in the unit interval. A permanent shift in $\log \Lambda$ changes $\log \ell$ by approximately $(1-\ell)^2/\ell \times d(\log \Lambda)$.

2.3. Summary

Balanced growth implies, by definition, that quantities such as output, consumption, capital, investment, and labor income all have a common trend. Growth models with this property typically assume that per capita labor input is stationary, implying that the common trend in these variables results solely from technical progress and population growth. We have shown that this is a questionable assumption for the postwar United States, and therefore that these quantities should be normalized by aggregate labor input if they are to reflect the technology trend alone. Low frequency movements in per capita hours of work are presumably driven by fiscal policy and demographics, factors reasonably assumed to be unrelated to technology.

In effect, the theory tells us that as far as low frequency behavior is concerned, we can divide output per capita Y/N into Y/L (labor demand) and L/N (labor supply). We refer to Y/L as labor demand because with Cobb–Douglas technology it is proportional to the marginal product of labor (MPL). Standard assumptions about preferences and technology imply that long-run labor demand is horizontal (though it may shift over time), while long-run labor supply is vertical. Thus a permanent shift in labor supply (as represented by L/N) does not lead in the long run to any change in the other quantity (MPL). Similarly, a permanent increase in Y/L (i.e. in the MPL), should not have permanent impact on L/N.

3. A common factor model

3.1. The regime-switching dynamic factor model

Our estimation strategy draws upon the regime-switching dynamic factor model recently proposed by Kim and Murray (2002) and Kim and Piger (2002), among others. The essence of this approach is to examine a number of related economic time series and to use their comovements to identify two shared factors: a common permanent component and a common transitory component. In addition, we follow these authors in allowing for regime changes in both components. We are also motivated by similar considerations as in Rotemberg (2003), who argues that the technology trend should be smooth (or, at least smoother than measured productivity) and relatively distinct from cyclical movements around the trend.

The regime-switching specification has several attractive features. First, the permanent component allows us to account for sustained changes in trend growth without making the growth process itself nonstationary.[8] Second, the transitory component allows for asymmetries in business cycles that others have claimed to be important.[9] The regime changes in the transitory component capture the idea proposed in Friedman's (1969, 1993)

[8]Cf. Cogley (2005), whose drifting parameters have a unit root.

[9]See, for example, Kim and Piger (2002), Kim et al. (2002), and Beaudry and Koop (1993).

"plucking model" model that economic fluctuations are largely permanent during expansions and transitory during recessions. Finally, the regime-switching specification is a straightforward way of estimating both the timing and expected duration of periodic changes in the processes generating the two components.

This approach also has some limitations, of course. It requires specifying a small number of discrete states for continuous economic variables such as trend productivity growth. Nonetheless such models have been shown to be able to capture features of the data that standard linear models can not (see, for example, Hamilton, 2004), and we will argue that trend productivity growth appears especially amenable to this approach. Ultimately these sorts of modeling choices cannot be justified on *a priori* grounds, but rather on the insights they provide and on how the models perform both in- and out-of-sample.

As we have seen, the neoclassical growth model provides the basis for treating a number of growing time series as having a common stochastic trend. If these series have some independent sources of measurement error or other transitory deviations, there may be some gain to using all three to extract a more precise estimate of the common permanent component.[10] Consequently we adopt a multivariate, common factor approach to estimate a single permanent component from multiple series. We also use that approach to estimate a common transitory component, although without the benefit of any theoretical restrictions.

Following Kim and Murray (2002), we can describe the regime-switching dynamic factor model as follows.[11] Suppose we consider J time series indexed by i. Let Q_{it} denote the logarithm of the ith individual time series. It is assumed that the movements in each series are governed by the following process:

$$Q_{it} = \gamma_i X_t + \lambda_i x_t + z_{it}, \quad i = 1, \ldots, J, \tag{6}$$

where X_t denotes a permanent component that is common to all series, x_t denotes a common transitory component, and z_{it} is an idiosyncratic error term. The parameter γ_i (the permanent "factor loading") indicates the extent to which the series moves with the common permanent component. Similarly, the parameter λ_i indicates the extent to which the series is affected by the transitory component.

The common permanent component is assumed to be difference stationary, but subject to the type of regime-switching proposed by Hamilton (1989) in which there are periodic shifts in its growth rate:

$$\Delta X_t = \mu(S_{1t}) + \phi_1 \Delta X_{t-1} + \cdots + \phi_p \Delta X_{t-p} + \upsilon_t, \upsilon_t \sim iid\ \ N(0,1) \tag{7}$$

$$\mu(S_{1t}) = \begin{cases} \mu_0 & \text{if} \quad S_{1t} = 0, \\ \mu_1 & \text{if} \quad S_{1t} = 1, \end{cases} \tag{8}$$

$$\Pr[S_{1t} = 0 | S_{1,t-1} = 0] = q_1, \quad \Pr[S_{1t} = 1 | S_{1,t-1} = 1] = p_1, \tag{9}$$

[10] This is not to say that the information is completely independent, and indeed the source of errors in one series may be present in the other series as well. For example, an inaccurate price deflator could result in common mismeasurement across multiple series. Nonetheless, the theory suggests that considering these series together may provide better information about underlying trends than consideration of any of them in isolation.

[11] Additional details are in the Appendix. Kim and Murray (2002) discuss how the regime-switching dynamic factor model can be cast in a state-space representation and estimated using the Kalman filter. We follow the approach of Kim (1994) and use an approximation to construct the likelihood function.

where S_{1t} is an index of the regime for the common permanent component. The transition probabilities q_1 and p_1 indicate the likelihood of remaining in the same regime. Under these assumptions, the common permanent component X_t grows at the rate $\mu_0/(1-\phi_1-\cdots-\phi_p)$ when $S_{1t}=0$, and at the rate $\mu_1/(1-\phi_1-\cdots-\phi_p)$ when $S_{1t}=1$.

The common transitory component x_t is stationary in levels, but also subject to regime-switching:

$$x_t = \tau(S_{2t}) + \phi_1^* x_{t-1} + \phi_2^* x_{t-2} + \cdots + \phi_p^* x_{t-p} + \varepsilon_t, \varepsilon_t \sim iid \ N(0,1), \tag{10}$$

$$\tau(S_{2t}) = \begin{cases} 0 & \text{if} \quad S_{2t}=0, \\ \tau & \text{if} \quad S_{2t}=1, \end{cases} \tag{11}$$

$$\Pr[S_{2t}=0|S_{2,t-1}=0]=q_2, \quad \Pr[S_{2t}=1|S_{2,t-1}=1]=p_2, \tag{12}$$

where S_{2t} is an index of the regime for the common transitory component, with transition probabilities q_2 and p_2. The parameter τ represents the size of the "pluck," with $\tau<0$ implying that the common transitory component is plucked down when $S_{2t}=1$. Note that because the specification does not pin down the level of x, we are free to normalize by setting $\tau(0)=0$.

The permanent and transitory regimes are assumed to be independent of each other. While the two regimes are not directly observable, it is nevertheless possible to estimate the parameters of the model and to extract estimates of the common components. An important byproduct from the estimation procedure is that we can draw inferences about the likelihood that each common component is in a specific regime at a particular date. The restriction of unit variance for the error terms of the two processes is an identifying restriction, since X and x are of indeterminate scale.

Finally, the idiosyncratic components are assumed to have the following structure:

$$z_{it} = \psi_{i1} z_{i,t-1} + \psi_{i2} z_{i,t-2} + \cdots + \psi_{ip} z_{i,t-p} + \eta_{it}, \eta_{it} \sim iid \ N(0,\sigma_i^2), \quad i=1,\ldots,J, \tag{13}$$

where all innovations in the model are assumed to be mutually and serially uncorrelated at all leads and lags. We assume that z_{it} is stationary, i.e. the roots of $1-\Psi_i(L)$ all lie inside the unit circle, so these are transitory shocks to the levels of the variables.

To relate all of this back to the growth model from Section 2 of the paper, the permanent component X corresponds to the stochastic trend term for technology from the growth model, which we saw is common to Y/L, W/L, C/L, I/L and K/L. Because capital stock measures are only available annually, and investment is a much more volatile series than the others, in our analysis we will focus on just the first three of these variables, which we label 1, 2, and 3, respectively. The theory implies, therefore, that the factor loadings on the permanent component should satisfy $\gamma_1=\gamma_2=\gamma_3$. The transitory component x reflects the direct impact of transitory disturbances as well as transition dynamics from all shocks, but only to the extent they are linearly related across all of the series. The idiosyncratic component includes what is left of the transitory movements after the common component is subtracted. It will include measurement error and model noise. For example, a literal reading of the model is that $W/L=(1-\alpha)Y/L$, so the two series should be perfectly correlated. But W and Y are measured (to some extent) independently, and with error, and moreover the assumption of a Cobb–Douglas production function with constant parameters is undoubtedly not literally accurate. These factors will also contribute to idiosyncratic variation.

3.2. Relationship to previous work

A number of recent papers have focused on the trend in U.S. productivity growth and the related issue of changes in its behavior. Hansen (2001) applies endogeneous structural break tests to post-war data. His results provide evidence of a break in the behavior of productivity growth in the 1990s, with much weaker evidence of a second break in the series. However, Hansen's conclusions are based on a univariate analysis. Moreover, his measure of productivity is for the manufacturing/durables sector rather than the more conventional measure for the nonfarm sector. Estrella (2004) uses spectral methods to decompose variation in productivity growth into short-, medium-, and long-term components. He finds an economically significant increase in the long-term, or low frequency component in the late 1990s, but (in contrast to our findings) predicts a relatively fast reversion to the long-run average growth rate.

The work of French (2001) and Roberts (2000) is more closely related to this study. French estimates a Markov switching model for total factor productivity using data from 1959–1999. His findings document a shift in trend growth in 1973–74, and also suggest a possible slight upward shift in trend growth toward the end of the sample. Roberts employs a multivariate Kalman filter procedure that allows trend productivity growth to vary over the period 1960–2000.[12] His results indicate that trend productivity growth declined steadily from the early 1960s to the middle 1970s, and then remained at a fairly steady level through the early 1990s. Since then, trend productivity growth displayed a steady increase through 2000.

