Estimating changes in trend growth of total factor productivity: Kalman and H-P filters versus a Markov-switching framework

Mark W. French

September 6, 2001

Abstract
Trend growth in total factor productivity (TFP) is unobserved; it is frequently assumed to evolve continuously over time. That assumption is inherent in the use of the Hodrick-Prescott or Bandpass filter to extract trend. Similarly, the Kalman filter/unobserved-components approach assumes that changes in the trend growth rate are normally distributed. In fact, however, innovations to the trend growth rate of total factor productivity are far from normal. The distribution is fat-tailed, with large outliers in 1973. Allowing for those outliers, the estimated trend growth rate changes only infrequently. A nonlinear filtering approach is probably better suited to capturing the infrequent past and possible current shifts in trend growth of TFP. One such approach is the Markov-switching model, which is estimated and tested in this paper. The Markov-switching approach appears to have several advantages over repeated Andrews tests.

Board of Governors of the Federal Reserve System

Division of Research and Statistics, Mail Stop 80

20th and Constitution Ave., NW

Washington, DC 20551

I would like to thank Alex David, Jeremy Piger, David Reifschneider, Dan Sichel and Cheryl Varner for their helpful suggestions. The views expressed here are those of the author and do not necessarily reflect those of the Board of Governors or the Federal Reserve System.

I. Introduction and overview

Over the past five years, rapid growth of labor productivity helped hold down unit labor costs and thereby reduce inflationary pressures. However, productivity weakened a bit in the economic slowdown of early 2001, making it more difficult to estimate the trend pace of productivity growth. After the current downturn ends, will rapid labor productivity growth resume, thereby restraining the future growth of unit labor costs? The answer depends in part on the future rate of capital accumulation, but also depends in an important way on future growth of total factor productivity (TFP), that is, the "Solow residual" or "technical progress." Growth of TFP was impressive from mid-1995 through mid-2000. Was this just a temporary blip, or has there been a surge in the trend growth rate of TFP? This paper examines several methods of extracting the trend component of total factor productivity, in an attempt to find a "real-time" method that performs well in evaluating historical data and can be used as a guide to the current – and possibly future – trend in the economy's productive potential.

Section II of the paper lays out a simple state-space model of nonfarm business output that justifies the use of univariate filters to estimate the trend component of the Solow residual. Section III briefly comments on several linear univariate filters available for estimating trend. Section IV presents estimates using three of these linear filters.

Of the three linear filters, the Kalman filter works best in modeling TFP trends; however, none of the linear filters is fully adequate in modeling total factor productivity. The problem appears to be that innovations to trend growth of TFP are not normally distributed; in particular, significant changes in trend productivity growth appear to occur only at infrequent intervals. The state-space/Kalman filter approach instead assumes normally distributed innovations to trend growth. Similarly, the Hodrick-Prescott smoothing parameter implicitly assumes a path for trend TFP that is differentiable at any given point and whose growth rate evolves continuously.

To capture the discrete nature of significant shocks to trend TFP, I turned to a logical nonlinear alternative – a variant on the Hamilton (1989) Markov-switching model. Section V presents estimation results for such a model of trend TFP growth. This model's mean estimate for the path of trend TFP seems to capture nicely the well-documented slowdown in productivity growth in the early 1970s.

II. Calculating trend productivity consistent with a Cobb-Douglas production function

Given estimates of the capital stock, labor quality, and the equilibrium labor share, calculating trend labor productivity amounts to calculating the trend in the Solow residual. One can do this filtering exercise in several ways, but all start from the same basic model:

- Nonfarm business output (Q) is the product of trend (Q_{trend}) and gap (Q_{cycle}).
- Hours supplied (H) is the product of trend hours supplied (H_{trend}) and a stationary cyclical term (H_{cycle}).
- Trend labor (L_{trend}) is the product of trend hours supplied (H_{trend}) and the quality of each labor hour (L_{qual}).
- A two-factor Cobb-Douglas production function holds in equilibrium.[1]
- The Solow residual (S) includes cyclical factors as well as trend in total factor productivity (TFP_{trend}).

These five definitions are expressed in equations 1 - 5:

(1) $Q = Q_{trend} * Q_{cycle}$

(2) $H = H_{trend} * H_{cycle}$

(3) $L_{trend} = H_{trend} * L_{qual}$

(4) $Q_{trend} = TFP_{trend} * L_{trend}^{\alpha} * K^{1-\alpha}$

(4a) $Q = S * (H * L_{qual})^{\alpha} * K^{1-\alpha}$

The observed Solow residual is calculated by inverting equation 4a. Equations 1 - 4a imply the following relationship between the observed Solow residual and the underlying trend in total factor productivity:

(5) $S = [Q_{cycle} / H_{cycle}^{\alpha}] * TFP_{trend}$

[1]The measure of nonfarm business output is an extension of the BLS published measure, as defined in the appendix. Labor quality is taken to be exogenous; likewise, the capital stock is assumed given. For this exercise, I used the BLS capital-stock series, interpolated to a quarterly frequency, as described in the appendix. Also, I have used an equilibrium labor share of 68 percent, similar to the long-run average of the series over history.

The objective is to extract TFP_{trend} , the unobserved trend in total factor productivity shown in equation 5, from the Solow residual.

III. Univariate methods of extracting trend

A variety of approaches have been used over the years to estimate potential output and trend productivity. Laxton and Tetlow (1992) give a helpful description of the evolution of these approaches. Peak-to-peak interpolation was common in the 1960s (and was even used by DeLong and Summers as late as 1988). From the early to mid-1970s, the deterministic trend was commonly used in describing labor productivity or trend output. From 1977 to 1982, several authors began to decompose the sources of trend output growth using aggregate production functions; in essence, the problem shifted from estimating trend labor productivity to estimating trend total factor productivity. Over this period, polynomial or split time trends were most commonly used. In the 1980s and thereafter, trends in output and productivity have frequently been estimated using a stochastic-trend model. (Perron (1989, 1997) is a noteworthy exception; he argued that stochastic trend had not been shown superior to the split trend model.) All of these methods, discussed in more detail below, assume that output or productivity can be separated into distinct trend and cyclical components.

Peak-to-Peak Interpolation

This method has the advantage of simplicity. However, it relies on the questionable assumption that the economy operates at the same level of excess demand at all cyclical peaks. For long expansions, this assumption is harmless, and the procedure works reasonably well. For short expansions or periods shortly after recessions, it can be misleading. A version of peak-to-peak interpolation was used by Braun (1990).

Ad hoc segmented trend

Like the Bureau of Labor Statistics, Oliner and Sichel (2000) show trends in productivity by calculating averages over several intervals: before 1973, 1973-90, 1990-95, and 1995 onward. The choice of intervals is a bit arbitrary, but this approach is certainly easy to implement. One problem for purposes of this study is that the above time periods were chosen to match breaks in labor productivity rather than in total factor productivity. The 1990 break in labor

productivity had little to do with total factor productivity, according to Sichel. Thus, for total factor productivity, the relevant break points chosen in their paper are 1973 and 1995.

Single deterministic trend

This method also has the advantage of simplicity. However, the evidence is strong (even when one accounts for such factors as labor quality and capital deepening) that the trend growth rate of TFP has not been constant over time. There was a persistent change in average growth of output and productivity beginning in the early 1970s. The test suggested by Andrews (1993) makes this fairly clear: An Andrews test for a significant shift in the growth of total factor productivity indicates a break in average growth in the third quarter of 1973, significant at the 90 percent confidence level.[2] Are there other breaks in trend TFP growth over the past 40 years? As a crude check, I made a level adjustment to average TFP growth post-1973, equalizing the mean growth rates of the two subsamples. I then performed an Andrews test on the adjusted overall sample. The most likely remaining break point was in 1991:Q2, but the Andrews test statistic was less than half the 90 percent critical value of 7.17; the 1991 break also failed to pass a less stringent (but inappropriate) t-test at the 90 percent level. I take this result as suggestive that there are no statistically significant breaks in the TFP sample other than the one in 1973.

Many economists (for example, Oliner and Sichel (2000) and Gordon (2000)) would nonetheless say that there have been changes in output and productivity trends more recently than 1973. At a minimum one needs a methodology that accounts for the possibility of such additional changes. Three such encompassing methodologies are the Hodrick-Prescott (H-P) filter, Christiano's bandpass filter, and trend-cycle decomposition with state-space modeling and the Kalman filter.

