Distinguishing trends from cycles in productivity

James A Kahn and Robert Rich,[1]
Federal Reserve Bank of New York

1. Introduction and summary

Productivity growth has long been a focus of both policymakers and economists. At least since Solow's (1956) pioneering work on long-term growth, economists have understood that productivity growth is the only source of sustainable increases 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 assessment of trend productivity growth. But on a quarterly basis, productivity growth is extremely volatile. It is also cyclical, typically rising during an expansion and falling at the onset of a recession. Thus it is often only with hindsight that any change in its long-term trend can be discerned.

It is widely believed that this difficulty contributed significantly to the economic instability of the 1970s, as policymakers and economists were unaware of the slowdown in productivity growth for many years, and only much later were able to date the slowdown at approximately 1973. Policymakers thus overestimated potential GDP (at least so the conventional wisdom goes) and left interest rates too low, and double digit inflation followed not long after.

In recent years, attention has turned once again to productivity because of speculation that growth rates may be picking up again. The growth rate of non-farm output per hour increased by approximately 1% beginning in 1996 relative to the period 1991-95, and by about 1.3% relative to 1973-95. The acceleration of productivity puts its growth rate during this five-year period close to where it was during the most recent period of strong growth, from roughly 1948 to 1973. This has provoked a debate over whether we can expect an extended period of more rapid productivity growth. Robert Gordon (2000), for example, attributes about half of the acceleration to a "cyclical" effect, and much of the remainder to measured productivity growth in the technology sector. Others (eg Stiroh, forthcoming) find evidence that productivity growth has spilled over into other sectors through capital deepening.

Much of the difficulty in evaluating the arguments in this debate relates to the issue of separating trend from cycle in the data. For example, if Gordon had assumed an acceleration in trend GDP, then he would have found a smaller or non-existent output gap, and consequently would have been unable to attribute the productivity acceleration to cyclical effects. Thus without more information, either story (new economy: accelerating productivity and output; or old economy: increased productivity growth confined to information technology (IT) sector, all else is cyclical) seems consistent with the productivity data. This is a problem that plagues any effort to distinguish trend from cycle in a single time series over a relatively short period of time. Moreover, even apart from the difficulty of distinguishing trend from cycle, it is reasonable to question whether much of anything can be learned from five or six years of data on a series as volatile as productivity growth.

In this paper we attack this problem with a multivariate approach, drawing on standard neoclassical growth theory to help us identify variables other than productivity - namely consumption and labour compensation - that should help to identify the trend in long-term growth. We treat that trend as a process with two "regimes", high- and low-growth, with some probability of switching between the two at any point in time. We model the business cycle as a stationary 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).

[1] The authors may be contacted by e-mail at james.kahn@ny.frb.org and robert.rich@ny.frb.org. 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 and Christian Murray for assistance with programs used in performing the computations.

There are several advantages to this approach. First, we argue that aggregate productivity data alone do not provide as clear or as timely a signal of structural change 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 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 Hansen (2001).

Our analysis picks up striking evidence of a switch in the second half of the 1990s to a higher long-term growth regime, as well as a switch in the early 1970s in the other direction. While these conclusions may come as no big surprise, our analysis has further implications. First, one could not conclude that there was a switch to a higher regime on the basis of productivity data alone, or even with the addition of a variable to control for the business cycle. Only the corroborating evidence from consumption and labour compensation can swing the balance in favour of a regime switch. Second, the ability to discern the switch to a higher-growth regime based on our analysis has appeared relatively recently. Only with data up to the end of 1999 could one conclude that with fairly high probability a switch had occurred.

2. Background: the neoclassical growth model

Some 40 years ago, Nicholas Kaldor (1961) established a set of stylised facts about economic growth that have guided empirical researchers ever since. His facts are: (1) labour and capital's income shares are relatively constant; (2) growth rates and real interest rates are relatively constant; (3) the ratio of capital to labour 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: that 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 either the level or growth rate of labour productivity has no permanent impact on employment.

Of course, closer inspection suggests that none of the above "stylised facts" is literally true. Indeed, the premise of much work on US productivity is that growth was systematically higher from 1948 to 1973 than it was over the subsequent 20-plus years. But they still provide a starting point for modelling economic growth. That starting point is generally referred to as the neoclassical growth model. The linchpins of this approach are typically a constant returns-to-scale Cobb-Douglas production technology in capital and labour, constant elasticity-of-substitution preferences for consumption, and exogenous labour-augmenting technological progress.

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

$$Y = K^{\alpha}(LE)^{1-\alpha}$$

and preferences defined in terms of a present discounted value of utility:

$$U = \ln(C/N) + \Lambda v(1-\ell)$$

where $\ell \equiv L/N$ represents the proportion of available hours devoted to work. The term $\Lambda v(1-\ell)$ therefore represents the utility of leisure, where v is a concave function, and Λ is a taste parameter that might shift over time.

Let $c \equiv C/(NE)$, $y \equiv Y/(NE)$, and $k \equiv K/(NE)$. If we assume for the moment that E grows at a constant rate g and that Λ is constant, then the economy described by the above preferences and technology has a steady state in which y, k, ℓ and c will be constant (see Technical Appendix A). This implies that the capital-labour ratio $K/L = kE$ and the ratio of consumption to hours of work $C/L = cE$ will both grow at rate g, as will output per hour Y/L. Also, if we assume that labour is compensated according to its marginal product, then labour compensation per hour W will simply equal $(1 - \alpha)Y/L$ and will also grow at the same rate. To summarise:

Result 1. *With constant technological progress (g) and no shifts in labour supply (Λ), the capital-labour ratio, consumption relative to hours worked, real labour compensation per hour, and output per hour should all have the same long-term trend, with growth rate g*.

(*Proof:* See Proposition 1 and its Corollary in Appendix A.)

The preceding result neatly characterises the relationships between variables of interest under the assumption that g and Λ are constant over time. Of course, this paper is predicated on the possibility that both (and g in particular) may have experienced important changes in the postwar period. To dispense with the easy case first, level changes in Λ have no impact on the conclusions of the previous paragraph: the permanent components of the capital-labour ratio, output per hour, consumption relative to hours, and labour compensation per hour, should all continue to move together except for transitory variation in response to changes in labour supply. For example, if Λ falls, the result will be an increase in L/N, but K and C will adjust proportionally.

Result 2. *Exogenous variation in labour supply (Λ) implies only transitory differences in the behaviour of the four variables considered in Result 1*.

(*Proof:* See Proposition 2 in Appendix A.)

Result 3. *Exogenous variation in g has only a transitory impact on labour supply.*

(*Proof:* See the Lemma in Appendix A.)

These two results are particularly helpful for understanding the relationship between employment and long-term growth: any impact in either direction is transitory. Trends in employment are primarily a function of demographic and "taste" factors.[2] The practical implication of this is that we can safely detrend employment without fear of discarding important information about underlying trends in per capita growth.

If the underlying growth rate of the economy g changes, the situation is a little more complicated with regard to the other variables. A sustained change in g implies a sustained change in the capital-labour ratio. This is essentially because a change in the growth rate requires a change in the real interest rate in the same direction. Faster growth, for example, raises future consumption relative to current levels, and therefore discourages saving and capital accumulation. The result is that the capital-labour ratio, while ultimately growing faster as well, also experiences a downward shift in its level. A corollary of this effect is that the output-capital ratio shifts upwards. This effect can be seen rather strikingly in Graph 1A, which depicts the output-capital ratio over the period 1948-95. This ratio should shift in the same direction as the underlying growth rate, which it clearly does. Moreover, the size of the shift is roughly what one would expect from the model using plausible parameter values.