While our approach is similar in some respects to the approaches adopted by French and Roberts, there are several important differences. We extend the empirical framework of French by bringing additional information from other related variables to bear on the detection of discrete shifts in the behavior of productivity growth. In addition, we allow both the trend and cyclical components of productivity growth to be subject to regime shifts. In contrast to Roberts, we identify the split between trend and cycle through direct estimation of their shared influences on the data rather than through a decomposition that relies on the specification and estimation of a set of auxiliary equations. More importantly, we will argue that Roberts's use of per capita data (instead of data expressed on a per hour basis) results in his estimate of trend productivity growth being a mixture of the trends associated with technology and labor supply.

Regarding the previously cited work by Kim and Murray (2002), Kim and Piger (2002), Kim and Nelson (1999), and Kim et al. (2002), there are several important differences. Kim and Murray adopt a less structural approach, with a set of non-cointegrated indicator variables. Kim and Nelson focus on the plucking model of business cycles, and do not have regime-switching in the drift term, only in the transitory component. Kim and Piger constrain the permanent component to equal consumption of nondurables and services, and also constrain the regime switches to be the same for the permanent and transitory components. All of these authors look at aggregate or per capita data. Consequently, they

[12] French (2001) and Roberts (2000) assume that the logarithm of labor productivity is $I(2)$ to allow for changes in the growth rate. When we applied Dickey-Fuller (1979) tests to productivity growth, we strongly rejected the hypothesis of a unit root. Roberts (2000) and Staiger et al. (2001) have noted, however, that unit root tests have a high false-rejection rate when applied to series such as productivity growth. Specifically, if the estimated variance of the growth rate process is small, then a series can spuriously appear to be stationary.

do not filter out movements in labor supply, which (as we will argue) leads them to find substantial cyclical movement in their permanent components.

Finally, Cogley (2005) estimates trend GDP growth using information contained in consumption (or more precisely, the GDP-consumption ratio). Cogley uses Bayesian methods to estimate parameter drift in a bivariate VAR in the two variables. He finds a modest but transitory increase in trend GDP growth in the 1990s. His findings are not directly comparable to those in this paper, as he does not distinguish between changes in trend productivity growth and growth in hours of work, a distinction we argue is crucial for sorting out the structural shifts underlying the data. Nonetheless his use of consumption data to help identify the trend in GDP growth is in the spirit of our multivariate approach.

4. Results

4.1. Data

Our data consist of quarterly observations of nonfarm sector output, labor productivity, real compensation per hour (nominal compensation relative to the nonfarm output deflator), and hours of work. The nonfarm sector was chosen because of the availability of consistent data for all of these series. We also use aggregate data on real consumption expenditures. While a series that converts expenditures on durables to service flows would be preferable, at the time of these computations such a series was not available for the whole sample period, and where it was, it exhibited very similar behavior to total consumer expenditures. Another issue with using this consumption series is that it is for the entire U.S. economy, whereas the other series represent the nonfarm sector only. This will only create a problem if there are significantly different trends, but the results did not appear to be sensitive to these choices. Unless stated otherwise, all variables are in logarithms, multiplied by 100, with first differences interpreted as quarterly growth rates in percent. For calculating per capita quantities we used the resident population (interpolated from annual to quarterly).

As we have seen, balanced growth implies common trends across a large number of aggregates. When these aggregates are divided by labor hours, we can interpret the resulting trend as the technology factor X, which the model implies is orthogonal to labor supply. There is no analogous set of variables that can be used to help estimate a stochastic trend in labor supply—if L/N is nonstationary, theory does not readily offer other variables with which it should be cointegrated.[13] Thus the common factor model will not be of direct use for estimating it. Higher frequency movements in hours of work would, however, presumably be useful for capturing the transitory component. So we do want to include an hours measure in our system, but without having to estimate a second stochastic trend for labor supply. Thus we detrend the hours series using the Hodrick-Prescott (1997) filter and include it in our system with a zero loading on the permanent component.[14] Our benchmark specification, therefore, has $Q_1 = Y/L$, $Q_2 = W/L$, $Q_3 = C/L$, and $Q_4 = \hat{L}$, where the "^" indicates the H-P filtered series.

[13]There could be labor supply related variables in addition to population that are cointegrated with L. We do not explore this here, but see Mulligan (2002).

[14]Using the first difference of hours (in logs) rather than the H-P filtered series yielded very similar results.

Table 1
Unrestricted cointegration rank test

Hypothesized number of cointegrating equations	Eigenvalue	Trace statistic[a]	p-values
None**	$\hat{\lambda}_1 = 0.148$	64.73	0.0000
At most 1**	$\hat{\lambda}_2 = 0.113$	29.18	0.0003
At most 2	$\hat{\lambda}_3 = 0.012$	2.57	0.1088
Hypothesized number of cointegrating equations	**Eigenvalue**	**Maximal eigenvalue statistic[b]**	**p-values**
None**	$\hat{\lambda}_1 = 0.148$	35.54	0.0003
At most 1**	$\hat{\lambda}_2 = 0.113$	26.61	0.0004
At most 2	$\hat{\lambda}_3 = 0.012$	2.57	0.1088

Note: The sample is 1947:Q1–2002:Q4.
*Denotes rejection of the hypothesis at the 5% level.
**Denotes rejection of the hypothesis at the 1% level.
[a] $\lambda_{\text{trace}} = -T\sum \ln(1-\hat{\lambda}_i)$
[b] $\lambda_{\max} = -T\ln(1-\hat{\lambda}_{r+1})$.

To examine the cointegration properties of Y/L, W/L, and C/L more rigorously we conducted multivariate unit root tests based on the procedure developed by Johansen (1991, 1995). Table 1 describes the results, based on quarterly data over the sample period 1947:Q1 to 2002:Q4, including a constant in the cointegrating relationship. The combination of trace and maximal eigenvalue tests suggests that there are two cointegrating equations, implying a single common trend, as the theory suggests. Letting $Q = (Q_1\ Q_2\ Q_3)$, and letting $\beta^i = (\beta_1^i\ \beta_2^i\ \beta_3^i)$, $i = 1,2$ denote the cointegrating vectors, we estimated the two cointegrating relationships $Q'\beta^1$ and $Q'\beta^2$ and found:

$$\beta^1 = \{1.0 \quad 0.0 \quad -1.018\}, \quad \beta^2 = \{0.0 \quad 1.0 \quad -0.984\}, \tag{14}$$

where we normalized the equations with respect to β_1^1 and β_2^2. Because the estimates of β_3 differ from -1, we also conducted tests of the null hypothesis that the two cointegrating equations are (1.0 0.0 -1.0) and (0.0 1.0 -1.0). The calculated value of the χ^2 statistic with two degrees of freedom is 8.04, with an associated p-value of 0.018. Thus, although the results confirm the theory qualitatively, the quantitative implication that $\beta_3 = -1$ in both vectors is rejected. This does not necessarily mean, however, that we will reject $\gamma_1 = \gamma_2 = \gamma_3$ in (6), which is the real implication of the theory, and one that we will test when we estimate the benchmark model.

Initial estimates of more general specifications suggested that the common permanent component should include one lagged value of ΔX_t, and the common transitory component should include two lagged values of x_t. The common idiosyncratic component should include one lagged value of z_{it} for output per hour, labor compensation per hour, and consumption normalized by hours, and two lags for detrended hours. We restricted the estimated factor loadings on the permanent component for productivity, real compensation per hour, and consumption per hour to be equal (i.e. we set $\gamma_1 = \gamma_2 = \gamma_3$), and set the value of the permanent factor loading for detrended hours, γ_4, equal to zero.

We consider a number of alternative specifications, but for the moment focus on two. In one, we eliminate the regime shift in the transitory component, thereby imposing symmetry on business cycles. We will refer to this as the "no pluck" specification. In the other, we use per capita rather than per hour variables, i.e. we set $Q_1 = Y/N$, $Q_2 = W/N$, $Q_3 = C/N$, where N is the U.S. resident population.[15] We will refer to this as the "per capita" model, as opposed to the benchmark or "per hour" specification. As discussed in more detail below, if L/N were stationary, then the per capita and per hour specifications should yield similar estimates of the permanent component.

One final modeling issue relates to the synchronization of the data. The three trending variables were selected primarily on the basis of their having a common permanent component. They are not, however, necessarily "coincident indicators" with respect to the transitory component. In theory it would be possible to allow for a more general lead/lag structure in our system, but this would greatly increase the number of parameters to estimate. As an alternative, we first examined the cross-correlations of the four series in the benchmark specification. We found that the first two variables (productivity and labor compensation per hour) both tended to lead the other two series by about three quarters. To capture this asynchronization in the estimation we lagged Q_{1t} and Q_{2t} by three quarters in the system described above.[16]

4.2. Parameter estimates

The first column of Table 2 provides the parameter estimates for our benchmark model with the four variables as described above.[17] The data cover 1947:Q1–2002:Q4. Because output per hour and labor compensation per hour are lagged three quarters, their growth rates run from 1947:Q2–2002:Q1, while the growth rates of the consumption and detrended hours variables run from 1948:Q1–2002:Q4.

The model yields precise estimates of most of the parameters of interest: The factor loadings on both the permanent and transitory components, the transition probabilities, and the shift parameters associated with the regimes (μ_0, μ_1, τ) all enter significantly. The difference between the high- and low-growth regimes works out to be $\gamma(\mu_0-\mu_1)/(1-\phi) = 0.353$. This corresponds to approximately 1.4% on an annualized basis, very close to the difference between the 1947–73 and 1973–96 growth rates of productivity.