Hodrick-Prescott filter

The Hodrick-Prescott filter estimates trend by smoothing – in effect, by taking a

[2]To carry out this test, I regressed BLS-based estimates of quarterly growth in total factor productivity against a constant (no evidence of autocorrelation in the errors), over a full sample from 1959:Q2 through 1998:Q4 and over two subsamples. Following Stock and Watson (1998), I chose the split point between the two regressions to start 15 percent into the sample and end 15 percent from the end of the sample. Thus, the split point between the two subsample periods ranged from 1965:Q2 to 1993:Q2. I used critical values for the Quandt Likelihood Ratio from Andrews (1993). The maximum value of the likelihood ratio was 8.5754, implying a value for the Andrews lambda of 7.62, greater than the 90 percent critical value of 7.17.

weighted moving average of the original series, where the moving average is symmetric and centered. In practice, one must choose how smooth the resulting trend should be. Hodrick and Prescott suggest a smoothness parameter of 1600 for quarterly data, and that parameter value is most commonly used in practical applications. In some applications, however, researchers have used a parameter value of 6400 or higher; as the smoothness parameter becomes very large, the H-P filter constructs a linear trend.

The Hodrick-Prescott filter is quite simple to implement; however, it also has potentially large drawbacks. One would like a filter to generate estimates of the cycle which are close to the actual values for the cycle. Unfortunately, the HP filter is optimal in this sense only under two fairly restrictive assumptions. First, the data must, a priori, be known to have an I(2) trend. Otherwise, the filter will generate shifts in trend growth rates when they don't exist in the raw data. If instead there are only one-time permanent shocks to the level of trend, or a constant (split) trend growth rate, or both – the H-P filter distorts the cyclical properties and the higher moments of the data in a significant way – a point discussed by King and Rebello (1993) and Cogley and Nason (1995). This restriction on usage of the H-P filter (to I(2) cases only) is fairly serious. If one is not committed to a time-varying trend growth rate, other approaches are preferable to H-P, because they allow for the possibility of a fixed (split) trend growth rate.

Second, the H-P filter is optimal only if the cycle consists of white noise or if the identical dynamic mechanism propagates changes in the trend growth rate and in the innovations to the business cycle component. King and Rebello (1993) judge both these conditions "unlikely to be even approximately true in practice". Of course, even the best of alternatives to the H-P filter are unlikely to be optimal either, as misspecification is a virtual certainty. In fact, Baxter and King (1999) point out the similarity between the H-P(1600) filter and a bandpass filter with cutoff at 16 or 32 quarters. They point out, however, that the H-P filter does a relatively poor job of picking out the evolution of low-frequency (i.e. trend) movements, those greater than 32 quarters.

Another oft-mentioned drawback to the H-P filter is that in most published applications, the choice of the smoothness parameter typically has been rather ad hoc. Hodrick and Prescott (1980) showed that under certain restrictive assumptions, the optimal smoothing parameter equals the ratio of the variance of shocks to the cycle divided by the variance of shocks

to (growth rate of) trend.[3] They held a "prior view that a five percent cyclical component is moderately large, as is a one-eighth of one percent change in the rate of growth in a quarter. This led us to . . . 1600 as a value for our smoothing parameter." In 1990, Kydland and Prescott suggested that with a value of 1600, "the implied trend path for the logarithm of real GNP is close to the one that students of business cycles and growth would draw through a plot of the series."

Calculation of the Hodrick-Prescott variance ratio over the period 1960:Q4 through 1990:Q4 suggests that a choice of smoothing parameter of 1600 is too low for total factor productivity. First, I did a rough estimate of the denominator – the variance of shocks to the growth rate of trend TFP – following Stock and Watson (1998). For total factor productivity from 1959 through 1999, the Stock-Watson method yielded a standard deviation of shocks to trend growth of 0.035 percentage point (quarterly rate).[4] Next I calculated the numerator of the Hodrick-Prescott ratio – the square of the standard deviation of TFP around trend. The standard deviation of TFP around stochastic trend is likely to be somewhat less than the standard deviation around a linear deterministic trend (3.64 percentage points) and is probably closer to the standard deviation around a trend split in 1973 (2.21 percentage points). Thus, as a very rough cut, the optimal smoothing parameter for TFP is probably less than $(3.64/.035)^2$ – that is, 10816 – and is probably close to $(2.21/.035)^2$, or 3987.

Uncertainty about an appropriate full-sample smoothing parameter is not the only problem in using the H-P filter. The filter has no good method for determining the appropriate gain on new datapoints at the end of the sample, when the filter is no longer two-sided – a marked drawback relative to the state-space/Kalman filter approach, at least in theory. As reviewed in

---

[3]In their construction of the optimal smoothing parameter, they first removed a linear deterministic trend from the data. Then, to calculate the remaining variance of trend, they assumed that the cyclical component was white noise and that the only shocks to trend were shocks to the growth rate of trend.

[4]I began with the data for growth of the Solow residual from 1959:Q2 to 1998:Q4, and truncated the first and last 15 percent of the sample, as Stock and Watson recommend. I then performed a sequence of Chow tests for a regression of TFP growth against a constant. (The partial autocorrelograms of the regression errors suggested that they were approximately white noise, by and large, needing no autoregressive correction.) For this sample, the Stock-Watson test statistic (using the Quandt Likelihood Ratio, the supremum of all the F-statistics for the Chow tests) turned out to be 8.58. (The supremum occurred at the break between the periods 1973:Q3 and 1973:Q4.) This value implied a ratio of I(2) to I(1) standard deviation of .0485. Stock and Watson combined their ratio with a range of estimates for the I(1) standard deviation. I approximated the I(1) standard deviation as the standard error of the estimate of a full-sample regression of growth of TFP against a constant, a post-1973:Q3 dummy, a recession dummy, and a rapid recovery dummy (with ARMA(1,1) errors assumed). (The two business cycle dummies are defined in section V below). The product of the resulting standard error (.715 percent) and .0485 yielded the estimated I(2) standard deviation of .035 percentage point (quarterly rate).

Hamilton (1994, p. 385), the Kalman filter is the optimal one-sided filter, with optimal estimate of gain, for any process that linearly links the observed variables with state variables and exogenous and/or lagged dependent variables (i.e. a linear state-space process, with any order of integration). It is optimal among all filters, linear or nonlinear, if in addition the state-space model's errors are Gaussian.

One can hope that the distortions created by the H-P filter are small: if they are, its simplicity is a major recommendation. Furthermore, even if the H-P filter is suboptimal, it is not clear that the alternatives are better – provided one assumes an I(2) trend a priori. Misspecification problems can make any approach to estimating trend suboptimal – even the two purportedly superior filters discussed next, the Bandpass and Kalman filters.

Bandpass filter

In a recent paper, Christiano and Fitzgerald (1999) discussed a bandpass filter that they claim performs inherently better than the H-P filter, especially when the filters become nearly one-sided near the end of the sample period. I used Fitzgerald's Matlab software to filter the log of TFP. I set the parameters of the bandpass filter so as to make all cycles eight years and longer part of trend TFP, to come close to the shape of H-P (1600) trend. In effect, the Bandpass filter's assignment of all low-frequency cycling to the trend ensures that the Bandpass trend will be I(2), like the H-P filter.[5] Thus we get the same caveat as with the H-P filter: if the underlying trend might not be I(2), other filtering methods may be preferable. As discussed in section IV, results using the Bandpass filter were similar to those using the H-P (1600) filter.

Kalman filter / state-space model

This approach uses historical data to estimate the share of the variance (of the Solow residual) to assign to trend rather than cycle. In principle, the approach can also determine whether or not the I(2) component of trend is significantly different from zero – a major advantage over the H-P and Bandpass filters. In practice, the state-space/Kalman filter approach has an important drawback: its variance estimates are biased toward zero when the true variance is small but nonzero, as demonstrated by Stock and Watson (1998). More concretely, they found that this

[5]Christiano and Fitzgerald (1999, p. 11) note that the Bandpass filter typically estimates I(2) trends better than I(1) trends (conversely for the cycle).

approach is biased against identifying shifts in the growth rate of trend productivity or output. However, they offer an easy solution to this problem. They carried out sequential Chow tests on the intercept of an ARMA(4,0) model of GDP growth.[6] They compared the maximum of the F-statistics from these Chow tests (that is, an Andrews test statistic) against a table of critical values, and used the result to calculate the ratio of the standard deviation of the I(2) shocks relative to the standard deviation of the I(1) shocks.

I used the above Stock-Watson procedure, and imposed this ratio on the relative variances directly in state-space estimation.[7] One can instead follow the procedure described in footnote 4 to obtain nonzero estimates of the separate standard deviations of the I(1) and I(2) shocks to trend. That procedure yielded smaller estimates of the two standard deviations, as noted in footnote 4.

IV. Filtering results: A comparison of estimated trends in TFP

In this section I compare trends in total factor productivity estimated using three different filters – against each other and against the observed Solow residual.

Hodrick-Prescott estimates

Figures 1a shows an estimate of trend TFP derived using the optimal smoothing parameter for the H-P filter generated earlier by the Stock-Watson (1998) procedure: 3987. Figure 1b shows an estimate derived using the traditional H-P smoothing parameter, 1600. The trend line shown in figure 1b seems less likely to be accurate, because it puts a fair amount of the 1980-82 recession into estimated TFP trend. This result is problematic.[8] Figure 1c shows, for comparison, the estimate of trend TFP derived using a smoothing parameter of 6400. This value is above the Stock-Watson optimum, but the estimate does have the advantage of putting even less of the recession of the early 1980s into trend. However, with this smoothing parameter the slowdown

[6]With break points starting 15 percent into the sample, and ending 85 percent into the sample.