The level shifts in output per hour, labour compensation per hour and consumption per hour, on the other hand, should all be the same, so that those three variables should have the same permanent component even in the face of regime shifts in growth (Graph 1B).

Result 4. *If the economy experiences a sustained change in its underlying growth rate, output per hour, hourly labour compensation and consumption/hour will display a common trend, while the capital-labour ratio will experience a level shift relative to the other series.*

(*Proof:* See Proposition 3 and its Corollary in Appendix A.)

The upshot of this foray into the neoclassical growth model is that other data, notably labour compensation and consumption relative to hours, may provide auxiliary information about the trend in output per hour. 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

[2] Of course, it is possible that growth affects demographics, and vice versa. For example, sustained higher growth seems to be related to lower fertility rates. But such relationships are unlikely to be apparent over a few decades.

suggests that considering these series together should provide better information about underlying trends than consideration of any of them alone.

Graph 1A

The US output-capital ratio, 1947-94

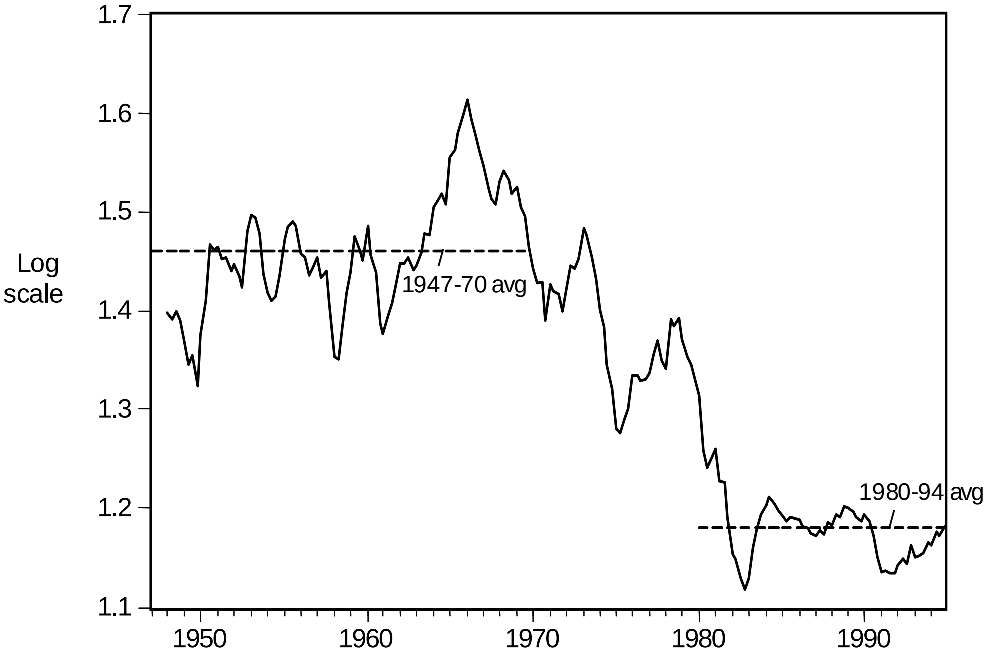

Graph 1B

The relative importance of level shifts in *K*, *C* and *W*

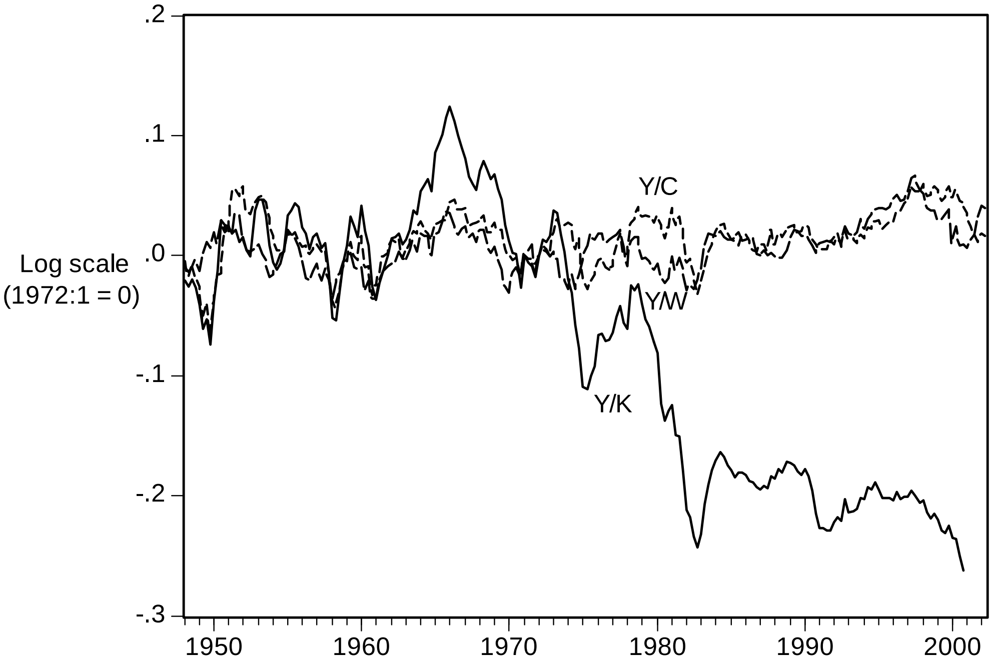

3. Econometric specification

Our estimation strategy draws upon the regime-switching dynamic factor model recently proposed by Kim and Murray (2002). The essence of our approach is to examine a number of related economic time series and to use their co-movement in order to identify two shared factors: a common permanent component and a common transitory component. The common permanent component relates to the movements in the series associated with their secular growth rate. On the other hand, the common transitory component relates to movements in the series associated with the business cycle. These components provide a basis to distinguish between long- and short-run changes in the series as well as to gauge their relative importance over particular historical episodes.

While there are other time series models that also include permanent and transitory components, the Kim and Murray (2002) model is unique in that it allows for the possibility that the components may be subject to changes in regime. That is, there may be periodic changes in the underlying process generating the common permanent and common transitory components.

For the purpose of analysing the trend and cyclical movements in productivity, the regime-switching in the model has several attractive features. First, the regime-switching common transitory component allows us to account for possible asymmetries across booms and recessions. In particular, we can capture the idea proposed in Friedman's (1964, 1993) "plucking model" that economic fluctuations are largely permanent during expansions and largely transitory during recessions. Second, the regime-switching permanent component allows us to account for changes in secular growth. Consequently, we cannot only address the issue of a post-1973 productivity growth slowdown, but also a post-1995 productivity growth resurgence. Last, while the model allows for periodic changes in the process generating the common permanent and common transitory components, the timing of the switches does not need to be specified prior to estimation. Rather, the timing of the switches is determined as one of the outcomes from the estimation procedure.