The transition probabilities for the permanent regimes imply an expected duration of $1/(1-q_1) = 100$ quarters for the high-growth regime, and 59 quarters for the low-growth regime. They also imply that the unconditional probability of being in the high growth regime is $(1-p_1)/(2-p_1-q_1) = 0.63$, suggesting that the economy would be in the high growth regime on the order of 63% of the time. For the period from 1947 to 2002, this would be 35 years, consistent with the view that the postwar high growth regimes were pre-1973 and post-1996. The AR coefficient on the permanent component is estimated to be -0.245, suggesting that growth innovations in one quarter tend to get partially offset in the following one.

[15]We obtained very similar results using the population over age 15.

[16]This procedure results in the last three observations of variables 1 and 2 not being used to estimate the parameters of the model. We do, however, incorporate them in the period-by-period assessments of the state variables and regime probabilities, as described below.

[17]Estimates of the idionsyncratic variances are omitted from Table 2 for the sake of brevity.

Table 2
Estimation of model

Coefficient	4 equation system	*No pluck system*	Per capita system	2 equation system
q_1	0.990 (0.011)	0.991 (0.011)	***	0.993 (0.008)
p_1	0.983 (0.015)	0.983 (0.015)	***	0.985 (0.013)
q_2	0.984 (0.014)	–	0.972 (0.017)	0.959 (0.022)
p_2	0.515 (0.133)	–	0.553 (0.171)	0.551 (0.153)
ϕ	−0.245 (0.118)	−0.250 (0.118)	0.543 (0.090)	−0.735 (0.151)
ϕ_1^*	1.416 (0.087)	1.466 (0.068)	1.392 (0.069)	1.370 (0.133)
ϕ_2^*	−0.557 (0.079)	−0.601 (0.067)	-0.523 (0.062)	−0.567 (0.113)
ψ_{11}	0.892 (0.040)	0.880 (0.045)	0.918 (0.034)	0.873 (0.051)
ψ_{21}	0.896 (0.046)	0.907 (0.043)	0.918 (0.033)	–
ψ_{31}	−0.585 (0.104)	−0.588 (0.102)	−0.195 (0.163)	–
ψ_{41}	1.440 (0.053)	1.448 (0.053)	1.555 (0.104)	1.438 (0.068)
ψ_{42}	−0.519 (0.038)	−0.524 (0.039)	−0.605 (0.081)	−0.517 (0.049)
γ	0.319 (0.034)	0.318 (0.034)	0.446 (0.044)	0.263 (0.124)
λ_1	0.239 (0.047)	0.267 (0.051)	0.643 (0.059)	0.091 (0.084)
λ_2	0.133 (0.030)	0.146 (0.033)	0.617 (0.053)	–
λ_3	−0.486 (0.045)	−0.544 (0.044)	0.219 (0.036)	–
λ_4	0.410 (0.044)	0.453 (0.047)	0.549 (0.043)	0.229 (0.206)
μ_0	0.596 (0.130)	0.582 (0.130)	−0.005 (0.049)	0.755 (0.419)
μ_1	−0.782 (0.165)	−0.792 (0.168)	***	−1.211 (0.689)
τ	−2.688 (0.791)	–	−2.640 (0.482)	−6.013 (5.229)

Note: The estimation also produces estimates of the variances of the idiosyncratic errors, not reported here. *** indicates parameters that could not be estimated because they are not identified. Standard errors are in parentheses.

The transitory process is estimated to be a "hump-shaped" autoregressive process typical of the business cycle (see, e.g., Blanchard, 1981), but with a statistically significant negative pluck, i.e. a relatively short-lived reduction in the level of the transitory component presumably associated with recessions. The magnitude of this downward shift is related to $\lambda_i\tau$. Productivity, for example, would decline by $\lambda_1\tau = 0.642$, which corresponds to about 2.6% annualized. The expected duration of the pluck regime is given by $1/(1-p_2)$, or 2.06 quarters. The transition probabilities imply that the pluck regime occurs less than 4% of the time. Since about one-sixth of the quarters since 1947 have been associated with recessions, and the average recession has lasted roughly 3.5 quarters, these results suggest that not all recessions will coincide with transitory regime shifts.

The four λ_i coefficients (the factor loadings on the transitory process) are all estimated very precisely, and have the expected signs. Note that λ_3 the factor loading for consumption/hours, is negative, reflecting the fact that hours are more cyclical than consumption. The positive estimates for the factor loadings on real hourly compensation and productivity indicate that those two variables are positively related to the transitory component, albeit leading by three quarters relative to the other two variables due to their entering the system lagged by three quarters.

We tested the restriction $\gamma_1 = \gamma_2 = \gamma_3$ by estimating the benchmark model without these constraints and then conducting a likelihood ratio test. The test statistic has a value of 1.886, and is asymptotically distributed as χ^2 with two degrees of freedom. The critical value for rejecting the hypothesis at a 10% significance level is 4.61. This suggests that the basic theoretical implications with respect to the common permanent component are confirmed by the data, notwithstanding the estimates of the cointegrating vectors described in Section 4.1.

We also estimated the model without regime-switching in the transitory component, i.e. restricting $\tau = 0$. These results are shown in the second column of Table 2. The estimates of the other parameters turn out to be very close to the benchmark estimates, suggesting that at least insofar as estimating the permanent component is concerned, whether one takes into account nonlinearities in the business cycle does not matter much.

4.3. Growth regime assessments

Before further describing the permanent and transitory components of productivity growth, it is instructive to examine the inferred probability of being in the high- or low-growth state over time. There are two common approaches for examining regime probabilities. The first is in "real time": At each point in time, using only data through that point in time, what probability would one have assigned to being in the high-growth state? This question is usually answered with the so-called "zero-lag" or "unsmoothed" estimates of the regime (see Hamilton, 1994). The second way is retrospectively: Given all the data available through 2002:Q4, what can we say looking back over time about the likelihood of being in the high growth state? (Obviously the two assessments coincide at the end, i.e. as of 2002:Q4.) This is Hamilton's (1994) "full-sample smoother."

Two practical considerations arise with the real time approach. First, there is the matter of data revisions. The historical series we have today have been revised numerous times over the years, so the data truncated at some date do not correspond to what anyone would actually have known. Second, it may not be practical to have rolling estimates of parameters, especially if regime shifts occur infrequently. Over a shorter sample few if any regime switches will be observed, making it problematic to estimate transition probabilities with any degree of precision. Other parameters will also be difficult to estimate precisely as well. So for the moment we focus on the restrospective approach, but return to true real-time analysis in Section 5.

There is also a technical issue related to the timing of the data series. Since Q_1 and Q_2 are lagged by three periods, if we were simply to assess the state vector given data through a given observation number, we would be ignoring the most recent three data points for those two series. Specifically, a standard application of the Kalman filter at date t would yield the following assessment of the state vector given data through period t:

$$\hat{\xi}_{t|t} = E(\xi_t | Q_{1t}, Q_{2t}, Q_{3t}, Q_{4t}; Q_{1t-1}, Q_{2t-1}, Q_{3t-1}, Q_{4t-1}; \ldots). \tag{15}$$

In terms of the actual underlying data in our study, however, this would be expressed as

$$E(\xi_t | \tilde{Q}_{1t-3}, \tilde{Q}_{2t-3}, \tilde{Q}_{3t}, \tilde{Q}_{4t}; \tilde{Q}_{1t-4}, \tilde{Q}_{2t-4}, \tilde{Q}_{3t-1}, \tilde{Q}_{4t-1}; \ldots), \tag{16}$$

where the "$\sim$" denotes the variable indexed by its true time period. Fortunately it is relatively straightforward to "partially update" the state vector and regime assessments. This involves conditioning on the three subsequent observations of Q_1 and Q_2 at each

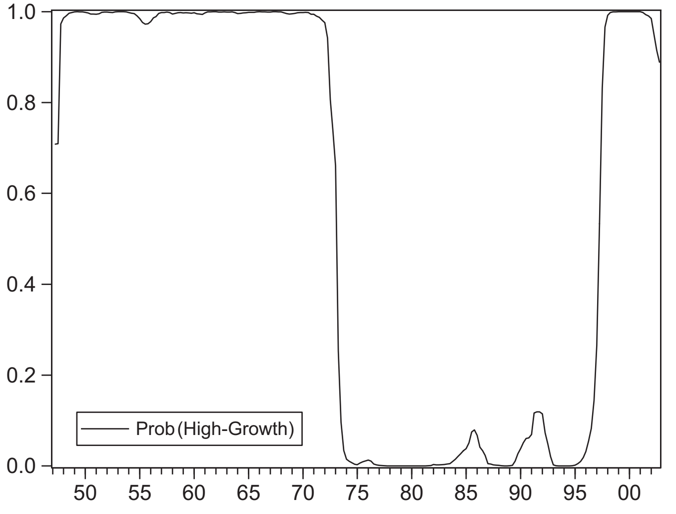

Fig. 2. Growth regimes.

point in time to obtain

$$\begin{aligned}\hat{\xi}_{t|t+3'} &= E(\xi_t|Q_{1t+3}, Q_{2t+3}, Q_{3t}, Q_{4t}; Q_{1t+2}, Q_{2t+2}, Q_{3t-1}, Q_{4t-1}; \ldots) \\ &= E(\xi_t|\tilde{Q}_{1t}, \tilde{Q}_{2t}, \tilde{Q}_{3t}, \tilde{Q}_{4t}; \tilde{Q}_{1t-1}, \tilde{Q}_{2t-1}, \tilde{Q}_{3t-1}, \tilde{Q}_{4t-1}; \ldots).\end{aligned} \tag{17}$$

Here we use $t+3'$ to denote Q_1 and Q_2 observed through $t+3$, and Q_3 and Q_4 observed through time t. This simply undoes the staggering so that the information set is appropriately aligned.[18]

These full-sample regime assessments are plotted in Fig. 2. The vertical axis is the probability of being in the high-growth regime. This retrospective assessment presents a very clear picture: The economy was in a high-growth state until the early 1970s, followed by a low-growth regime that lasted until the second half of the 1990s, at which point there was a return to the high-growth regime that persisted through the end of the sample. It should be noted that the full-sample probabilities converge to the zero-lag probabilities at the end of the sample, as there are no subsequent observations to inform the assessment. Consequently they tend to reflect this greater uncertainty by being farther from one or zero. Aside from this, the lack of ambiguity in the regime assessments is striking. For example, after being less than 0.05 in 1996:Q1, the high-growth regime probability surpassed 0.95 in 1997:Q4, and did not subsequently fall below 0.95 until the last few observations.