[7]The estimate of the I(1) standard deviation was .90 percent (quarterly rate) and that of the I(2) standard deviation was .0485*.90, or .044 percent (quarterly rate).

[8]Of course recessions can permanently reduce labor productivity by temporarily reducing the growth of the capital stock. Likewise, prolonged unemployment during recessions can erode workers' skills, thus reducing workers' productivity when they finally find new jobs. In principle, however, both these factors should be captured in estimates of labor quality and capital stock without affecting the estimate of trend TFP.

of trend TFP growth in the 1970s is estimated to be four or five years later than the split point having the highest Andrews test statistic.

Bandpass estimate

For the Bandpass estimate of trend shown in figure 2, I chose the parameters of the filter so as to put into trend only those cycles greater than or equal to eight years in length. As can be seen in figure 3, the H-P trend (smoothing parameter = 1600) and the Bandpass trend look very similar.

H-P versus Bandpass estimates

The key comparison for real-time analysis is how well each filter estimates trend toward the end of the sample period, relative to an "ex post" estimate calculated with full knowledge of actual outcomes at least three years in the future. Such an assessment requires calculating the two filters over historical subsamples that are truncated at various points before the final three years of history, and then comparing the difference between trend calculated over the subsample versus trend calculated over the subsample plus three succeeding years. The filter that, on average, has the smaller revision in the final period of the original subsample, given three further years of data, is preferable. Such historical simulations using the H-P and Bandpass filters revealed no obvious superiority of one filter or the other: The two filters had about equal numbers of "wins" in the various subsamples selected. One could carry out a full rolling test of this proposition over all truncation points, but the lack of a clear winner seemed evident enough without consuming the extra time to cover all endpoints.[9] In any case, neither filter did particularly well in generating end-of-sample estimates of trend. For both filters, the final trend point was typically revised sharply, given three more years of data.

Kalman estimate

For the univariate Kalman procedure, I first imposed a prior restriction on the ratio of the variances of I(1) shocks and the variance of I(2) shocks, as suggested by Stock-Watson. The

[9] A similar rolling test of the Kalman filter is just a matter of rolling estimation of a straight-line trend. As such, the Kalman filter would almost surely emerge a winner over the other two filters by this criterion of minimal endpoint revision, as a straight-line trend with no splits typically changes very little with the addition of new datapoints.

estimated model is shown in the following table:

UNOBSERVED COMPONENTS (Stochastic Trend) MODEL OF TFP	
(All state variables are in logs. Sample period 1960:Q4 to 1999:Q4) (t-statistics in parentheses)	
TFP Cycle = .94 TFP Cycle [t-1] + I(0) $Shock_{TFP}$ (2.4)	Standdev(I(0) $Shock_{TFP}$) = .92/100 (1.0)
Level of TFP Trend = Level of TFP Trend [t-1] + Trend growth rate of TFP + I(1) $Shock_{TFP}$	Standdev(I(1) $Shock_{TFP}$) = .9204/100 (1.0)
Trend Growth Rate of TFP = Trend Growth Rate of TFP [t-1] + I(2) $Shock_{TFP}$	Standdev(I(2) $Shock_{TFP}$) = .9204*.0485/100

Given this model, the Kalman filter and smoother generated the estimate of trend shown in figure 4. The estimated trend line has two problems, one fairly obvious and the other subtle but much more important. The first problem is that the estimate of trend appears to be biased downward in the period 1963 to 1979. The more important problem is discussed in the next subsection.

Failure of the stochastic trend model of TFP

A fundamental problem exists with the univariate stochastic trend/Kalman model of TFP: underlying innovations in the trend growth rate are not normally distributed.[10] Evidence of this problem appears in at least three ways:

1. **Andrews tests.** The Andrews tests discussed in section III (on which the Stock-Watson test relies) indicated a significant break in trend TFP growth at the third quarter of 1973 but showed no compelling evidence of breaks in any other quarters. Furthermore, once a dummy was included in the Andrews test[11] (to allow for the

[10]The infrequent occurrence of significant innovations in trend growth creates a similar problem for the other two linear filters used here.

[11]Bai (1997) carried out a similar sequential Andrews test procedure.

1973 split in growth rates), the Stock-Watson estimate of the ratio of remaining I(2) to I(1) variance was microscopic.

2. **Outliers and Kurtosis.** Given the results of the two Andrews tests, one might expect the significant 1973 break found by the Andrews test to be matched by an outlier in the distribution of estimated shocks to the Kalman trend growth rate – and indeed there were standardized shocks of 2.7 and 2.6 standard deviations in the first and second quarters of 1973. In an iid normal distribution, such large successive shocks would be extremely unusual in a sample of only 157 observations. The 1973 outliers suggest that the unobserved component/Kalman model is misspecified in its treatment of trend growth rates. Indeed, the excess kurtosis of the standardized shocks to the trend growth rate was 1.78; the adjusted kurtosis test statistic of Harvey and Koopman (1992) indicates more than 99 percent probability of the presence of non-normality (that is, a fat-tailed distribution).[12]

3. **With the addition of a single dummy variable, the I(2) state-space model reduces to the special case of a split trend.** A further test of the stochastic trend issue is equally suggestive. As a first step, use the Kalman filter to estimate the TFP model – imposing only the relative restriction on the two standard deviations, **not** separate restrictions on each. This procedure yields fairly large estimates of the two standard deviations relative to the doubly restricted example above; the I(1) standard deviation is now estimated at .90 percent and the I(2) at .044 percent. Now add to the trend growth equation a dummy variable equal to 1 for 1974 onward. Within six iterations of the likelihood maximization procedure, the estimate for the I(1) standard deviation drops from .72 to .04; over many interations, it gradually approaches 0.

The conclusion from these alternative Kalman estimates: If the 1973 shock to TFP growth is omitted from the set of trend growth rate innovations, there is no statistically significant evidence for stochastic shifts in trend TFP during the sample period. If that shock is not omitted,

[12]To demonstrate excess kurtosis, higher values of the excess kurtosis measure are needed than in the standard test, due to the autocorrelation in residuals created by the smoothing procedure (Harvey and Koopman, 1992). Notwithstanding, the evidence for kurtosis was very strong.

the distribution of innovations to trend growth is not normal. If the distribution is not normal, there is no assurance that the TFP model's parameter estimates or the estimates of the state are optimal. Quite the opposite – the stochastic trend model tends to assign innovations continuously to trend TFP growth, even though there is no evidence that such innovations to trend TFP growth are anything but rare.

A deterministic split trend model of TFP may be superior to the unobserved components model, for the reasons discussed above. It has one major drawback, however, as mentioned earlier: It cannot help determine whether a new trend shift may be occurring near the end of the sample. An extension of the split trend model **can** make such a determination: the Markov switching model presented in section V.

Discussion

None of the three linear filters was completely satisfactory. The bandpass and H-P filters produced similar results. Given the "optimal" H-P smoothing parameter of 3987, both filters imposed I(2) swings on the TFP trend. Neither did particularly well in generating end-of-sample estimates of trend: in both cases, given three more years of data, the final trend point was typically revised sharply. The eight-year bandpass filter picked up a lot of the 1980-82 downturn and put it into trend – as did the H-P (3987) filter, albeit to a lesser extent. An argument could be made that the 1980-82 recession was long enough to destroy the human capital of laid-off workers – but if and when those laborers later returned to work at a lower real wage, the loss in human capital should, in principle, have been picked up by a deceleration of or decline in the "quality of labor" term, rather than in TFP. In any case, I would not have expected the destruction of human capital to have created such a large downturn in measured TFP trend (nearly 3 percent) as seen with the Bandpass filter. More plausible, perhaps, is the notion that the 1979-81 surge in U.S. energy prices made a substantial portion of the measured capital stock obsolete – with the obsolescence showing up as a decline in TFP.

After the tests discussed in the preceding subsection, I added a 1973 dummy variable to the Kalman trend for TFP, and the stochastic trend collapsed to a deterministic trend line split in the third quarter of 1973. Next I followed Oliner & Sichel (2000) and added a further dummy variable to allow for an additional break in trend from 1995 onward. The resulting point estimate of the 1995 rise in TFP trend is economically significant, though not statistically

significant. Of course, the problem remains that the Kalman filter, even with the splits added, will assume that no variance in trend growth exists beyond the exogenous historical splits; thus, the Kalman filter will tend to put end-of-sample shocks into cycle rather than trend – the opposite of the H-P filter (1600), which tends to put too much of any new shock into trend.