Following Kim and Murray (2002), we can describe the regime-switching dynamic factor model as follows.[3] Suppose we consider a number of time series indexed by i. Let ΔY_{it} denote the growth rate of the i th individual time series. It is assumed that the movements in each series are governed by the following process:

$$\Delta Y_{it} - D_i + \gamma_i \Delta\Pi_t + \lambda_i \Delta x_t + z_{it}$$

where D_i is the average growth rate of the series, $\Delta\Pi_t$ denotes the growth rate of a permanent component that is common to all series, Δx_t denotes the growth rate of a common transitory component, and z_{it} denotes an idiosyncratic component. 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.[4]

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

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

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

$$\Pr[S_{1t} = 0 | S_{1,t-1} = 0] = q_1 \text{ and } \Pr[S_{1t} = 1 | S_{1,t-1} = 1] = p_1$$

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 Π_t grows at the rate $\mu_0/(1-\phi_1-...-\phi_p)$ when $S_{1t} = 0$, and at the rate $\mu_1/(1-\phi_1-...-\phi_p)$ when $S_{1t} = 1$.

---

3 Additional details are provided in Appendix B.

4 As we demonstrate shortly, the formulations for the common permanent component and common transitory component have different implications for their effects on the individual time series. Specifically, the level of the common permanent component can increase over time, while the level of the common transitory component is stationary.

The common transitory component x_t is also subject to regime-switching, as in Kim and Nelson (1999):

$$x_t = \tau(S_{2t}) + \phi_1^* x_{t-1} + \phi_2^* x_{t-2} + \ldots + \phi_p^* x_{t-p} + \varepsilon_t \text{ and } \varepsilon_t \sim \text{iidN}(0,1)$$

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

$$\Pr[S_{2t} = 0 | S_{2,t-1} = 0] = q_2 \text{ and } \Pr[S_{2t} = 1 | S_{2,t-1} = 1] = p_2$$

where S_{2t} is an index of the regime for the common transitory component, with transition probabilities q_2 and p_2. The permanent and transitory regimes are assumed to be independent of each other. The parameter τ represents the size of the "pluck," with $\tau < 0$ implying that the common transitory component is plucked down during a recession. Finally, the idiosyncratic components are assumed to have the following structure:

$$z_{it} = \psi_{i1} z_{i,t-1} + \psi_{i2} z_{i,t-2} + \ldots + \psi_{ip} z_{i,t-p} + \eta_{it} \text{ and } \eta_{it} \sim \text{iidN}(0, \sigma_i^2)$$

where all innovations in the model are assumed to be mutually and serially uncorrelated at all leads and lags.

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.[5] An important by-product 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 inferences can be based on a "real time" assessment, which provides an estimate of the current regime as of date *t* based only on information up to date *t*, or a "retrospective" assessment that incorporates information for the entire sample.[6] Of course the retrospective assessment is a more reliable estimate of past regimes because it incorporates subsequent data, but the real-time assessment is useful as a gauge of what analysts and policymakers could realistically have known at the time.

Our data consist of four quarterly series on labour productivity, real compensation per hour, consumption deflated by hours of work, and hours of work itself. All of the data except aggregate consumption refer to the private non-farm sector, and all series cover the period 1947 Q1 to 2002 Q2.

The selection of the series and their subsequent transformation for the estimation are motivated by consideration of their usefulness for identification of the common permanent component and common transitory component. With regard to the hours series, as we have seen, there is little reason to believe that its underlying trend should be fundamentally related to that of the other three series. Consequently, where we include this series in the analysis to help refine the estimate of the common transitory component, we detrend it using the Hodrick-Prescott (1980) filter.

With regard to the dynamic specification of the common and idiosyncratic components, our diagnostic checking procedure suggested that the common permanent component should include one lagged value of $\Delta\Pi_t$, the common transitory component should include two lagged values of x_t, and that the idiosyncratic component should include one lagged value of z_{it} for each series. We also drew upon theory to impose restrictions on the γ_i which represent the permanent factor loadings in the model. Specifically, we restricted the estimated permanent factor loadings for productivity, real compensation per hour, and consumption per hour to be equal. In addition, we set the value of the permanent factor loading for detrended hours to be equal to zero.

5 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. More details of how we applied their methods to the present study are available on request.

6 The designation "real time" is potentially misleading on two counts. First, the parameter estimates are based on the entire sample. Second, the data may have undergone subsequent revisions that would have been unavailable as of the time of the assessment.

4. Results

As described above, our benchmark model uses the three variables that growth theory suggests should have approximately the same permanent components: real output per hour (variable 1), real labour compensation per hour (variable 2), and consumption relative to hours (variable 3). We also include detrended hours (variable 4) to help capture the transitory or business cycle component of the data. As we have seen, while the capital-labour ratio should also exhibit the same growth as output per hour within a growth regime, both theory and the data suggest that it can jump substantially between regimes, so we exclude it from our system.

One final specification issue relates to the synchronisation of the data. The three trending variables were selected primarily with a view to helping identify the common permanent component. They are not necessarily "coincident indicators" with respect to the transitory component. Since allowing for a more general lead/lag structure in our system would greatly increase the complexity of the model as well as the number of parameters to estimate, we instead examined the cross-correlations of the four series. We found that the first two variables (productivity and labour compensation per hour) tended to lead the other two series over the business cycle by about three quarters. To capture this in the estimation, we lagged variables 1 and 2 by three quarters in the system described above.[7]

The first set of results in Table 1 provides the parameter estimates for our benchmark model with these four variables. As mentioned, we restrict γ (the loading factor on the common permanent component) to be the same for the three trending variables, a restriction that we test and fail to reject for this specification. The sample covers 1947 Q1-2002 Q2. The results indicate that the model yields fairly precise estimates of most of the parameters of interest: the loading factors on both the permanent and transitory components, the transition probabilities, and the shift parameters connected with the regimes (μ_0, μ_1, τ) are all statistically significant. The difference between the low and high permanent regimes works out to be $\gamma(\mu_0 - \mu_1)/(1-\phi) = 0.367$, which corresponds to roughly 1.5% on an annualised basis. The transition probabilities for the permanent regimes imply that the unconditional probability of being in the high-growth regime is 0.593, which, as we shall see, reflects the inference that the economy was in the high-growth regime 59.3% of the time between 1947 and the present.

The transitory process is estimated to be a "hump-shaped" autoregressive process typical of the business cycle, but with a significant negative pluck. The four λ_i coefficients, which represent the loading factors on the transitory process, are all estimated very precisely, and with the expected signs. Note that λ_3, the loading factor for consumption/hour, is negative, reflecting the fact that hours is more cyclical than consumption. The positive estimates for the loading factors on the real compensation and productivity variables mean that those two variables associate positively with the transitory component, but leading by three quarters relative to the other two variables (since they enter the system lagged by three quarters).

Before further describing the permanent and transitory components of productivity growth, however, it is instructive to examine the inferred probability of being in the high- or low-growth state over time. There are two ways to look at this. First, in "real time", ie given only the information analysts would have had at each point in time, what probability would they have assigned to being in the high-growth state? The second way is retrospectively: given what has happened up to 2002 Q2, 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, ie as of 2002 Q2.) The retrospective version is plotted in Graph 2A. The vertical axis is the probability of being in the high-growth regime. The data present a very clear picture: the economy was in a high-growth state until the early 1970s, followed by a roughly 20-year low-growth regime, followed by a switch back to high growth in the mid-1990s. Perhaps the only surprise here is how unambiguous the current assessment is: the probability that the economy was in a sustained high-growth regime (given what we now know[8]) by 1998 Q4 is estimated to be about 0.94. Note that because productivity enters lagged three quarters, this would date its acceleration at 1998 Q1.