Thus we find (given data through 2002) that the break in trend occurred sometime in 1997—a year or so later than the conventional wisdom, but broadly consistent with it. Nonetheless it is reasonable to ask why the procedure selects 1997 as the year of the regime shift and not 1995 or 1996, despite three very strong quarters of growth beginning in 1995:Q4. First, note that short stretches of strong growth are not unusual, and that the three strong quarters beginning in 1995:Q4 were followed by three weak quarters. Second, the strong quarters came in the wake of an economic slowdown, a time when productivity

[18]Additional details on the partial updating procedure are provided in the Appendix.

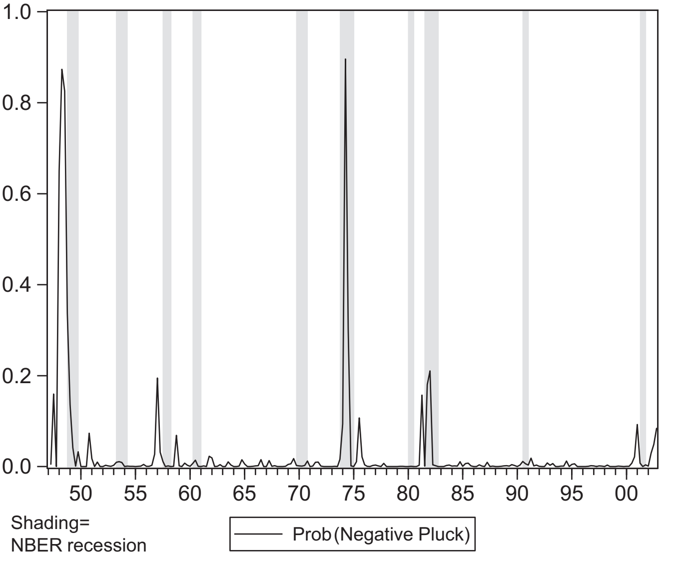

Fig. 3. Regime shifts in the transitory component.

typically accelerates. Beginning in 1991:Q2, for example, productivity growth was in excess of 2.5% for seven consecutive quarters, and averaged over 4%. Similarly strong growth was observed coming out of the 1981–82 recession. While there was no recession in 1994-95, GDP growth did slow substantially. Thus even with hindsight, the three strong quarters in late 1995 and early 1996 do not suggest a change in trend. The consumption and labor compensation variables followed very similar patterns. It was only beginning in 1997:Q2 that consistently strong numbers began to appear that differed from the normal cyclical pattern, given that the economy was well into a strong expansion.

For completeness, we also examine the assessment of the transitory regimes. Some (e.g. Beaudry and Koop, 1993; Kim and Piger, 2002) have argued for the importance of taking business cycle asymmetries into account. Consistent with the evidence from Table 2 discussed earlier, however, Fig. 3 suggests that asymmetries may not be very important, at least in data that are detrended by hours of work. We find that while the more prominent spikes all coincide with NBER-defined recessions, in only two cases does the probability of a negative pluck exceed 0.5. One reason that regime-switching may not appear so significant in the transitory component is that the magnitude of the pluck for these series may vary across recessions. Note, for example, that the 1990–91 recession does not register in this picture. The idea of a "pluck" is a sharp downturn followed by an equally sharp recovery sufficient to get the economy back to trend. The 1990–91 recession was characterized by a relatively mild downturn followed by an unusually slow and gradual recovery.

5. Discussion: properties of the permanent component

The common permanent component, as depicted with labor productivity in Fig. 4, clearly indicates changes in trend in the early 1970s and mid-1990s, with little evident

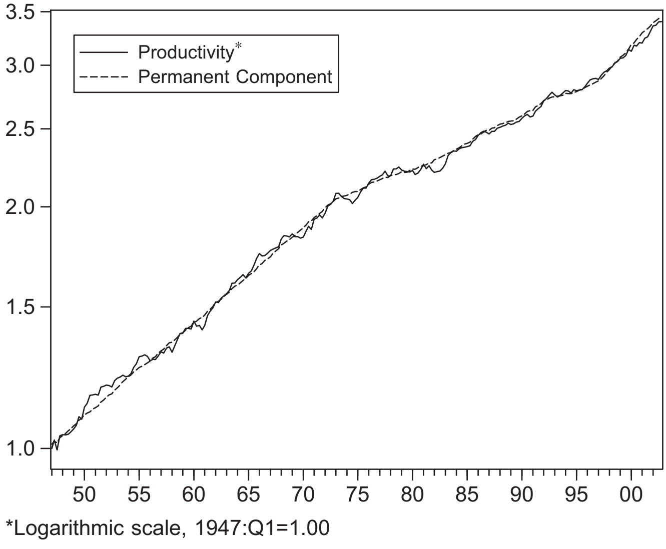

Fig. 4. Trend productivity.

cyclicality. Its standard deviation is only one-fourth that of productivity growth. This contrasts with many recent estimates of the permanent component of output. For example, Kim and Murray's (2002) estimated permanent aggregate component shows substantial downward movement during recessions. Kim and Piger (2002) equate the permanent component of output to consumption of nondurables and services, which is much more cyclical and volatile than our permanent component. Kim et al. (2002) estimate a common permanent component for output and consumption, with regime-switching, and argue that the permanent component contributes meaningfully to business cycle fluctuations.

When we estimate the model with output, compensation, and consumption on a per capita (instead of on a per hour) basis, we get very different results, as indicated in the third column of Table 2. Under this specification, the growth regime shift disappears (note that the estimated values of μ_0 and μ_1 are identical, so the transition probabilities are indeterminate). To better understand this result, note that the dependent variables in the per capita specification differ by the logarithm of hours per capita. We can see how the models are related if we begin with the equation for output per hour from the per hour specification, and then add the equation for detrended hours:

$$Y_t - L_t = \gamma X_t + \lambda_1 x_t + z_{1t}, \tag{18}$$

$$\hat{L}_t = \lambda_4 x_t + z_{4t}, \tag{19}$$

$$Y_t - \left(L_t - \hat{L}_t\right) = \gamma X_t + (\lambda_1 + \lambda_4) x_t + z_{1t} + z_{4t}. \tag{20}$$

The resulting dependent variable is output relative to the trend in hours. That trend in turn can be divided into population growth and the low-frequency movement in Λ (i.e. the trend in L/N) depicted earlier in Fig. 1, which we will denote by $\bar{\Lambda}_t$. Consequently output per

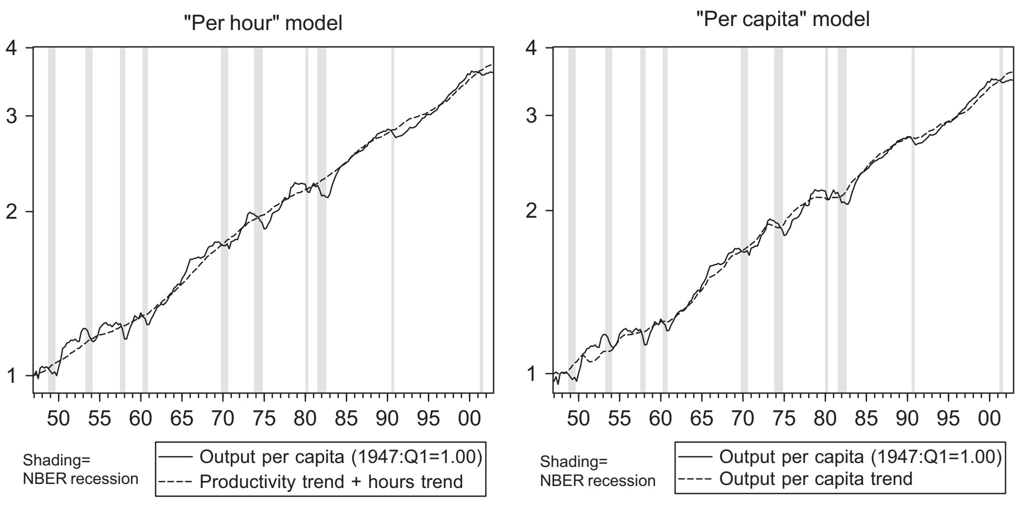

Fig. 5. Comparing the estimated trends in output per capita.

capita should be

$$Y_t - N_t = \bar{\Lambda}_t + \gamma X_t + (\lambda_1 + \lambda_4)x_t + z_{1t} + z_{4t}, \tag{21}$$

where $\bar{\Lambda}_t$ is just $L_t - \hat{L}_t - N_t$. And just as $W_t{-}L_t$ and $C_t{-}L_t$ have the same form as (18), $W_t{-}N_t$ and $C_t{-}N_t$ have the same form as (21).

Thus the permanent component in the per capita specification is $\bar{\Lambda}_t + \gamma X_t$ rather than just γX_t, i.e. it is the sum of the labor supply and technology trends. While there are other differences between the two specifications, a plausible explanation for the absence of a detectable regime shift in the permanent component of the per capita specification is simply that movements in $\bar{\Lambda}_t$ obscure the regime shifts in X_t. That in itself would not be a problem if the estimated permanent component were similar to the one obtained from combining the separate labor supply trend $\bar{\Lambda}_t$ (i.e. the HP filter of hours of work less population depicted in Fig. 1) and permanent component γX_t from the per hour specification. This does not turn out to be the case, however, as Fig. 5 illustrates. The estimated trends in per capita output from the two specifications clearly differ qualitatively, with the trend from the per capita specification showing much greater cyclical movement as well as more volatility in general. In addition, the standard deviation of the permanent component is almost twice as large as that from the per hour model. Thus it would appear that failure to examine the two trends individually results in the obscuring of fundamental properties of the permanent component of output per capita. It does so by making it difficult to detect the breaks in the technology trend, and by assigning more of the cyclical variation to the permanent component.