I conclude that none of the filtering methods does a particularly good job of estimating any shifts in TFP trend near the end of the sample period – Another approach is needed. I tried incorporating extra information, by joint Kalman filtering of all the elements of the production side (TFP, hours, employment, labor force, and so forth). The results are not included here due to space constraints, but are available from the author on request. In any case, the multivariate Kalman filter with a dummy for a 1973:Q3 break in TFP still yields a split deterministic trend for TFP, just as the univariate version does.[13] This leaves us still unable to catch an end-of-sample shift in trend TFP growth.

The biggest remaining problem is to find a structure for modeling trend in total factor productivity that will capture the infrequent historical shift(s) – <u>and will also eventually detect any recent shift in trend TFP growth.</u> The Hamilton (1986) state-space framework is fundamentally inappropriate for the problem, as it is clear that TFP does not follow the standard stochastic trend/unobserved components model. Given that TFP does not follow the standard stochastic trend, the Hamilton (1986) model cannot efficiently detect late-sample shifts in trend growth of TFP. In fact, a stochastic trend model will probably detect too many such breaks.

Although the linear filters failed to successfully model shifts in trend TFP growth, it may be possible to model such shifts with greater success using an alternative to linear filtering. One approach was suggested by Filardo and Cooper (1996). They checked for a recent break in trend growth rate (given an early-sample break as well) using sequential Andrews tests, with critical values generated by the bootstrap.[14]

Hamilton's Markov-switching model (1994, pp. 685-702) is another alternative. The Markov-switching estimates provide more information than the estimates based on Andrews

---

[13]However, other supply-side trends fit better into the Kalman filter framework. Both the labor force participation rate and the workweek have an I(2) component, according to the Stock-Watson (1998) test. Thus the Stock-Watson restriction on relative variances was imposed in the joint state-space estimation and Kalman filtering process. The resulting estimated innovations to trend growth appeared normally distributed for both these variables.

[14]My thanks here to Norm Morin, a colleague at the Board who showed by example how to implement such an approach.

tests. The Markov approach provides a time series of estimates of the conditional probability of being in each trend-growth-rate regime, and the conditional mean of the estimated trend growth rate at each point in time. A key question is whether the Markov-switching approach is less slow than the Andrews-test approach in recognizing a given break in trend. This Markov approach is reviewed in the next section.

V. Modeling discrete innovations to TFP Growth: Markov switching

Results described earlier suggest that there was a slowdown in the trend growth of TFP sometime in the early 1970s. My state-space/Kalman model found no evidence of other significant shifts in trend TFP growth through the early 1990s. That is, when a dummy was included in the state-space model to reflect a trend split in 1973:Q3, the estimated variance of innovations to TFP growth dropped to zero. The resulting estimate of trend for TFP was simply a split deterministic trend line.

However, few economists would claim that 1973 was the last time trend TFP growth will ever change. More explicitly, many economists believe a pickup in TFP growth has occurred over the past four or five years. Thus, we need a model that allows for the possibility of further (albeit infrequent) changes in trend growth since the early 1970s. Such a model should evaluate the likelihood that any further trend change has occurred in recent quarters. A deterministic split trend model of TFP is unsatisfying because it will not search for further changes in trend growth. In contrast, the Markov-switching model of Hamilton (1989) fits the bill quite well. It can encompass split trend models while providing estimates of the probability of having arrived in a new growth regime. In this paper, I use a generalized version of the switching model in Hamilton (1994, chapt. 22), modifying slightly Hamilton's 1989 Gauss program to allow for dummy variables in the observation equation.

Hamilton's switching model as applied to trend TFP growth assumes that the observed growth rate of TFP will depend on the state of the economy (fast trend growth or slow trend growth) and on the deviation of observed growth from trend last quarter plus an error term. In the version used here, the conditional mean of TFP growth also depends on the stage of the business cycle. Evidence has accumulated steadily over the past decade that the stages of the cycle include not just recession and expansion, but a rapid recovery phase after recessions. For TFP, a rapid recovery phase seems quite clear after all six recessions in the sample period except the most

recent in 1990-91. In the following subsections, I present the estimation results, first using a simpler two-stage model and then using a "plucking" model (that is, a model with the three-stage business cycle that is more in line with current thinking). The latter gives results about the timing of switches in trend growth rates that are more in line with the standard view that trend growth slowed somewhere around the early 1970s.

With a two-stage model of the business cycle

This model assumes that there are two states of TFP growth: fast trend growth and slow trend growth. Furthermore, it assumes that trend growth is slowed by recessions. I define notation as follows:

y_t is observed TFP growth at time t;

s_t is the state of the economy:

$s_t = 1$ when trend TFP growth is rapid

$s_t = 0$ when trend TFP growth is slow

d_t is an exogenous dummy variable

$d_t = 1$ during NBER recession quarters, and 0 otherwise[15]

u_1 is the mean growth of TFP in State 1

u_2 is the mean growth of TFP in State 2

u_3 is the mean reduction in TFP growth during recession quarters

p_{11} is the probability of remaining in State 1 (rapid trend growth) for another quarter

p_{22} is the probability of remaining in State 2 (slow trend growth) for another quarter

p_{12} is the probability of moving to fast trend growth from slow trend growth the previous quarter

p_{21} is the probability of moving to slow trend growth from fast trend growth the previous quarter

Then the model can be summarized as follows:

[15] I modified the NBER definition slightly. If TFP growth was also negative in the quarter or quarters immediately preceding the recession, those quarters were included as part of the exogenous "TFP recession". Using similar logic, if TFP growth was faster than 3 percent (annual rate) in the last quarter of the NBER recession, I did not include it as part of the exogenous TFP recession. These adjustments added eight additional recession quarters and subtracted one (q1 1961).

$$(y_t - u_{s,t} - u_3 * d_t) = f(y_{t-1} - u_{s,t-1} - u_3 * d_{t-1}) + e_t$$

That is, this quarter's deviation from trend growth will depend on last quarter's deviation from trend growth, after adjusting each trend growth rate downward if a recession is present.[16]

Finally, the state s_t is assumed to follow a Markov chain with transition probabilities given by the following matrix:

$$\begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$$

Hamilton (1989) and Hamilton (1994, chapt. 22) show how to use maximum likelihood techniques to estimate the parameters of the system, and also show the filter equations that generate the conditional mean estimate of the state at each point in time. In this example, TFP growth is defined as earlier in this paper and the sample period is 1959:Q2 through 1999:Q4. Estimated parameter values are as follows (t-statistics in parentheses):

$$(y_t - u_{s,t} + \underset{(9.3)}{1.4} * d_t) = \underset{(1.8)}{-.147} (y_{t-1} - u_{s,t-1} + 1.4 * d_{t-1}) + e_t \quad , \text{ where } \underset{(8.9)}{var_e = .59}$$

Trend TFP growth is:

2.8 percent per annum during the fast-growth period (outside of recession)

0.9 percent per annum during the slow-growth period (outside of recession)

5.4 percentage points (annual rate) less than the above figures during recessions

The estimated matrix of transition probabilities is

$$\begin{bmatrix} .992 & .010 \\ .008 & .990 \end{bmatrix}$$

Notice that the coefficient of the autoregressive term in the observation equation is very small;

[16]Of course, this equation can't be estimated directly, because we cannot observe the state of the economy s and thus we do not know the true value of the trend growth rate u at any given point in time. However, given initial estimates of parameter values we can calculate the error term e for each of the four compound states (that is, combinations of current and lag-one states). Next, using the filtering procedure described in Hamilton (1994, pp. 690-696), we can estimate quarter after quarter the probability of being in either of the two regimes at successive times t and t-1. Conditional on these four estimated probabilities, we can construct the likelihood of observing the actual vector of output growth rates y. The standard optimizing routine BFGS is then used to choose parameter values that maximize that likelihood.

deviations from trend (allowing for recessions) are almost white noise. Also, the transition probabilities imply that, given exogenous recessions, a switch from slow to fast trend growth (or vice versa) occurs about every twenty-five years, on average. (Of course if one allowed more states for TFP growth, this figure might well drop.)

Running the data through a Markov filter and smoother gives the probability of being in the fast trend state or the slow trend state in any given quarter. Figure 7 shows the conditional mean estimate of trend growth after downward adjustments for recession periods. The mean estimate equals the probability of being in a fast expansion multiplied by the unconditional mean growth rate for the fast growth state, plus the probability of being in a slow expansion multiplied by the unconditional mean growth rate for the slow state, after downward adjustments for recession periods. The smoothed probabilities suggest that TFP was in the fast trend growth regime with at least 90 percent probability from the beginning of the sample in 1959 until just beyond the end of the recession in late 1982. **The two-stage model of the cycle leads to the interpretation that the slow average TFP growth from late 1973 through 1983 was a result of the three deep recessions during that period, not a slowdown of trend TFP growth.** This implausible result occurs because the two-stage model assumes that rapid TFP growth immediately following recessions is just part of the overall rapid trend growth during expansions that had been seen before the early 1970s. The conditional mean of the smoothed Markov trend declines steadily from the time of the 1973-75 recession, but it is not until 1984:Q1 that the model estimates that the economy was in a slow trend-TFP-growth regime.