---

7 A drawback to this procedure is that the last three observations of variables 1 and 2 are not used in assessing the trend, even though they are known. We are looking at methods for addressing this problem.

8 But ignoring the last three observations of productivity and labour compensation.

Table 1

Estimation of model

Coefficient	Four-equation system	Two-equation system
p_1	0.989 (0.012)	0.992 (0.011)
q_1	0.984 (0.014)	0.992 (0.011)
p_2	0.982 (0.016)	0.952 (0.038)
q_2	0.555 (0.280)	0.509 (0.153)
ϕ	–0.381 (0.147)	–0.096 (0.162)
ϕ_1^*	1.424 (0.097)	1.439 (0.075)
ϕ_2^*	–0.576 (0.084)	–0.565 (0.072)
ψ_{11}	–0.246 (0.079)	–0.710 (0.233)
ψ_{21}	–0.011 (0.087)	–
ψ_{31}	–0.636 (0.087)	–
ψ_{41}	0.494 (0.070)	–0.026 (0.279)
γ	0.273 (0.041)	0.802 (0.123)
λ_1	0.212 (0.041)	0.156 (0.038)
λ_2	0.118 (0.031)	–
λ_3	–0.444 (0.049)	–
λ_4	0.434 (0.047)	0.581 (0.058)
μ_0	0.822 (0.228)	0.206 (0.126)
μ_1	–1.035 (0.271)	–0.173 (0.118)
τ	–2.926 (0.769)	–2.128 (0.458)

Note: The estimation also produces estimates of the variances of the idiosyncratic errors, not reported here.

As the previous discussion makes clear, however, this does not mean that as of the end of 1998 we actually could have made such an optimistic assessment. Graph 2B plots the "real time" version of Graph 2A, which indicates that as of the end of 1998 we would have only assessed the probability of being in the high-growth regime as 0.3. Thus it is only with hindsight that we think the probability of being in the high-growth regime was very high at that point in time. And it is easy to understand why, looking at Graph 2B. In real time there were a lot of "false alarms", post-1973 episodes when it appeared that productivity growth might have shifted into high gear. These tended to occur during the early phase of expansions, when productivity growth did in fact increase, and when it was too soon to tell - notwithstanding an effort to disentangle cycle from trend - whether the higher growth rate would be sustained.

Graph 2A

Retrospective assessment of growth regimes

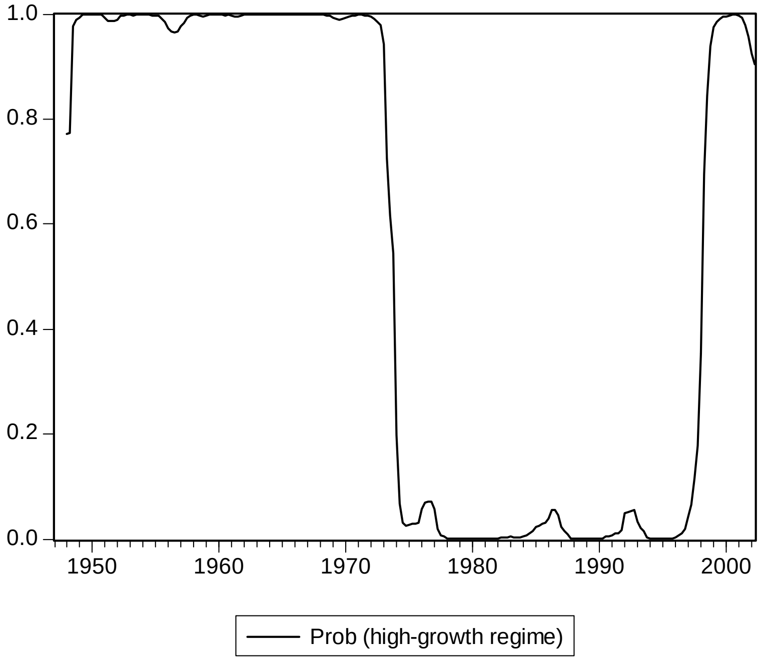

Graph 2B

Real-time assessment of growth regimes

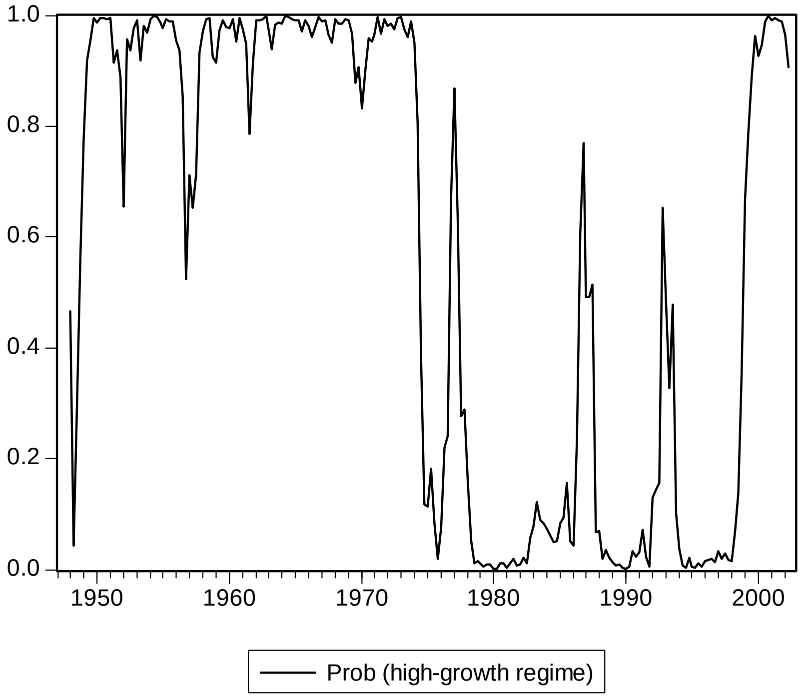

Nonetheless, even in real time it would have been clear from these techniques based on data only up to 1999 that there had been a change in regime back to high growth - this after only four years of relatively strong productivity growth in the middle of an expansion. While certainly the idea of a "new economy" with strong productivity growth had gained many adherents well before the end of 1999, there were also plenty of naysayers, and few of the optimists would have ventured to base their views on objective statistical analysis.[9]

For completeness, we can also look at the assessment of transitory regimes. The (retrospective) probability of being in the "plucked down" state is plotted in Graph 3A. Here the probability assessments are a little more ambiguous, which is not surprising considering that the regimes themselves are relatively transitory. Nonetheless, the more prominent spikes all coincide with NBER-defined recessions, though several recessions (most notably the 1990-91 recession) are missed entirely, and several others show up with a relatively low probability. It is perhaps instructive 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 characterised by a relatively mild downturn followed by an unusually slow and gradual recovery. If we instead look at the total common transitory component (which includes an autoregressive process as well as the pluck) in Graph 3B, we see that this recession does in fact register (as do all of the postwar recessions), albeit very mildly.