This last point recalls Perron's (1989) argument that evidence of unit roots can be a consequence of series with occasional structural breaks. The idea is that if an estimation procedure ignores a structural break in a linear trend, for example, then deviations from a linear trend may be so persistent as to suggest unit root behavior. In the present example, the per capita specification fails to find a regime switch in its permanent component. Our conjecture is that the resulting linear trend gives rise to greater persistence, and results in some transitory movements being labeled as permanent by the Kalman filter. In support of this it is worth noting that the estimate of the AR(1) parameter ϕ goes from -0.245 in the per hour specification to 0.543 in the per capita model. The persistence of the idiosyncratic components is higher as well.

The contrast between the per capita and per hour results also provides an explanation for Hamilton's (1989) surprising finding that growth regime shifts characterize business cycles rather than longer-term trends. Our results suggest that the treatment of hours of work is crucial. When we normalize the trending series by aggregate hours so as to remove the influence of variable labor supply trends, growth regime shifts are no longer cyclical phenomena, but instead reflect changes in long-term productivity trends. This finding is strengthened by looking at common trends in a multivariate analysis.

To summarize, the per hour specification leads us to characterize aggregate output as being made up of three distinct and more or less independent components: a stationary "business cycle" component, a permanent "technology trend" component and a low frequency "labor supply" component. The business cycle component exhibits the standard hump-shaped behavior emphasized by Blanchard (1981). The technology component is close to piecewise linear, with two breaks, one around 1973, and the other around 1997. The labor supply component is J-shaped, declining in the post-war era until the early 1960s and then rising from the mid-1960s up through the present.[19]

6. A closer look at the specification

6.1. The importance of additional variables

We have stressed the importance of information gleaned from several series in identifying the common permanent component. It is reasonable to ask how important this consideration actually is. To answer this question, we estimated the same econometric model, but with only one trending series, nonfarm output per hour, and one variable to capture the business cycle, detrended hours (variables 1 and 4 from the previous analysis). Thus we are looking to estimate trend growth with productivity data alone, using the detrended hours series to control for the business cycle. The result of this exercise is the fourth set of estimates in Table 2. Note first that the estimates of the transition probabilities are very similar to the earlier estimates, suggesting that the fundamental properties of the regime-switching dimension of the model are similar.

We also estimated two additional specifications, adding back alternately either C/L or W/L to the 2-equation system. The results are illustrated in Fig. 6, in which all four specifications' retrospective regime assessments are plotted against each other. There appears to be a big impact from adding either variable to the system, with the labor compensation variable clearly outperforming consumption/hours in terms of mimicking the 4-equation estimates of the regime probabilities. The change in going from three to four equations is relatively modest, especially if the fourth variable is consumption/hours.

[19]There are undoubtedly deeper explanations for the trend depicted in Fig. 1, related to separate treatment of men's and women's labor supply, the return to schooling, the structure of retirement benefits, and so on. See, for example, Ramey and Francis (2006). These are obviously beyond the scope of this paper.

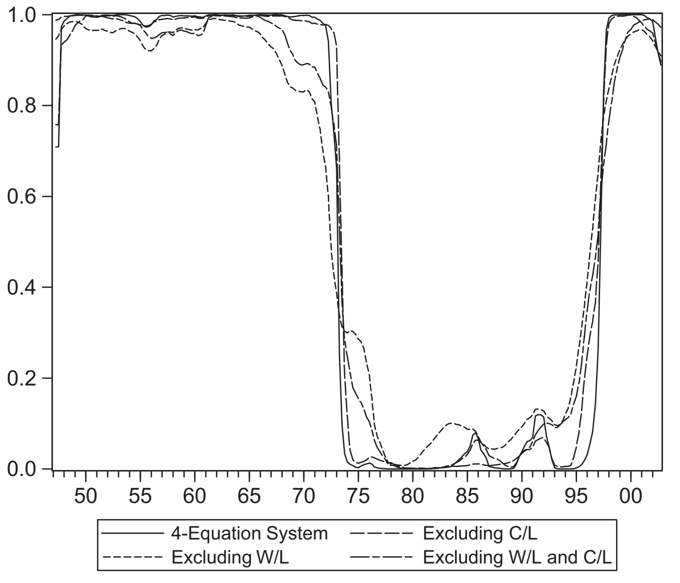

Fig. 6. Relative importance of variables for regime assessments.

Perhaps the most striking finding is that, at least judged by Fig. 6, consumption contributes least to the success of the model in identifying regime shifts. Even in the specification that excludes compensation per hour, the consumption variable does not appear to help detect the regime switches. This may reflect some sort of adjustment cost or recognition lag that results in a delayed response of consumption. It may also reflect factors that affect consumption but that are unrelated to changes in nonfarm productivity growth—for example, changes in government spending or taxes, changes in demographics, and the like. While these factors were not enough to overturn the cointegration implications of the theory, they could weaken the immediate connection to productivity growth. The relation of hourly labor compensation to labor productivity is more robust, depending primarily on the stability of labor's share of income.

Finally, we considered a variety of other specifications. We examined the specification suggested by the embodied technical progress model (e.g. Greenwood et al., 1997), which has a trend in the price of investment goods relative to consumption goods. This approach calls for expressing all variables in a common unit, for example by using the deflator for consumption goods. Surprisingly, the tests reject the $(1, 0, -1)$, $(0, 1, -1)$ cointegrating vectors more decisively than in the benchmark model, with estimates of the coefficient on C/L in the vicinity of -0.8. This rejection is puzzling because the embodied progress model is a generalization of the model from Section 2. Resolving this puzzle is beyond the scope of the paper, and in any case is not crucial to estimation of the embodied progress model, as we can estimate it without imposing the $\gamma_1 = \gamma_2 = \gamma_3$ constraint. When we do this we obtain similar results to the benchmark.

We also estimated the model with alternative measures of consumption, for example including just nondurables and services, or adding government purchases. The case for these is weaker, in our view, since theory provides no basis for ruling out trends in the ratio of nondurables and services expenditure to total consumption, or of government expenditure to consumption. Indeed, we find that the cointegrating relationship is less clear: While there is still an indication of a single common trend, the coefficients of the cointegrating relationships are again far from 1 and -1. Nonetheless the central finding of clear evidence for the two productivity growth regime switches was robust to these alternatives as well.

6.2. *The importance of regime-switching*

Because our empirical framework allows for regime shifts in the permanent and cyclical components, it would be instructive to offer evidence in support of this feature of the model. Formal testing of the Markov-switching model is not straightforward due to the presence of nuisance parameters that are not identified under the null hypothesis of parameter constancy. In addition, the information matrix is singular under the null hypothesis. In response to these challenges, a few papers have proposed tests for Markov-switching. Hansen (1992) considers a grid search method over the nuisance parameters to derive the supremum value of a likelihood ratio test statistic. Alternatively, Garcia (1998) studies the asymptotic distribution of a supremum-type likelihood ratio test. A drawback to each of these testing procedures is that they require estimating the model under the alternative hypothesis of Markov-switching which may be cumbersome. In addition, neither paper investigates the local power of the tests.

Recently, Carrasco et al. (2004) have proposed a new testing procedure for Markov-switching. Their approach not only is attractive in that the tests are asymptotically optimal, but also is quite tractable in that it only requires estimation of the model under the null hypothesis of parameter stability. For our purposes, we implemented their testing procedure by setting $\mu_1 = \mu_2 = \mu$ as well as $\tau = 0$ in our benchmark specification and then considering an alternative model with Markov-switching in the growth rate of the permanent component.[20] The test statistic turns out to be 19.07, whereas the 5% critical value is only 4.01. Thus, we find overwhelming evidence in support for including regime-switching in our dynamic factor model.

Another metric—less formal, but arguably of great practical importance—for the significance of regime-switching is the extent to which it helps forecast out of sample We will examine the benchmark model against alternatives along this dimension when we turn to real time data in the next section.

7. Real time estimates

In this section we explore the question of at what point the methods employed in this paper would have successfully detected the regime switch in productivity growth in real time, that is, with only the information available at a given point in time. We will also compare the out-of-sample forecast performance across alternative specifications. We collected real-time nonfarm productivity and hours data for vintages going back as far as

[20]The specification of the alternative model is based on two considerations. First, the results from Section 4 strongly suggest that the regime switches in the permanent component are a much more important feature of the data than those associated with the cyclical component. Second, the parameter τ cannot be identified in the absence of regime switches for the cyclical component.

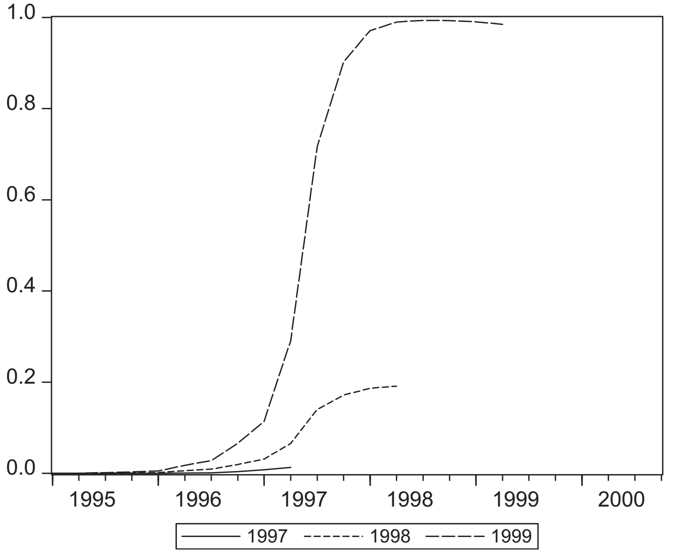

Fig. 7. Real-time estimates of high-growth regime probabilities by data vintage.

available. It turns out that vintages going back to the early 1990s are readily available from the Bureau of Labor Statistics web site. We also obtained consumption data from the Federal Reserve Bank of Philadelphia's real time database (see http://www.phil.frb.org/econ/forecast/reaindex.html). The results enable us to see how quickly the model would have detected what proved to be a change in trend productivity beginning in the second half of the 1990s.[21]

It turns out that the benchmark model would have picked up the change in trend within roughly a year of when (with hindsight) it occurred. Fig. 7 provides the estimated regime probabilities using data from vintages 1997, 1998, and 1999 (as of August of each year). By 1998 the estimates show a distinct increase in the high-growth regime probability, and by 1999 the estimates are essentially indistinguishable from those shown in Fig. 2. Moreover, as we shall see next, the changes in the regime probabilities translated to a discrete jump in the forecast of long-term productivity growth.