Even the simple model with a two-stage business cycle does have one important positive property. The smoothed probability of a new slow-growth regime exceeds 50 percent in 1984:Q1. In contrast, the unsmoothed "real-time" probability of a new slower-growth regime did not exceed 50 percent until 1986:Q3. That is, **ex post it appears that the switch to a slower growth regime occurred two and one-half years before it would have been discovered in real time by the switching model cum filter. This might seem to be a fairly long lag, but it is actually a good deal better than the results of sequential Andrews tests for the trend break.**

With a three-stage model of the business cycle

Like the two-stage model, this model assumes that there are two states of TFP growth – fast trend growth and slow trend growth – and that trend growth is slowed by recessions.

But this three-stage "plucking" model further assumes that growth is typically more rapid than trend expansion growth during periods immediately following recessions. More precisely, a new dummy variable, $d2_t$, is added to the model. The new "rapid recovery" dummy takes a value of 1 in the period following the trough quarter through the quarter in which TFP exceeds its pre-recession peak.[17] With t-statistics in parentheses, the observation equation becomes:

$$(y_t - u_{s,t} + \underset{(6.9)}{1.1} * d_t - \underset{(4.1)}{0.9} * d2_t) = \underset{(2.4)}{-.19}\,(y_{t-1} - u_{s,t-1} + 1.1 * d_{t-1} - 0.9 * d2_{t-1}) + e_t \text{, where } var_e = \underset{(8.8)}{.54}$$

Figure 8 shows the conditional mean estimate of trend growth, after adjusting for recessions and rapid recovery periods.

Estimation results for the plucking model can be summarized as follows. Trend TFP growth is:

2.3 percent per annum during the fast-growth period (outside of recessions or recoveries);

1.0 percent per annum during the slow-growth period (outside of recessions or recoveries);

4.4 percentage points (annual rate) less than the above figures during recession; and

3.4 percentage points (annual rate) more than the above figures during recoveries.

The estimated matrix of transition probabilities is

$$\begin{bmatrix} .987 & .008 \\ .013 & .992 \end{bmatrix}$$

[17] The recession ending in the third quarter of 1980 was followed by only two quarters of rapid recovery before a new recession. In those two quarters, activity did not return to its previous peak. Thus, following the end of the next recession in Q4 of 1982, recovery continued rapidly well after activity had recovered to the peak of mid-1981 – since activity was still well below the peak leading into the 1980 recession. Thus, I set the rapid recovery dummy to one for three additional quarters (1983:Q4 through 1984:Q2) where growth remained well above trend, yet activity remained below the peak which preceded the 1980 recession.

Following the 1990/91 recession, growth was not especially fast in the unusually long period before activity returned to its pre-recession peak. Thus the rapid recovery dummy was not set to 1 for this period.

Figure 9 shows the smoothed conditional probabilities of being in a given state in a given quarter. The curve suggests that TFP was in the fast trend growth regime with at least 90 percent probability from mid-1961 (two years after the beginning of the sample) until the first quarter of 1973. By the second quarter of 1974, the model estimates that we were (more likely than not) in a slow trend-TFP-growth regime. The latter regime persists, with minor shifting in estimated conditional mean, through the end of the sample period in 1999.

Figure 8b compares estimates of trend by the Markov "plucking" model and the deterministic split trend model. Although trend growth during **expansions** is noticeably different for the two models, the **mean** trend growth rate (**including trend downturns** in the Markov model) is virtually identical. The deterministic split trend model gives trend growth rates of 1.7 percent through 1973:Q3 and 0.4 percent thereafter. The Markov model's average trend growth rates over the same periods are 1.7 percent and 0.5 percent.

The plucking model, like the two-stage model, catches regime switches slowly but nonetheless faster than alternative methods. With the plucking model, the smoothed probability of a new slow-growth regime exceeds 50 percent in 1974:Q2. In contrast, the unsmoothed "real-time" probability of a new slower-growth regime did not exceed 50 percent until 1977. That is, ex post it appears that the switch to a slower growth regime occurred 3 ½ years before it would have been discovered in real time by the switching model *cum* filter. This is a year or two better than the unpublished results that Norman Morin (of the Federal Reserve Board) found in 1999 for his Andrews tests, following a slightly larger shock to trend productivity growth in 1973.

Was there an increase in trend TFP growth in mid-1995, as some have asserted? The Markov "plucking" model suggests at best a small rise. From the beginning of 1995 through the end of 1999, the model's conditional mean estimate rises only 0.1 percentage point (from 1.01 percent to 1.11 percent) The reason for so little change is that the Markov model has a much smaller slowdown in trend TFP expansion in the decade or two before 1995 than many analysts have. A large portion of the weakness in TFP growth from 1979 through 1995 is attributed by the model to cyclical considerations. In a sense, other estimates of trend TFP growth have risen since 1995 to match the Markov model's estimate.

VI. Pagan's moment tests of the Markov switching model

Adrian Pagan (2001) points out that the Markov switching model has been misused

in many published studies. In particular, model-generated estimates of the observed variable often had grossly different population mean and variance than the mean and variance of the data actually observed – suggesting that the observed data were most unlikely to be consistent with the Markov model. Fortunately, in the case of trend TFP growth the Markov model cannot be rejected based on these moment tests, as discussed in Appendix 1.

VII. Conclusion

Although trend TFP growth slowed in 1973-74, there is no clear evidence of continuously evolving trend growth in total factor productivity – a fact that calls into question the usual stochastic trend models or the use of Hodrick-Prescott filtering. Of all the methods tried, the "plucking" version of the Markov switching model seems most promising for both modeling the historical path of TFP growth and checking for more recent shifts in trend. That model comfortably captures the split in trend in 1973-74 and indicates a possible slight upward shift in trend growth for more recent years to about 1.1 percent.

Is it possible to detect shifts in trend growth of TFP more quickly? Perhaps by bringing additional information to bear from other supply-side variables, such as workweek and participation rates. But these other supply-side variables have normal innovations to trend growth whereas TFP has non-normal innovations. It turns out, however, that one can still jointly model the components of aggregate supply. Lam (1990) has presented a procedure for estimating a switching model for one variable jointly with more traditional unobserved components models for other variables in a nonlinear multivariate state-space framework. It would be interesting to see what such a joint model of the supply side would say about the possibility of an uptick in trend TFP growth around 1995.

Appendix 1: Pagan's moment test for assessing Markov switching models

Pagan's test is as follows. Start by conducting monte carlo simulations drawing from the random distributions in the model, taking all the estimated parameters as given. For the model in this paper, there are two such distributions. First, the error term in the observation equation at the top of page 16 is normally distributed. The second distribution is associated with the four transition probabilities p_{ij} described on page 16. Let ε_t be drawn from a uniform distribution between 0 and 1. If TFP growth was in state j at time t-1, the state at time t is

$$1 \quad \text{if} \quad \varepsilon_t < p_{1j} \text{ , or}$$

$$2 \quad \text{if} \quad \varepsilon_t \geq p_{1j}$$

For each time t (in succession) draw once from the uniform distribution and once from the normal distribution. Use these two draws to generate a simulated value of the state and of the growth rate of TFP at time t. Repeat the process for time t+1, contingent on the simulated state at time t and simulated growth rate of TFP at time t. Continue until simulated values has been generated for the growth rate of TFP for each time period in the observed sample. Calculated the mean and standard deviation of the simulated growth rates of TFP. Iteration 1 is now finished.

Repeat the steps in the previous paragraph many times (10,000 in this paper). This generates 10,000 simulated values of the mean and standard deviation of TFP growth. Take the average of these 10,000 means and 10,000 standard deviations, to get the population mean and standard deviation of TFP growth, to compare with the observed sample mean and standard deviation of TFP growth. Is the difference between observed mean and population mean for TFP growth (call it τ_{mean}) large or (hopefully) small? Similarly, is the difference between observed standard deviation and population standard deviation for TFP growth (call it τ_{std}) large or small? To make this judgment, Pagan suggests calculating the standard deviation of the 10,000 simulated values of the mean TFP growth, and of the 10,000 simulated values of the standard deviation of TFP growth. Divide the τ for each moment by the corresponding standard deviation of the 10,000 simulated moments. The ratio is approximately t-distributed. If the t-statistic for the mean or variance of MFP growth is large, one can reject the Markov switching model.

For the Markov model in this paper, the t-statistic for mean growth of TFP was 0.1, and the t-statistic for the standard deviation of tfp growth was 1.2. The sample moments are reasonably close to the population moments. In neither case are there sufficient grounds to reject the Markov-switching model.