If we examine the permanent component of productivity growth directly (Graph 4A), we see that the permanent component clearly indicates changes in trend in the early 1970s and mid-1990s, with no apparent cyclical residue.[10] At the same time, however, the model still assigns a lot of variation to the "idiosyncratic" component (Graph 4B). Presumably this reflects the fact that productivity movements are often out of phase with other cyclical variables, though it also reflects our choice of variables to include in the system. This choice was geared entirely towards variables that would help to pick up the permanent component, regardless of their cyclical behaviour.

Finally, from a methodological perspective, one may ask how important is the use of additional series in helping to identify changes in trend growth. To answer this, we use the same econometric model, but with only two series, non-farm output per hour and 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 second set of estimates in Table 1. 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 will be similar. Second, while many other parameters are similar, it is clear that the nature of the permanent component is very different from the previous estimate. Not only is the factor loading much higher (0.802-0.273) - not surprising since it is the only trending variable - but its dynamics are different. This is indicated by the statistically insignificant estimated AR coefficient (ϕ) of -0.096, versus the estimate of -0.381 from the four-equation system that was both larger in absolute value and statistically significant.

9 Indeed, optimistic views go back as early as 1997 (see, for example, "Measuring productivity in the 1990s: optimists vs skeptics", by Louis Uchitelle, *New York Times*, 2 August 1997), though they appear to have been based more on scepticism about the data than on direct signals generated by the data (see Corrado and Slifman (1999)). Of course, pessimists have also based their views on scepticism about the data, eg Roach (1998).

10 This contrasts with Kim and Murray (2002), whose estimated permanent aggregate component shows substantial downward movement during recessions.