We now turn to out-of-sample forecast performance as another basis for comparing the benchmark model to alternative specifications. In addition to verifying the practical benefits of good forecasting, this diagnostic also helps to guard against overfitting, i.e. including additional parameters that help the model look better in-sample, but actually harm its forecasting performance out-of-sample. We compare real-time forecasts based on the 2-equation model, the 4-equation model, and the model with no regime-switching, using methods described in Hamilton (1994). Since the models with and without regime-switching are likely to differ most in the vicinity of what the former takes to be a switch, it

[21]We recognize that the advance selection of the number of regimes for a Markov switching model may be problematic for this sort of real-time analysis. For example, it might not make sense to specify a regime-switching model for real-time analysis of the early 1970s, when there would have been no real basis for thinking there was a low-growth regime. This objection is weaker for our application to the 1990s, since by then at least both regimes had been observed.

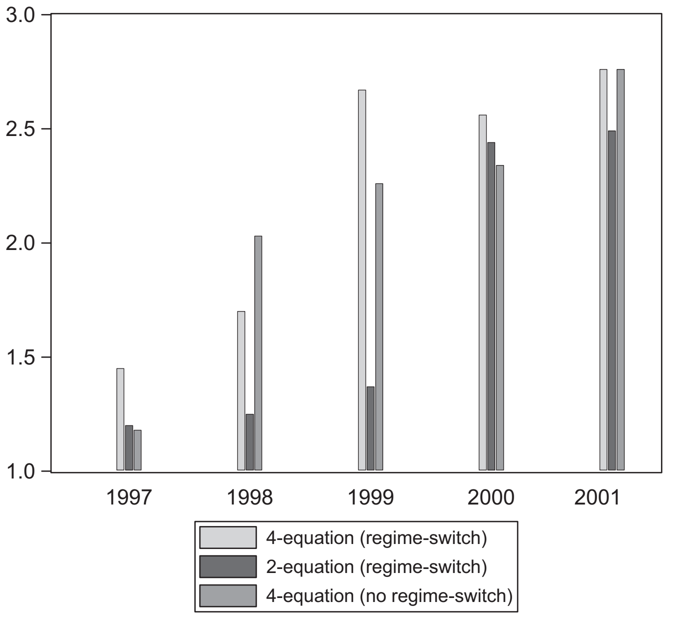

Fig. 8. Comparison of 4-year-ahead forecasts by data vintage.

is important to undertake the out-of-sample forecasting exercise during such a time period. Consequently, we use vintages from 1997 through 2001 (annually, as of August).

We first compare the 4-year-ahead point forecast of productivity growth (i.e. the forecast for 2001:Q3 from the 1997 vintage, for 2002:Q3 from the 1998 vintage, etc.) for each of the three specifications. The idea of looking four years ahead is to filter out differences in short-term forecasting accuracy, and focus primarily on the trend component. The results are depicted in Fig. 8. Clearly the 4-equation model with the regime switch was the first to register any kind of uptick (in 1997), and its forecasts are generally higher—and therefore better—during the transition from the low- to high-growth regime. It is also the first to jump above 2.5%, which it did in 1999. The 2-equation model clearly lagged behind the 4-equation models in picking up on the acceleration, while the no-regime-switch model generally lagged behind the regime-switch model (except in 1998), and only caught up in 2001.

We next look at the models' forecasting accuracy more generally. For the same five vintages, we construct out-of-sample forecasts through 2005:Q3, and compute root mean square errors for each of the three specifications. We do the same for the December 2005 vintage data used to obtain the parameter estimates in Table 2 (with the sample ending in 2002:Q4). Table 3 provides the results of this exercise. The root mean square error of the two-variable system is considerably larger for all but the 2000 vintage, suggesting that the four-variable systems are generally doing a better job. The benchmark model does only slightly better on average than the model with no regime-switch, following the pattern suggested by Fig. 8. Still, the fact that it does even slightly better, coupled with the results from the Carrasco et al. test, suggests that inclusion of the regime-switching component does not result in overfitting.

Thus even in real time it would have been clear by mid-1999, using the benchmark specification, that there had been a shift to significantly higher trend productivity growth.

Table 3
Out-of sample forecast performance

Root mean squared errors			
Data vintage	2-variable system, regime-switching	4-variable system, no regime-switching	4-variable system, regime-switching
1997[a]	3.12	3.13	2.98
1998[a]	3.37	2.86	3.02
1999[a]	3.54	2.92	2.84
2000[a]	2.62	2.70	2.63
2001[a]	2.60	2.46	2.45
2005[b]	2.71	2.44	2.40
Average	2.99	2.75	2.72

[a]Estimated on data available as of August 15 of each year, meaning through Q2, and forecast through 2005:Q3.
[b]Data available as of December 31, 2005, estimated through 2002:Q4, and forecast through 2005:Q3.

The simpler specifications considered above would have given a more ambiguous signal: The 2-equation model only picked up the shift a year later; the no-regime-switch model did a little better early on, but only gradually increased its forecast so as to catch up to the benchmark model by 2001. While certainly the idea of a ''new economy'' with strong productivity growth had gained many adherents well before 1999, there were also plenty of nay-sayers, and few of the optimists would have ventured to base their views on objective statistical analysis.[22]

8. Conclusions

The view that higher productivity growth is likely to be sustained has only really gained something approaching a consensus with the last recession. Prior to 2001, one could (and many did) easily argue that the increased growth rates experienced since 1996 were merely cyclical or otherwise ephemeral. That lack of agreement not only reflects the difficulty of separating a time series into its trend and cycle, but also the sensitivity of the results to various assumptions used in the decomposition. In the case of productivity, the problematic nature of the decomposition is only likely to be exacerbated by the inherent volatility of the series. Policymakers faced the same difficulty (albeit in the opposite direction) in the mid-1970s when the dramatic slowing of productivity growth coincided with a severe recession.

This paper analyzes productivity behavior by adopting a modeling strategy that integrates both theoretical considerations and recently developed statistical methods. The multivariate analysis exploits information from additional variables that growth theory implies should be cointegrated with the trend in productivity. We find strong support in the data for the notion that the economy (and productivity growth in particular) switched

[22]Indeed, optimistic views go back as early as 1997 (see, for example, *The New York Times*, August 2, 1997, ''Measuring Productivity in the 90's: Optimists vs. Skeptics,'' by Louis Uchitelle). The optimism appears, however, to have been based on something other than the productivity data themselves, about which there was much skepticism (see Corrado and Slifman, 1999). Of course, pessimists have also based their views on skepticism about the data, e.g. Roach (1998).

from a relatively low-growth to a high-growth regime in 1997. The annualized difference between the mean growth rates in the two regimes is estimated to be approximately 1.5%. We also show from estimates using real-time data that these techniques could have provided conclusive signals of the regime shift by 1999. Finally, from a methodological standpoint we argue that the incorporation of additional information from other time series is crucial to the strength of our conclusions.

The results also show that taking account of low-frequency movements in labor supply is decisive for detecting the regime shift in the permanent "technology" component. An alternative specification based on the assumption of stationary hours of work per capita fails to find any regime shift in the permanent component. This suggests an explanation for why other authors (e.g. Hamilton, 1989) who have applied regime-switching models to GDP data have found that growth regimes are associated with business cycle fluctuations rather than with low-frequency changes. It would appear that low-frequency movements in hours of work mask the regime changes in output per hour, so that the latter are hard to detect in GDP data.

Appendix A

A.1. State space model

We employ the following state-space representation for our model:

Measurement equation: $\Delta Q_t = H'\xi_t, \qquad \Delta Q_t \equiv (\Delta Q_{1t}, \ldots, \Delta Q_{4t})'$
Transition equation: $\xi_t = \alpha(S_t) + F\xi_{t-1} + V_t,$

with $E(V_t V_t') = \Sigma,$

and where (after we restrict $\gamma_1 = \gamma_2 = \gamma_3 \equiv \gamma$, and set $\gamma_4 = 0$)

$$H' = \begin{bmatrix} \gamma & \lambda_1 & -\lambda_1 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \gamma & \lambda_2 & -\lambda_2 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ \gamma & \lambda_3 & -\lambda_3 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & \lambda_4 & -\lambda_4 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}$$

$$\xi_t = \begin{bmatrix} \Delta X_t \\ x_t \\ x_{t-1} \\ z_{1t} \\ z_{1t-1} \\ z_{2t} \\ z_{2t-1} \\ z_{3t} \\ z_{3t-1} \\ z_{4t} \\ z_{4t-1} \end{bmatrix}, \quad \alpha(S_t) \equiv \alpha(S_{1t}, S_{2t}) = \begin{bmatrix} \mu_0(1 - S_{1t}) + \mu_1(S_{1t}) \\ \tau S_{2t} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad V_t = \begin{bmatrix} \upsilon_t \\ \varepsilon_t \\ 0 \\ \eta_{1t} \\ 0 \\ \eta_{2t} \\ 0 \\ \eta_{3t} \\ 0 \\ \eta_{4t} \\ 0 \end{bmatrix},$$

$$F=\begin{bmatrix}\phi_1&0&0&0&0&0&0&0&0&0&0\\0&\phi_1^*&\phi_2^*&0&0&0&0&0&0&0&0\\0&1&0&0&0&0&0&0&0&0&0\\0&0&0&\psi_{11}&0&0&0&0&0&0&0\\0&0&0&1&0&0&0&0&0&0&0\\0&0&0&0&0&\psi_{21}&0&0&0&0&0\\0&0&0&0&0&1&0&0&0&0&0\\0&0&0&0&0&0&0&\psi_{31}&0&0&0\\0&0&0&0&0&0&0&1&0&0&0\\0&0&0&0&0&0&0&0&0&\psi_{411}&\psi_{412}\\0&0&0&0&0&0&0&0&0&1&0\end{bmatrix},$$

$$\Sigma=\begin{bmatrix}1&0&0&0&0&0&0&0&0&0&0\\0&1&0&0&0&0&0&0&0&0&0\\0&0&0&0&0&0&0&0&0&0&0\\0&0&0&\sigma_1^2&0&0&0&0&0&0&0\\0&0&0&0&0&0&0&0&0&0&0\\0&0&0&0&0&\sigma_2^2&0&0&0&0&0\\0&0&0&0&0&0&0&0&0&0&0\\0&0&0&0&0&0&0&\sigma_3^2&0&0&0\\0&0&0&0&0&0&0&0&0&0&0\\0&0&0&0&0&0&0&0&0&\sigma_4^2&0\\0&0&0&0&0&0&0&0&0&0&0\end{bmatrix}.$$