Appendix 2: Data definitions

This paper uses data from the first quarter of 1959 through the fourth quarter of 1999. Data for almost all variables used in the above analysis were taken from quarterly and annual data published by the Bureau of Labor Statistics (BLS). They do not include new or revised BLS data published from April 2001 onward.

Labor quality

Quality-adjusted hours data (hours of all private nonfarm workers excluding government enterprises) were taken from raw data published by the BLS.[18] The quality-adjusted hours data were then divided by the unadjusted hours data, to get an index number representing "labor quality". The index number was forecasted forward based on the growth of the factor over the last historical year where adjusted BLS data are available (that is, 1997). In the work described above, I have treated the historical data series for labor quality as an exogenous variable.

Capital services

Data on capital services were generated as follows. Annual capital stock data produced by the BLS (available through 1997) were interpolated to a quarterly frequency. Then the quarterly capital stock estimates were extended forward on the basis of quarterly investment data since 1997.

The BLS defines private nonfarm business capital stock very broadly to include equipment, structures, inventories and land. The (Tornqvist) aggregate of these variables differs crucially from the BEA aggregate in that BLS uses **rental** rates of capital to chain-aggregate, while BEA uses **prices** to chain-aggregate. This difference matters when data on nonresidential structures (which have relatively low depreciation rates and thus low rental rates relative to price)

[18]The BLS quality adjustment of labor hours involves Tornqvist-weighting of the hours of different classes of workers, based on their share of total labor compensation. Tornqvist aggregation is similar to Fisher-aggregation in many ways: both have "ideal", "superlative" properties. The Tornqvist index averages a sector's weight in successive periods before taking a weighted geometric average of quantity growth in the various sectors, while the Fisher index aggregates growth across sectors with one period's weights (Laspeyres), then with the next period's weights (Paasche), and then geometrically averages the two aggregates. The Fisher index is exact assuming unitary elasticity of substitution among sectors (Cobb-Douglas), while the Tornqvist index is best for the more general Translog setup. In practice, this difference in methodology is probably not the most important cause of differences between the BEA and BLS aggregates.

are combined with data on equipment (which has high depreciation rates and thus high rental rates relative to price). Because the stock of equipment has grown faster than the stock of nonresidential structures, the BLS capital stock increases faster than the BEA capital stock. Also, the BLS assigns higher weights to computers than to other equipment compared with the BEA aggregate equipment data. This difference in weights also affects the estimated growth of the stock of equipment. The BLS approach is a theoretically preferred "service flow" approach to estimating an aggregate capital input to the production function. The BLS use of rental rates has a practical advantage as well: the resulting faster growth of estimated capital stock reduces the size of the Solow residual in the BLS estimates.

The BLS measure of aggregate capital differs from most in that it includes land. In this study, the real value of the stock of land for recent quarters is based on its average growth rate over the last two or three years of available data through 1997.

Nonfarm business output

Data on output were taken from the BLS. (The BEA measure assigns zero labor productivity growth in government enterprises, which is not an acceptable assumption for studies of productivity.) In contrast to the BEA definition, the BLS output definition excludes output of government enterprises: the resulting measure is called "private nonfarm business output". The BLS releases data for this narrower measure in the productivity and cost release, twice quarterly, one and one-half weeks after the advance and preliminary GDP release.[19]

[19]Second-revision data on nonfarm output are published in the first productivity and cost release for the following quarter -- i.e. with a month delay.

REFERENCES

Andrews, D.K. (1993), "Tests for Parameter Uncertainty and Structural Change with Unknown Change Point", Econometrica, vol. 61, pp. 821-856.

Bai, J. (1997), "Estimating Multiple Breaks One at a Time", Econometric Theory, vol. 13, no. 3, pp. 315-352, June.

Baxter, M. and R. King (1999), "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series", Review of Economics and Statistics, vol. 81, pp. 575-593, November.

Braun, S. (1990), "Estimation of Current-Quarter Gross National Product by Preliminary Labor-Market Data", Journal of Business and Economic Statistics, pp. 293-304, July.

Bureau of Labor Statistics (1997), Handbook of Methods, Chapter 10, "Productivity Measures: Business Sector and Major Subsectors", April.

Christiano, L. and T. Fitzgerald (1999), "The Band Pass Filter", NBER Working Paper 7257, July.

Cogley, T. and J. Nason (1995), "Effects of the Hodrick-Prescott filter on trend and difference stationary time series: Implications for business cycle research", Journal of Economic Dynamics and Control, vol. 19, pp. 253-278.

DeLong, J.B. and L. Summers (1988), "How Does Macroeconomic Policy Affect Output", Brookings Papers on Economic Activity 2, pp. 433-480.

Filardo, A. and P. Cooper (1996), "Cyclically-Adjusted Measures of Structural Trend Breaks: An Application to Productivity Trends in the 1990s", Federal Reserve Bank of Kansas City Research Working Papers RWP 96-14, December.

Gordon, R. (2000), "Does the 'New Economy' Measure up to the Great Inventions of the Past? ", Journal of Economic Perspectives, vol. 14, no. 4, pp. 49-74, Fall.

Hamilton, J. (1986), "A Standard Error for the Estimated State Vector of a State-Space Model", Journal of Econometrics, vol. 33, pp. 387-397.

Hamilton, J. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", Econometrica, vol. 57, pp. 357-384, March.

Hamilton, J. (1994), Time Series Analysis, Princeton University Press.

Harvey, A. and S. Koopman (1992), "Diagnostic Checking of Unobserved-Components Time Series Models", Journal of Business and Economic Statistics, pp. 377-389, October.

Hodrick, R. and E.C. Prescott (1980), "Post-war U.S. business cycles: An empirical investigation", Discussion Paper at Northwestern University and Carnegie-Mellon University.

King, R. and S. Rebello (1993), "Low frequency filtering and real business cycles", Journal of Economic Dynamics and Control, vol. 17, pp. 207-231.

Kydland, F. and E. Prescott (1990), "Business Cycles: Real Facts and a Monetary Myth", Federal Reserve Bank of Minneapolis Quarterly Review, vol. 14, no. 2, pp. 3-28, Spring.

Lam, P. (1990), "The Hamilton Model with a Generalized Autoregressive Component: Estimation and Comparison with Other Models of Economic Time Series", Journal of Monetary Economics, vol. 26, pp. 409-432.

Laxton, D. and R. Tetlow (1992), "A Simple Multivariate Filter for the Measurement of Potential Output", Bank of Canada Technical Report No. 59, June.

Oliner, S. and D. Sichel (2000), "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?", Journal of Economic Perspectives, vol. 14, no. 4, pp. 3-22, Fall.

Pagan, A. and R. Breunig (2001), "Some Simple Methods for Assessing Markov Switching Models", working paper, Economics Program, RSSS, Australian National University, April 11.

Perron, P. (1989), "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis", Econometrica, vol. 57, no. 6, pp. 1361-1401, November.

Perron, P. (1997), "Further Evidence on Breaking Trend Functions in Macroeconomic Variables, Journal of Econometrics, vol. 80, no. 2, pp. 355-385, October.

Stock, J. and M. Watson (1998), "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model", Journal of the American Statistical Association, pp. 349-358.