Graph 3A

Probability of a negative pluck

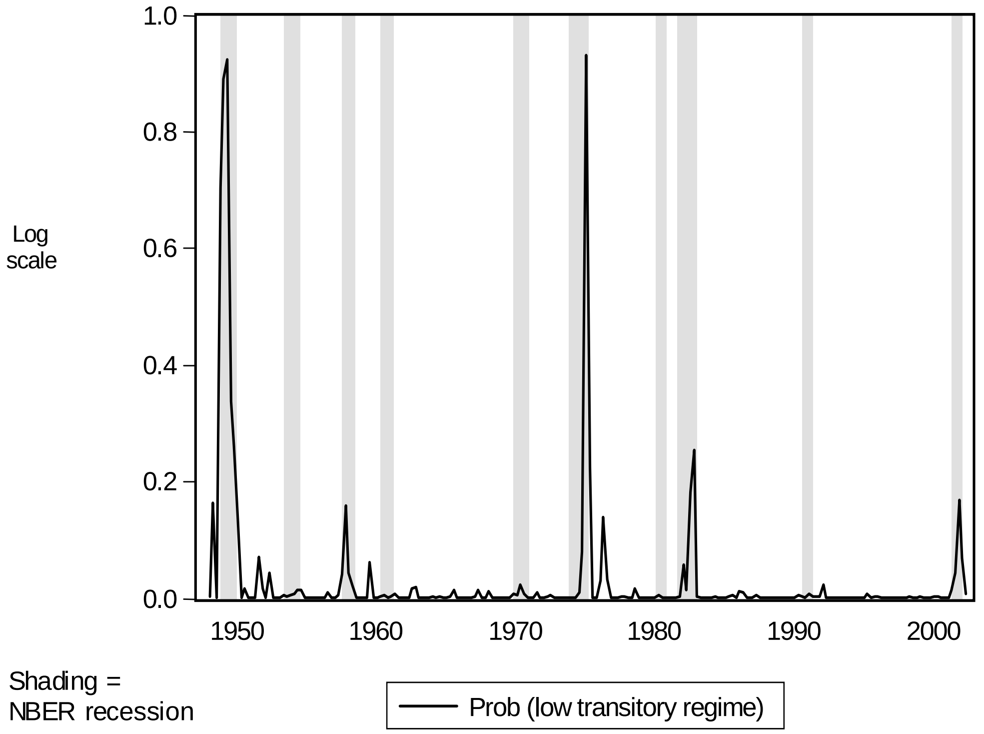

Graph 3B

Total common transitory component

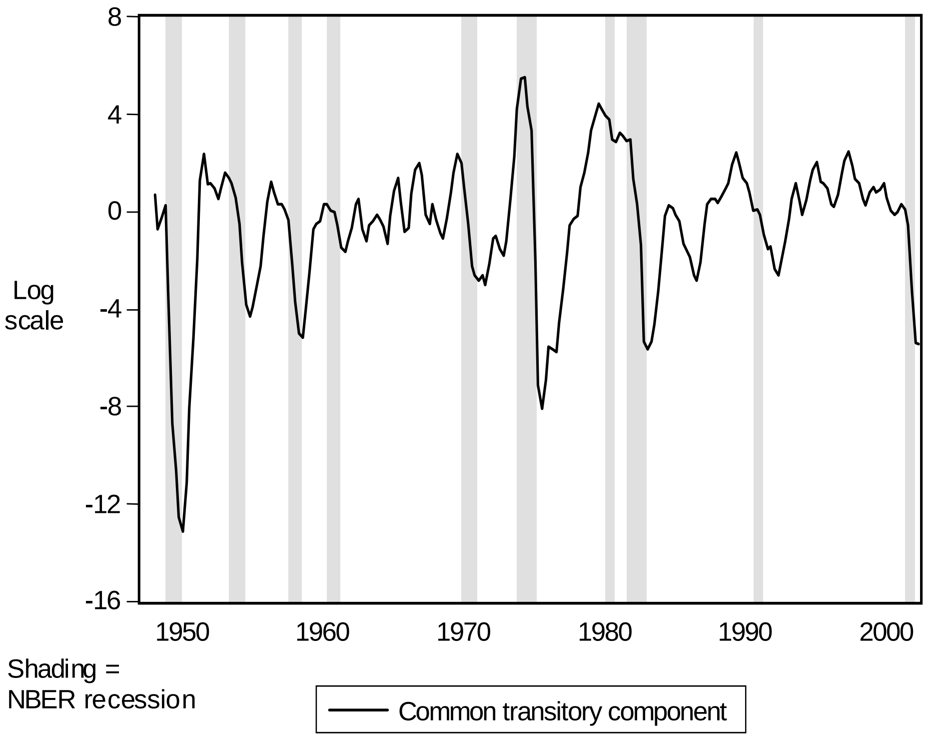

Graph 4A

Productivity and its permanent component

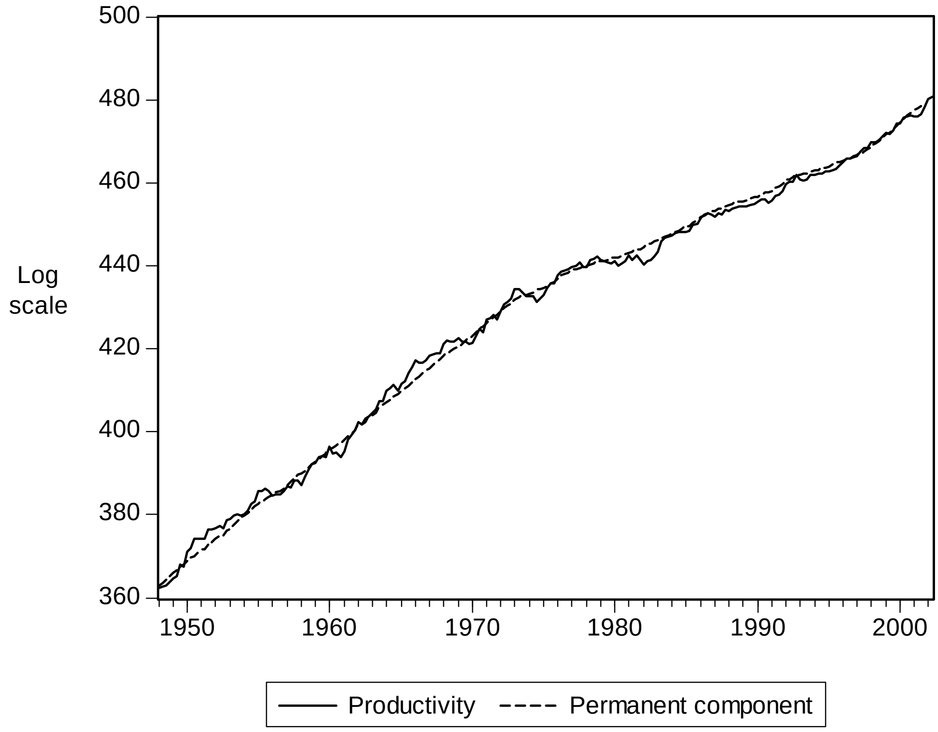

Graph 4B

Transitory and idiosyncratic components of productivity

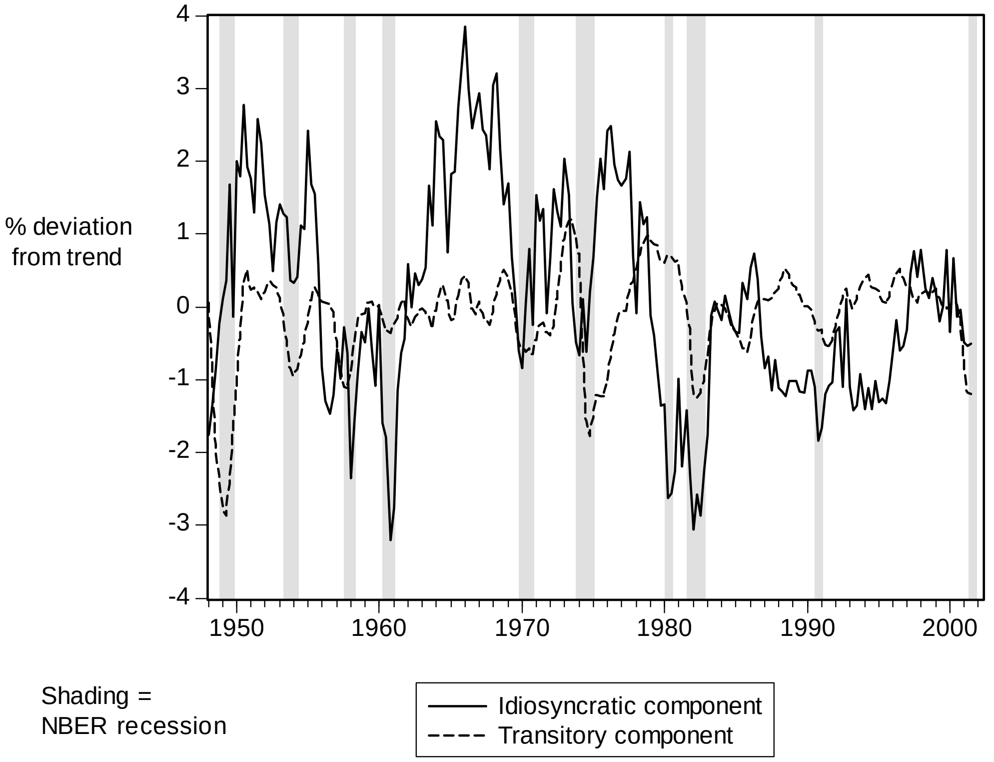

These differences are reflected in the retrospective regime assessments (Graph 5). Compared to the previous estimates, the estimated probability spends a lot more time in "grey areas" between transitions, and, most importantly, ends up well below 0.5, meaning that the two-variable system still cannot tell which growth regime the economy resides in. The difficulty in distinguishing trend from cycle is also manifested in the volatility of the estimated permanent component (Graph 6A). Now most of the movements in productivity are viewed as permanent, because there is little other information in the two-variable system to help filter out the noise. While the two-variable system arguably does a better job of getting at the transitory component (Graph 6B), as the remaining idiosyncratic component is smaller than in the four-variable system, from the point of view of identifying changes in underlying trends, the simpler approach appears to fall short.

Graph 5

The importance of additional variables

Prob (high-growth regime)

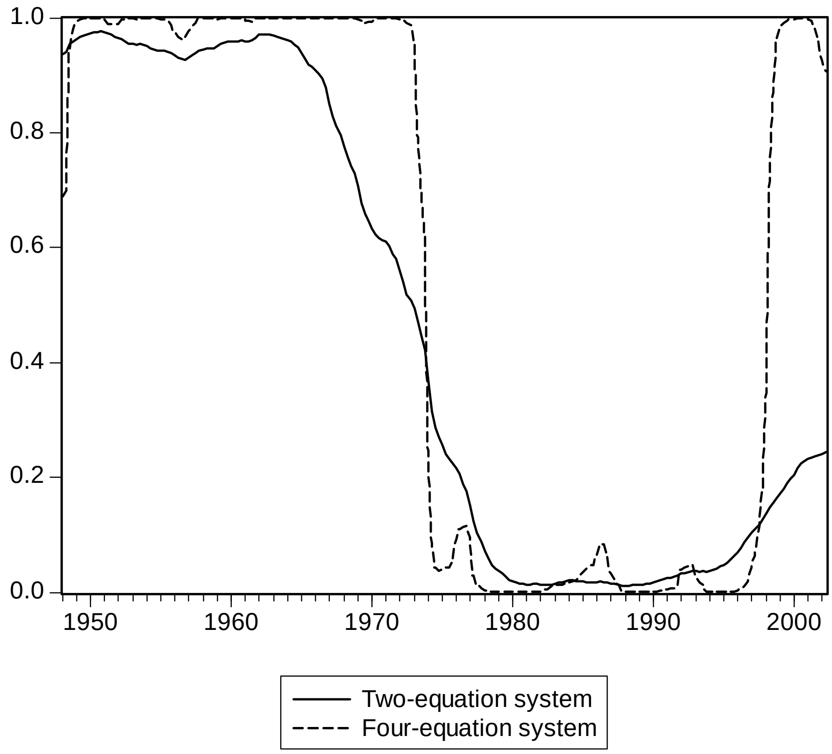

Of course, we have not shown that the two-variable system does a poorer job of characterising productivity trends. Perhaps the trends really are uncertain and the four-variable system is "overfitting" the data. To address this, we use the two models to forecast productivity at one- to four-quarter horizons.[11] The idea here is that if the four-variable system is overfitting the data, the root mean square error of its forecasts will be larger than that of the two-variable system. We constructed forecasts beginning in 1972 Q2 using the formulas outlined in Hamilton (1994). Consequently, the sample for the one-quarter forecast horizon covers the period 1971 Q3-2001 Q4, while the sample for the four-quarter forecast horizon covers the period 1972 Q2-2002 Q2. To obtain one-year horizon forecasts, we simply added the one- to four-quarter forecasts. Table 2 provides the results of this exercise. For both forecast horizons, the root mean square error of the two-variable system is larger, suggesting that the four-variable system is doing a better job, notwithstanding the fact that its primary intention is not short-term forecasting.