We next provide a brief overview of a filter developed by Kim (1994) that can be used for approximate maximum likelihood estimation of the state-space model with Markov switching. We focus our attention on the issue of drawing inferences about the unobserved regimes. For further details, interested readers are referred to Kim and Murray (2002).

To facilitate the discussion, we will represent the two unobserved Markov-switching variables S_{1t} and S_{2t} by a single Markov-switching variable defined such that

$$\begin{aligned}S_t&=1 \quad \text{if } S_{1t}=0 \text{ and } S_{2t}=0,\\S_t&=2 \quad \text{if } S_{1t}=0 \text{ and } S_{2t}=1,\\S_t&=3 \quad \text{if } S_{1t}=1 \text{ and } S_{2t}=0,\\S_t&=4 \quad \text{if } S_{1t}=1 \text{ and } S_{2t}=1,\end{aligned}$$

with

$$\Pr[S_t = j | S_{t-1} = i] = p_{ij}$$

and

$$\sum_{j=1}^{4} p_{ij} = 1.$$

Conditional on $S_t = j$ and $S_{t-1} = i$, the Kalman filter equations are given by

$$\begin{aligned}
\xi_{t|t-1}^{(i,j)} &= \alpha(S_j) + F\xi_{t-1|t-1}^{i} + V_t,\\
P_{t|t-1}^{(i,j)} &= FP_{t-1|t-1}^{i}F' + \Sigma,\\
\eta_{t|t-1}^{(i,j)} &= \Delta Q_t - H'\xi_{t|t-1}^{(i,j)},\\
f_{t|t-1}^{(i,j)} &= HP_{t|t-1}^{(i,j)}H',\\
\xi_{t|t}^{(i,j)} &= \xi_{t|t-1}^{(i,j)} + P_{t|t-1}^{(i,j)}H'[f_{t|t-1}^{(i,j)}]^{-1}\eta_{t|t-1}^{(i,j)},\\
P_{t|t}^{(i,j)} &= (I - P_{t|t-1}^{(i,j)}H'[f_{t|t-1}^{(i,j)}]^{-1})HP_{t|t-1}^{(i,j)},
\end{aligned}$$

where $\xi_{t|t}^{(i,j)}$ and $\xi_{t|t-1}^{(i,j)}$ are, respectively, an inference on ξ_t based on information through time period t (Ω_t) and $t-1$ (Ω_{t-1}), given $S_t = j$ and $S_{t-1} = i$; $P_{t|t}^{(i,j)}$ and $P_{t|t-1}^{(i,j)}$ are, respectively, the mean squared error matrix of $\xi_{t|t}^{(i,j)}$ and $\xi_{t|t-1}^{(i,j)}$, given $S_t = j$ and $S_{t-1} = i$; $\eta_{t|t-1}^{(i,j)}$ is the conditional forecast error of ΔQ_t based on information through time period t-1, given $S_t = j$ and $S_{t-1} = i$; and $f_{t|t-1}$ is the conditional variance of the forecast error $\eta_{t|t-1}^{(i,j)}$.

To keep the Kalman filter from becoming computationally infeasible, the following approximations are introduced to collapse the posteriors terms $\xi_{t|t}^{(i,j)}$ and $P_{t|t}^{(i,j)}$ into the posterior terms $\xi_{t|t}^{j}$ and $P_{t|t}^{j}$:

$$\xi_{t|t}^{j} = \frac{\sum_{i=1}^{4} \Pr[S_{t-1} = i, S_t = j|\Omega_t]\xi_{t|t}^{(i,j)}}{\Pr[S_t = j|\Omega_t]}$$

and

$$P_{t|t}^{j} = \frac{\sum_{i=1}^{4} \Pr[S_{t-1} = i, S_t = j|\Omega_t]\left\{P_{t|t}^{(i,j)} + (\xi_{t|t}^{j} - \xi_{t|t}^{(i,j)})(\xi_{t|t}^{j} - \xi_{t|t}^{j})'\right\}}{\Pr[S_t = j|\Omega_t]}.$$

The approximations result from the fact that $\xi_{t|t}^{(i,j)}$ does not calculate $E[\xi_t|S_{t-1} = i, S_t = j, \Omega_t]$ and $P_{t|t}^{(i,j)}$ does not calculate $E[(\xi_t - \xi_{t|t}^{(i,j)})(\xi_t - \xi_{t|t}^{(i,j)})'|S_{t-1} = i, S_t = j, \Omega_t]$ exactly. This is because ξ_t conditional on Ω_{t-1}, $S_t = j$, and $S_{t-1} = j$ is a mixture of normals for $t>2$.

To obtain the probability terms necessary to construct the approximations, the following three-step procedure is employed.

Step 1:

At the beginning of the tth iteration, given $\Pr[S_{t-1}=i|\Omega_{t-1}]$, we can calculate

$$\Pr[S_t=j,S_{t-1}=i|\Omega_{t-1}]=\Pr[S_t=j|S_{t-1}=i]\times\Pr[S_{t-1}=i|\Omega_{t-1}],$$

where $\Pr[S_t=j|S_{t-1}=i]$ is a transition probability.

Step 2:

We can then consider the joint density of ΔQ_t, S_t and S_{t-1}:

$$f(\Delta Q_t,S_t=j,S_{t-1}=i|\Omega_{t-1})=f(\Delta Q_t|S_t=j,S_{t-1}=i,\Omega_{t-1})\times\Pr[S_t=j,S_{t-1}=i|\Omega_{t-1}]$$

and then obtain the marginal density of ΔQ_t as

$$\begin{aligned}f(\Delta Q_t|\Omega_{t-1})&=\sum_{i=1}^{4}\sum_{j=1}^{4}f(\Delta Q_t,S_t=j,S_{t-1}=i|\Omega_{t-1})\\&=\sum_{i=1}^{4}\sum_{j=1}^{4}f(\Delta Q_t,S_t=j,S_{t-1}=i|\Omega_{t-1})\times\Pr[S_t=j,S_{t-1}=i|\Omega_{t-1}],\end{aligned}$$

where the conditional density $f(\Delta Q_t|S_t=j,S_{t-1}=i,\Omega_{t-1})$ is obtained using the prediction error decomposition:

$$f(\Delta Q_t|S_t=j,S_{t-1}=i,\Omega_{t-1})=(2\pi)^{\frac{T}{2}}|f_{t|t-1}^{(i,j)}|^{-\frac{1}{2}}\exp\left\{-\frac{1}{2}\eta_{t|t-1}^{(i,j)'}f_{t|t-1}^{(i,j)-1}\eta_{t|t-1}^{(i,j)}\right\}.$$

A byproduct of this step is that we can obtain the log likelihood function:

$$\ln L=\sum_{t=1}^{T}\ln(f(\Delta Q_t|\Omega_{t-1})),$$

which can be maximized with respect to the parameters of the model.

Step 3:

We can then update the probability terms after observing ΔQ_t and the end of period t:

$$\begin{aligned}\Pr[S_t=j,S_{t-1}=i|\Omega_t]&=\Pr[S_t=j,S_{t-1}=i|\Delta Q_t,\Omega_{t-1}]\\&=\frac{f(S_t=j,S_{t-1}=i,\Delta Q_t|\Omega_{t-1})}{f(\Delta Q_t|\Omega_{t-1})}\\&=\frac{f(\Delta Q_t|S_t=j,S_{t-1}=i,\Omega_{t-1})\times\Pr[S_t=j,S_{t-1}=i|\Omega_{t-1}]}{f(\Delta Q_t|\Omega_{t-1})}\end{aligned}$$

with

$$\Pr[S_t=j|\Omega_t]=\sum_{i=1}^{4}\Pr[S_t=j,S_{t-1}=i|\Omega_t].$$

The last term provides the "real-time" inference about the unobserved regimes conditional on only contemporaneously available information.

We can also derive smoothed values of ξ_t and S_t using all available information through period T. That is, we can construct $\xi_{t|T}$ as well as $\Pr[S_t=j|\Omega_T]$ which represent the "retrospective" assessments of the state vector and unobserved regimes. Because the inferences about the unobserved regimes do not depend on the state vector, we can first calculate smoothed probabilities. The smoothed probabilities can then be used to generate the smoothed estimates of the state vector.