Figure 1a

Actual vs. H-P Filtered (3987) Log of Total Factor Productivity

Figure 1b

Actual vs. H-P Filtered (1600) Log of Total Factor Productivity

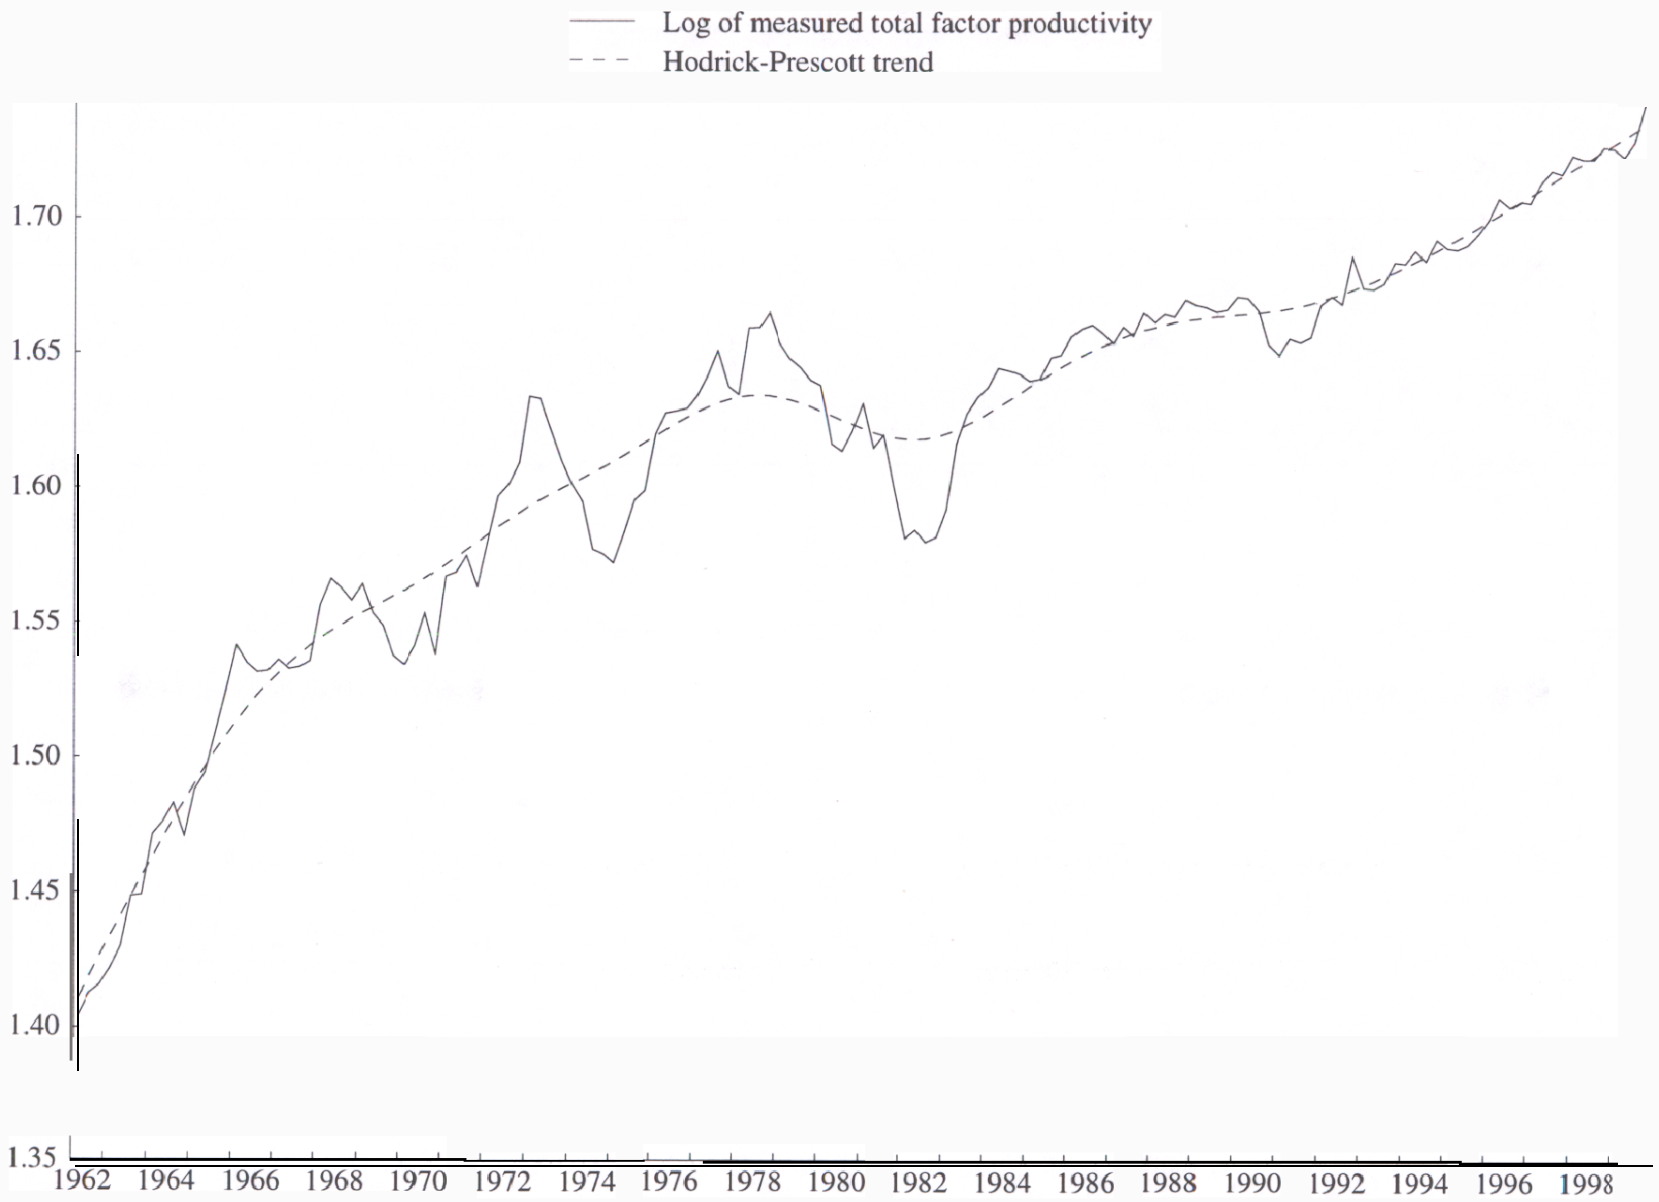

Figure 1c