11 Some caution may be needed with our use of the term "forecast". Because the model involves a three-quarter lag in the productivity growth series, only the observation relating to the four-quarter-ahead forecast would not be contained in the current information set.

Graph 6A

Productivity and its permanent component (two-equation system)

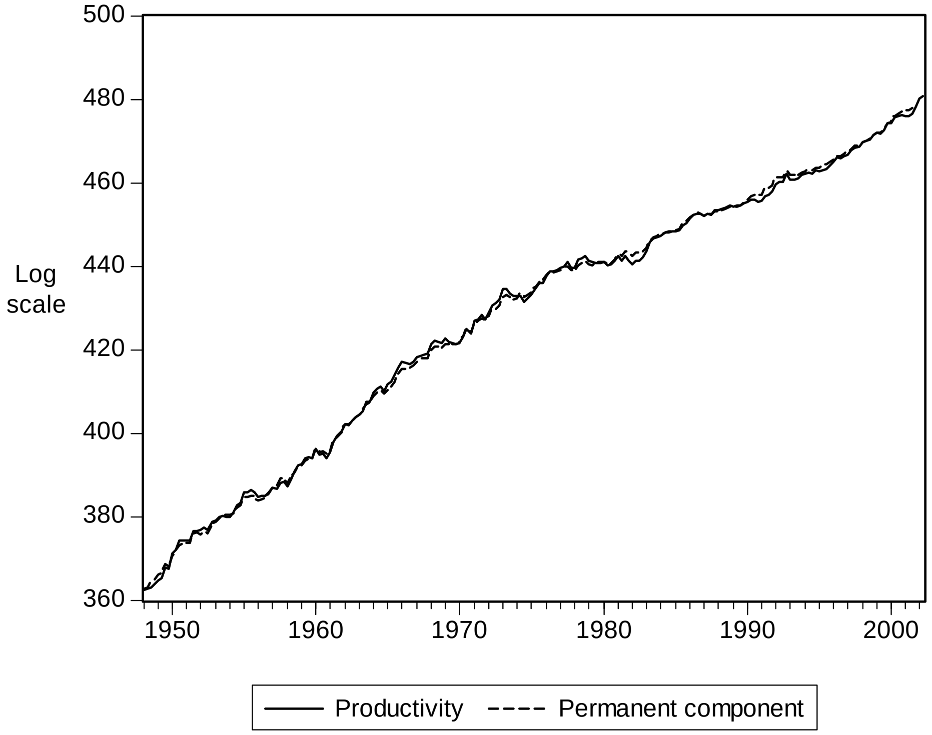

Graph 6B

Transitory and idiosyncratic components of productivity (two-equation system)

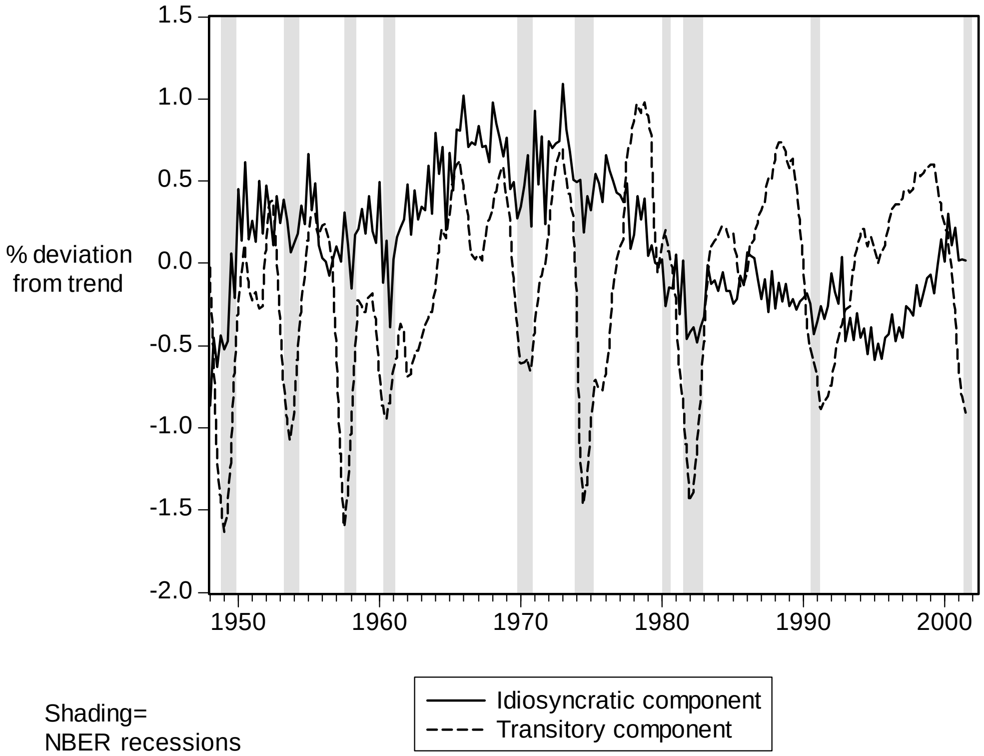

Table 2

Forecast performance

	Two-variable system	Four-variable system
	Root mean square errors	
One-quarter horizon	2.999	2.940
One-year horizon	1.503	1.432
	Non-farm productivity growth forecast (four quarters ahead; in percentages)	
2002 Q3	2.5	3.4

Note: All figures are at annualised percentage rates.

Table 2 also provides forecasts (four-step-ahead, because productivity enters the model lagged by three quarters) of productivity growth for 2002 Q3. The four-variable system predicts substantially stronger productivity growth (3.4% vs 2.5% on an annualised basis), primarily because of its assessment of a relatively high probability of being in a high-growth regime. Of course, productivity growth is a very noisy series, so a forecast for any given quarter cannot be made with much confidence, but the difference reflects different assessments of underlying trends, and provides some idea of the potential usefulness of the approach adopted in this paper.

5. Conclusions

The view that higher productivity growth is likely to be sustained has only really gained something approaching a consensus with the recent recession. Prior to 2001, one could more easily argue that the increased growth rates experienced since 1995 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. Economists faced the same difficulty (albeit in the opposite direction) in the mid-1970s when the dramatically slowing productivity growth coincided with a severe recession.

We explore the issue of the long-term trend in productivity by adopting a modelling strategy that integrates both theoretical considerations and recently developed statistical methods. In contrast to previous univariate analyses, we undertake a multivariate analysis in which we exploit information from additional variables that should be helpful in the identification of the trend and cycle in productivity. Specifically, we extend the data set to include consumption and labour compensation as well as detrended hours.