The smoothing algorithm for the probabilities will involve the application of approximations similar to those introduced in the basic filtering. The procedure can be understood by considering the following derivation of the joint probability that $S_{t+1} = k$ and $S_t = j$ conditional on full information:

$$\begin{aligned}
&\Pr[S_{t+1} = k, S_t = j|\Omega_T] \\
&\quad = \Pr[S_{t+1} = k|\Omega_T] \times \Pr[S_t = j|S_{t+1} = k, \Omega_T] \\
&\quad = \Pr[S_{t+1} = k|\Omega_T] \times \Pr[S_t = j|S_{t+1} = k, \Omega_t] \\
&\quad = \frac{\Pr[S_{t+1} = k|\Omega_T] \times \Pr[S_{t+1} = k, S_t = j|\Omega_t]}{\Pr[S_{t+1} = k|\Omega_t]} \\
&\quad = \frac{Pr[S_{t+1} = k|\Omega_T] \times \Pr[S_t = j|\Omega_t] \times \Pr[S_{t+1} = k|S_t = j]}{\Pr[S_{t+1} = k|\Omega_t]}
\end{aligned}$$

and

$$\Pr[S_t = j|\Omega_T] = \sum_{i=1}^{4} \Pr[S_{t+1} = k, S_t = j|\Omega_T].$$

The actual construction of the smoothed probabilities requires running through the basic filter and then storing the sequences $P^{(i,j)}_{t|t-1}, P^{j}_{t|t}$, $\Pr[S_t = j|\Omega_{t-1}]$ and $\Pr[S_t = j|\Omega_t]$. For $t = T-1, T-2, \ldots, 1$, the above formulas define a backwards recursion that can be used to derive the full-sample smoothed probabilities. It should be noted that the starting value for the smoothing algorithm is $\Pr[S_t = j|\Omega_t]$, which is given by the final iteration of the basic filter.

A.2. Partial updating

Suppose that additional observations become available, but only for some subset of the four data series represented by Q. Specifically, suppose that for the subset Q^1, data are available for periods 1 through $T+3$, whereas for Q^2, observations are only available through T. Let $T+1'$ denote the augmented information available through $T+1$, i.e. including Q^1_{T+1} but not Q^2_{T+1}.

Through T the standard Kalman updating algorithm applies (ignoring the regime-related term $\alpha(S_t)$ for brevity's sake), i.e.

$$\begin{aligned}
\hat{\xi}_{t+1|t} &= F\hat{\xi}_{t|t-1} + FP_{t|t-1}H(H'P_{t|t-1}H)^{-1}(\Delta Q_t - H'\hat{\xi}_{t|t-1}), \\
P_{t+1|t} &= F[P_{t|t-1} - P_{t|t-1}H(H'P_{t|t-1}H)^{-1}H'P_{t|t-1}]F' + \Sigma,
\end{aligned}$$

for $t \leqslant T$.

To update further, we simply recognize that

$$\Delta Q^i_t = H'^i\xi_t, \quad i = 1, 2,$$

where H^i is the appropriate submatrix of H. We then iterate beginning at $T+1$ according to

$$\begin{aligned}
\hat{\xi}_{T+1|T+1'} &= \hat{\xi}_{T|T-1} + P_{T|T-1}H^1(H'^1P_{T+1|T}H^1)^{-1}(\Delta Q^1_{T+1} - H'^1\hat{\xi}_{T+1|T}), \\
P_{T+1|T+1'} &= P_{T+1|T} - P_{T+1|T}H^1(H'^1P_{T+1|T}H^1)^{-1}H'^1P_{T+1|T},
\end{aligned}$$

and then

$$\hat{\xi}_{T+2|T+1'} = F\hat{\xi}_{T+1|T+1'},$$
$$P_{T+2|T+1'} = FP_{T+1|T+1'}F' + \Sigma.$$

The iteration can then proceed forward if subsequent observations on Q^1 become available.

In our case we essentially have three additional observations on two of the variables, since they appear lagged by three quarters. To take them into account at any point in time t, we can compute $\hat{\xi}_{t|t+3'}$, that is, the assessment of the state vector given the three additional observations of the two series that would otherwise be ignored because they are lagged by three quarters. To obtain $\hat{\xi}_{t|t+3'}$, we iterate forward to get $\hat{\xi}_{t+3|t+3'}$ and $P_{t+3|t+3'}$, and then iteratively "smooth" backwards using, e.g.

$$\hat{\xi}_{t+2|t+3'} = \hat{\xi}_{t+2|t+2'} + P_{t+2|t+2'}H^1(H'^1P_{t+3|t+2'}H^1)^{-1}(\hat{\xi}_{t+3|t+3'} - \hat{\xi}_{t+3|t+2'}).$$

The result is an improved estimate of the state vector, one that incorporates all information available at a given point in time.

References

Bai, J., Lumsdaine, R., Stock, J., 1998. Testing for and dating common breaks in multivariate time series. Review of Economic Studies 65, 395–432.

Beaudry, P., Koop, G., 1993. Do recessions permanently change output? Journal of Monetary Economics 31, 149–163.

Blanchard, O., 1981. What is left of the multiplier accelerator? American Economic Review 71 Papers and Proceedings, pp. 150–154.

Carrasco, M., Hu, L., Ploberger, W., 2004. Optimal Test for Markov Switching, unpublished manuscript, University of Rochester.

Christiano, L., Eichenbaum, M., Vigfusson, R., 2004. What Happens after a Technology Shock? unpublished manuscript.

Cogley, T., 2005. How fast can the new economy grow? A bayesian analysis of the evolution of trend growth. Journal of Macroeconomics 27, 179–207.

Corrado, C., Slifman, L., 1999. Decomposition of productivity and unit costs. American Economic Review Papers and Proceedings 89, 328–332.

Dickey, D., Fuller, W., 1979. Distribution of the estimates for autoregressive time series with a unit root. Journal of the American Statistical Association 74, 427–431.

Estrella, A., 2004. Decoding Productivity: Business Cycle Properties of Labor Productivity Growth. Federal Reserve Bank of New York, unpublished manuscript.

Francis, N., Ramey, V., 2005. Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. Journal of Monetary Economics 52, 1379–1399.

French, M., 2001. Estimating Changes in Trend Growth of Total Factor Productivity: Kalman and H-P Filters Versus a Markov-Switching Framework, Federal Reserve Board Staff Working Paper No. 2001-44.

Friedman, M., 1969. Monetary studies of the national bureau. In: The Optimum Quantity of Money and Other Essays. Aldine, Chicago.

Friedman, M., 1993. The "Plucking Model" of business fluctuations revisited. Economic Inquiry 31, 171–177.

Garcia, R., 1998. Asymptotic null distribution of the likelihood ratio test in Markov switching models. International Economic Review 39, 763–788.

Gordon, R., 2000. Does the 'New economy' measure up to the great inventions of the past? Journal of Economic Perspectives 14, 49–74.

Greenwood, J., Hercowitz, Z., Krusell, P., 1997. Long run implications of investment-specific technological change. American Economic Review 87, 342–362.

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

Hamilton, J., 1994. Time Series Econometrics. Princeton University Press, Princeton.
Hamilton, J., 2004. What's Real about the Business Cycle?, unpublished manuscript.
Hansen, B., 1992. The likelihood ratio test under non-standard conditions: testing the Markov switching model of GNP. Journal of Applied Econometrics 7, S61–S82.
Hansen, B., 2001. The new econometrics of structural change: dating changes in U.S. labor productivity. Journal of Economic Perspectives 15, 117–128.
Hodrick, R., Prescott, E., 1997. Post-war U.S. business cycles: an empirical investigation. Journal of Money, Credit and Banking 29, 1–16.
Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. Econometrica 59, 1551–1580.
Johansen, S., 1995. Likelihood-based Inference in Cointegrated Vector Autoregression Models. Oxford University Press, Oxford.
Kaldor, N., 1961. Capital accumulation and economic growth. In: Lutz, F., Hague, D.C. (Eds.), The Theory of Capital. St. Martin's Press, New York, pp. 177–222.
Kim, C.-J., 1994. Dynamic linear models with Markov-switching. Journal of Econometrics 60, 1–22.
Kim, C.-J., Murray, C., 2002. Permanent and transitory components of recessions. Empirical Economics 27, 163–184.
Kim, C.-J., Nelson, C., 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.
Kim, C.-J., Piger, J., 2002. Common stochastic trends, common cycles, and asymmetry in economic fluctuations. Journal of Monetary Economics 49, 1189–1211.
Kim, C.-J., Piger, J., Startz, R., 2002. Permanent and Transitory Components of Business Cycles: Their Relative Importance and Dynamic Relationship, Federal Reserve Bank of St. Louis Working Paper No. 2001-017B.
Mulligan, C., 2002. A Century of Labor-Leisure Distortions, National Bureau of Economic Research Working Paper No. 8774.
Perron, P., 1989. The great crash, the oil price shock, and the unit root hypothesis. Econometrica 57, 1361–1401.
Ramey, V., Francis, N., 2006. A Century of Work and Leisure, unpublished manuscript.
Roach, S., 1998. No Productivity Boom for Workers, Issues in Science and Technology.
Roberts, J., 2000. Estimates of the Productivity Trend Using Time-Varying Parameter Techniques, Federal Reserve Board Staff Working Paper No. 2001-8.
Rotemberg, J., 2003. Stochastic technical progress, smooth trends, and nearly distinct business cycles. American Economic Review 93, 1543–1559.
Sims, C., 2001. Comment on evolving post-world war II U.S. inflation dynamics. In: Bernanke, B.S., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2001. MIT Press, Cambridge.
Solow, R., 1956. A contribution to the theory of economic growth. Quarterly Journal of Economics 70, 65–94.
Staiger, D., Stock, J., Watson, M., 2001. Prices, Wages and the U.S. NAIRU in the 1990s. In: Krueger, A., Solow, R. (Eds.), The Roaring '90s: Can Full Employment Be Sustained. Russell Sage and Century Fund, New York.
Stiroh, K., 2002. Information technology and the U.S. productivity revival: what do the industry data say? American Economic Review 92, 1559–1576.
Stock, J., Watson, M., 1989. New indexes of coincident and leading economic indicators. NBER Macroeconomics Annual 1989, 351–393.
Stock, J., Watson, M., 2002. Macroeconomic forecasting using diffusion indexes. Journal of Business and Economic Statistics 20, 147–162.