For our empirical framework we adopt the regime-switching dynamic factor model recently proposed by Kim and Murray (2002). This approach has a number of attractive features. First, it allows for the estimation of a common permanent component and a common transitory component, consistent with our interest in the trend and cycle in the productivity data. Second, the model allows for rich dynamics and can account for periodic changes in the underlying processes generating the common components. The latter consideration is not only important for providing a better characterisation of the data, but is central to any discussion about a possible shift to the secular growth rate of productivity. Last, the nature and timing of the regime changes is determined as an outcome of the estimation procedure rather than imposed a priori. In fact, one could view the implied regime changes and their reasonableness as an additional metric by which to judge the adequacy of our modelling strategy.

We find strong support in the data for the notion that the economy (and productivity growth in particular) switched from a relatively low-growth to a high-growth regime in the mid-1990s. The annualised difference between the mean growth rates in the two regimes is estimated to be approximately 1.5%. We also show that these techniques could have provided conclusive signals of the regime shift by the end of 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.

In subsequent work we hope to improve on our method for dealing with the lead-lag relationships in the data. In this paper we simply entered productivity and labour compensation lagged by three quarters to reflect the fact that these two series appeared to lead the others by that amount over the business cycle. This results in our approach ignoring the last three quarters of data for these series, a potentially crucial omission in any effort to assess current conditions. We also plan to deal with other aspects of the data that we have neglected here - for example, the apparent decline in volatility of productivity growth over time. Finally, we are also constructing a truly “real time” data set for productivity, one that recreates the data actually available at each point in time, and can therefore more definitively answer the question of how well this methodology can detect shifts in underlying growth trends.

Appendix A

The first-order conditions for the maximisation problem are as follows:

$$c + \dot{k} + (n+g)k = F(k,\ell)$$
$$1/c = \lambda$$
$$\Lambda v'(1-\ell) = \lambda F_2(k,\ell)$$
$$\dot{\lambda}/\lambda = n + g - F_1(k,\ell)$$

where λ is the shadow value of the resource constraint embodied in the first equation, and the other variables are as defined in the paper. In a steady state both λ and k will be constant, so the system

$$c + (n+g)k = F(k,\ell) \tag{1}$$
$$1/c = \lambda \tag{2}$$
$$\Lambda v'(1-\ell) = \lambda F_2(k,\ell) \tag{3}$$
$$0 = n + g - F_1(k,\ell) \tag{4}$$

can be solved for the steady-state values of c, k, λ and ℓ for fixed values of g and Λ.

Proposition 1: *The per capita quantities C/N, K/N and Y/N all grow at rate g in the steady state*.

Proof: This follows from the constancy of c, k and y, and their definitions.

Corollary 1: Since L/N is constant, C/L, K/L and Y/L also grow at rate g in the steady state.

The following results relate to scenarios in which Λ and/or g undergo permanent changes.

Proposition 2: *The steady-state values of* c/ℓ, k/ℓ *and* y/ℓ *do not depend on* Λ.

Proof: The homogeneity of F implies that (1) and (4) can be expressed in terms of c/ℓ, k/ℓ and y/ℓ. Since Λ only enters (3), and $F_2(k/\ell)$ is only a function of k/ℓ, ℓ and λ can vary with Λ to keep $\Lambda v'(1-\ell)/\lambda$ constant, and c can vary with λ to keep (2) satisfied.

On the other hand, a permanent change in g clearly cannot be so "neutral", except with respect to labour supply, as the following results show.

Lemma: With Cobb-Douglas technology, the steady-state value of ℓ does not depend on g.

Proof: Let $y = k^{\alpha}\ell^{1-\alpha}$. Then $F_1 = \alpha(k/\ell)^{\alpha-1}$, $F_2 = (1-\alpha)(k/\ell)^{\alpha}$ and $y/\ell = F(k/\ell)/\ell = (k/\ell)^{\alpha}$. It follows from (1) and (4) that $F_2/c = (1-\alpha)/\alpha$ for any value of g. Substituting this result and (2) into equation (3) proves that ℓ is constant.

Proposition 3: *The steady-state values of* c/ℓ, k/ℓ, y/ℓ *and* F_2 *vary with* g.

Proof: Equation (4) implies immediately that k/ℓ depends (negatively) on g. Since y/ℓ is a function of k/ℓ, it changes as well. Equation (3) then implies that λ must change with g, since the Lemma shows that ℓ is constant. Equation (2) then implies that c/ℓ must vary with g.

Corollary 3: The permanent movements in c, F_2 and y with respect to changes in g are proportional to each other, but the change in k is not.

The intuition behind Proposition 3 (and the corollary) is simply that higher growth makes individuals feel wealthier, and therefore increases desired consumption. The interest rate (ie the marginal product of capital) must then rise to counter the increased demand.

Appendix B

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

Measurement equation: $\Delta y_t = H\xi_t$ and $\Delta y_t = (\Delta y_{1t}, \ldots, \Delta y_{4t})'$

Transition equation: $\xi_t = \alpha(S_t) + F\xi_{t-1} + V_t$

with $E(V_t V_t') = Q$

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

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

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

$$F = \begin{bmatrix} \phi_1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_1^* & \phi_2^* & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \psi_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \psi_{21} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \psi_{31} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \psi_{41} \end{bmatrix} \qquad Q = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_1^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_2^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_3^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sigma_4^2 \end{bmatrix}$$

References

Corrado, C and L Slifman (1999): "Decomposition of productivity and unit costs", *American Economic Review (Papers and Proceedings)*, no 89, May, pp 328-32.

Friedman, M (1964): "Monetary studies of the national bureau", *The National Bureau enters its 45th year, 44th Annual Report*, pp 7-25; reprinted in (1969): *The optimum quantity of money and other essays*, by Milton Friedman, Chapter 12, pp 261-84, Chicago: Aldine.

——— (1993): "The 'Plucking Model' of business fluctuations revisited", *Economic Inquiry*, April, pp 171-77.

Gordon, R (2000): "Does the 'new economy' measure up to the great inventions of the past?", *Journal of Economic Perspectives*, vol 4, no 14, Fall, pp 49-74.

Hamilton, J D (1989): "A new approach to the economic analysis of nonstationary time series and the business cycle", *Econometrica*, no 57, pp 357-84.

——— (1994): *Time Series Econometrics*, Princeton University Press.

Hansen, B E (2001): "The new econometrics of structural change: dating breaks in US labour productivity", *Journal of Economic Perspectives*, no 15, Fall, pp 117-28.

Hodrick, R and E Prescott (1980): "Postwar US business cycles: an empirical investigation", Carnegie-Mellon University working paper.

Kaldor, N (1961): "Capital accumulation and economic growth", in F Lutz and D C Hague (eds), *The Theory of Capital*, New York: St. Martin's Press, pp 177-222.

Kim, C-J and C Murray (2002): "Permanent and transitory components of recessions", *Empirical Economics*, no *27*, pp 163-83.

Kim, C-J and C R Nelson (1999): "Friedman's plucking model of business fluctuations: tests and estimates of permanent and transitory components", *Journal of Money, Credit and Banking*, 31: Part 1, pp 317-34.

Roach, S S (1998): "No productivity boom for workers", *Issues in Science and Technology*, Summer.

Solow, R M (1956): "A contribution to the theory of economic growth", *Quarterly Journal of Economics*, no 70, pp 65-94.

Stiroh, K: "Information technology and the US productivity revival: what do the industry data say?", *American Economic Review*, forthcoming.