Actual vs. H-P Filtered (6400) Log of Total Factor Productivity

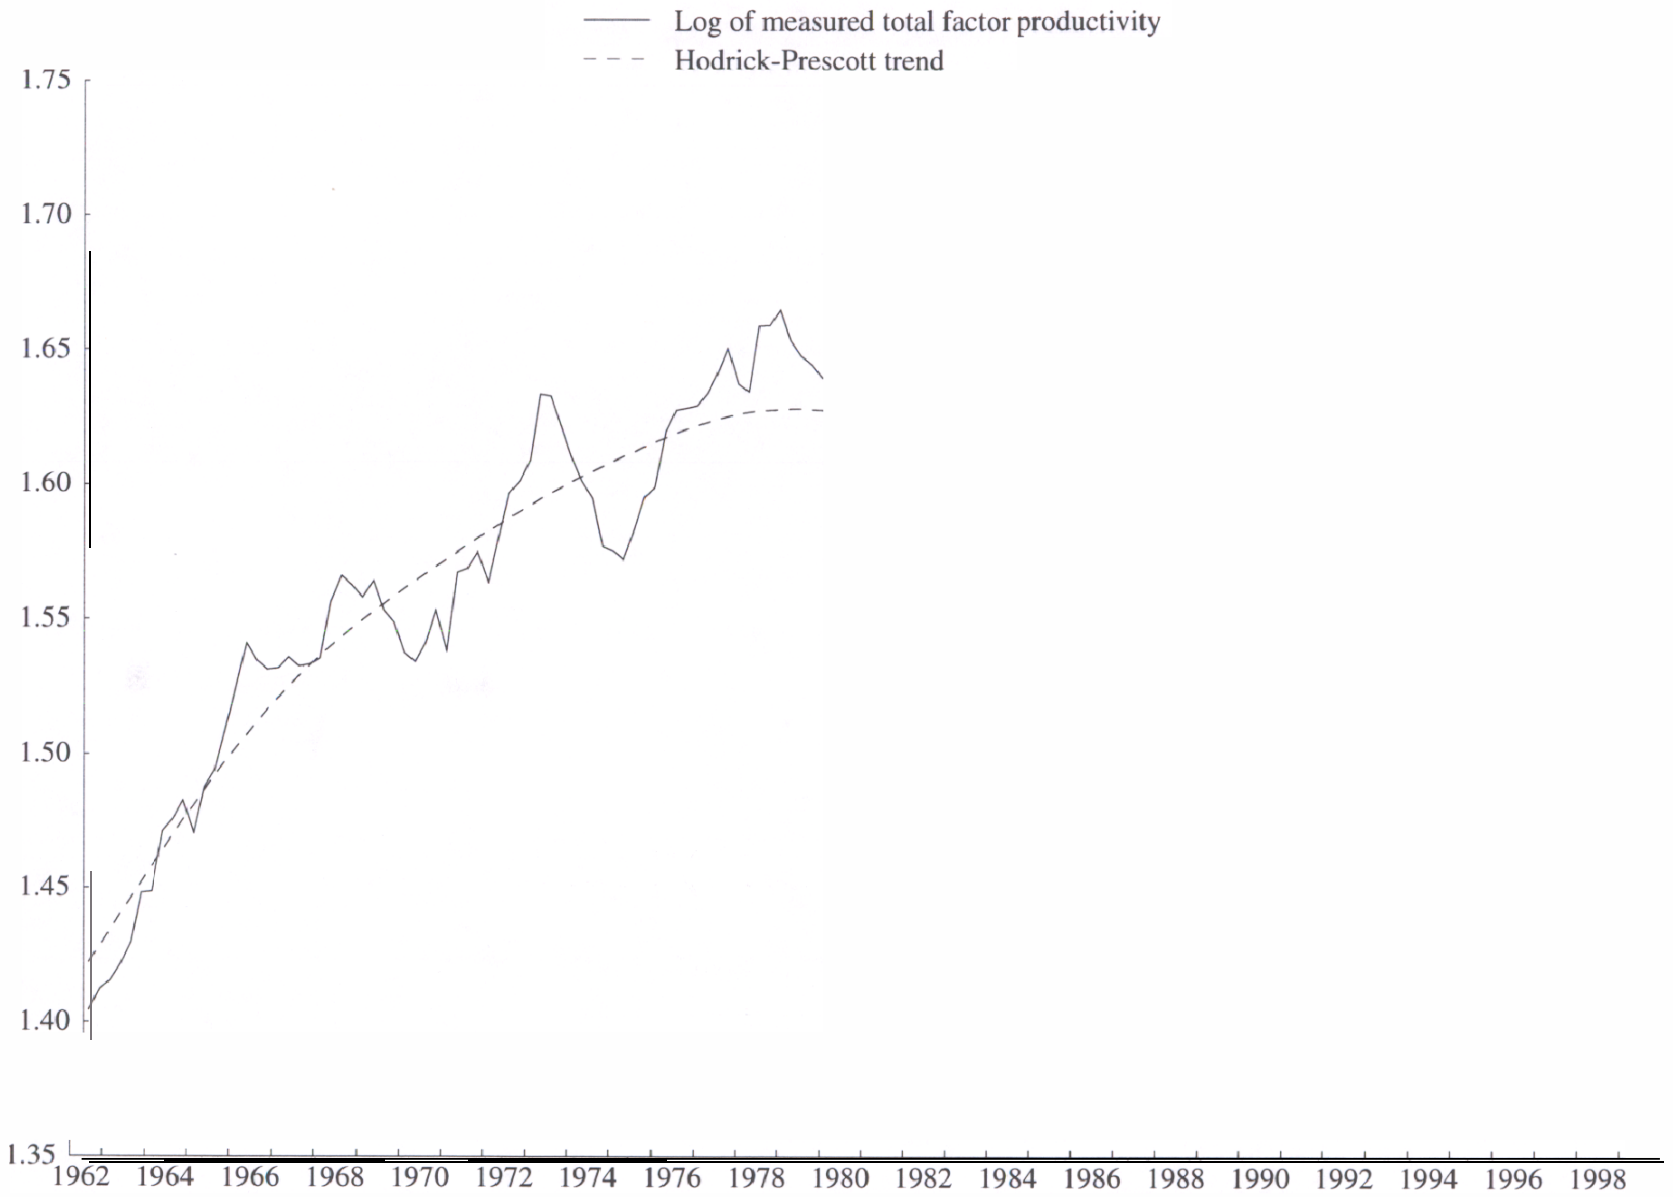

Figure 2

Actual Log TFP vs. Bandpass Trend

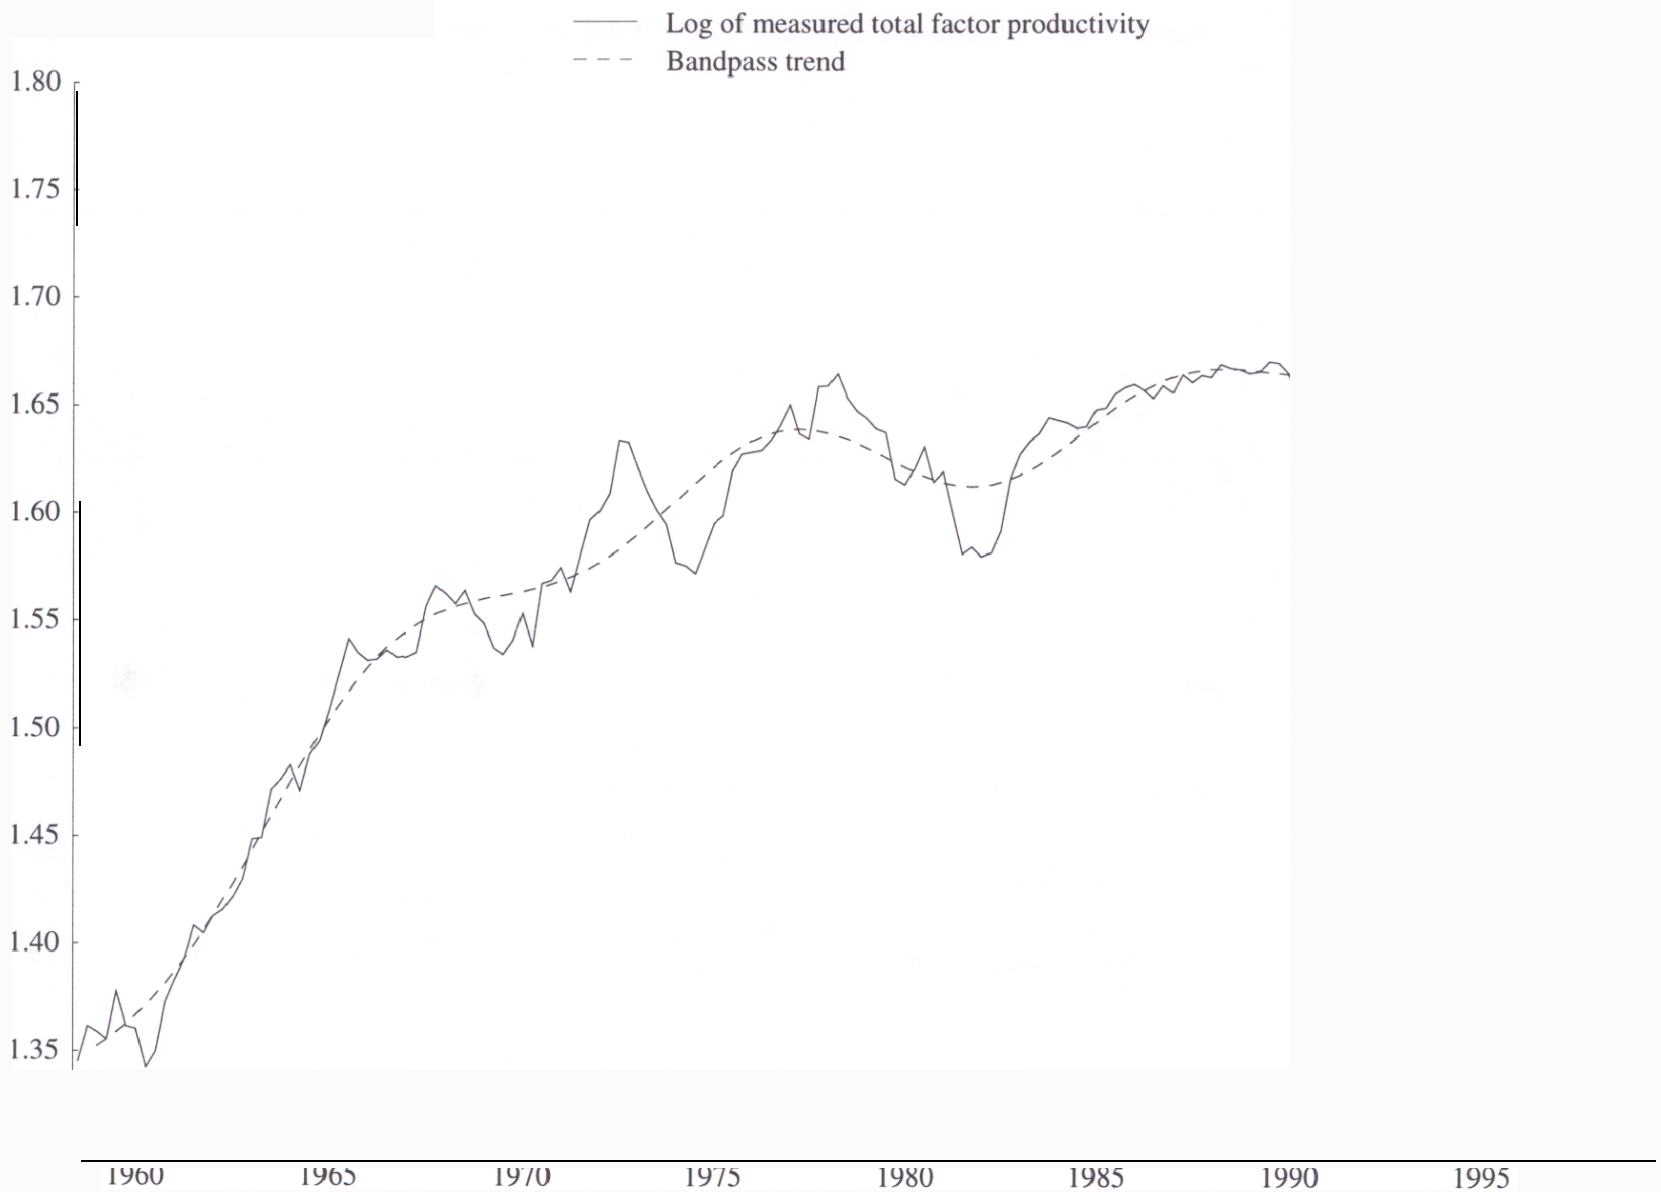

Figure 3

Two Measures of Log Trend TFP

Bandpass vs. H-P Filter (1600)

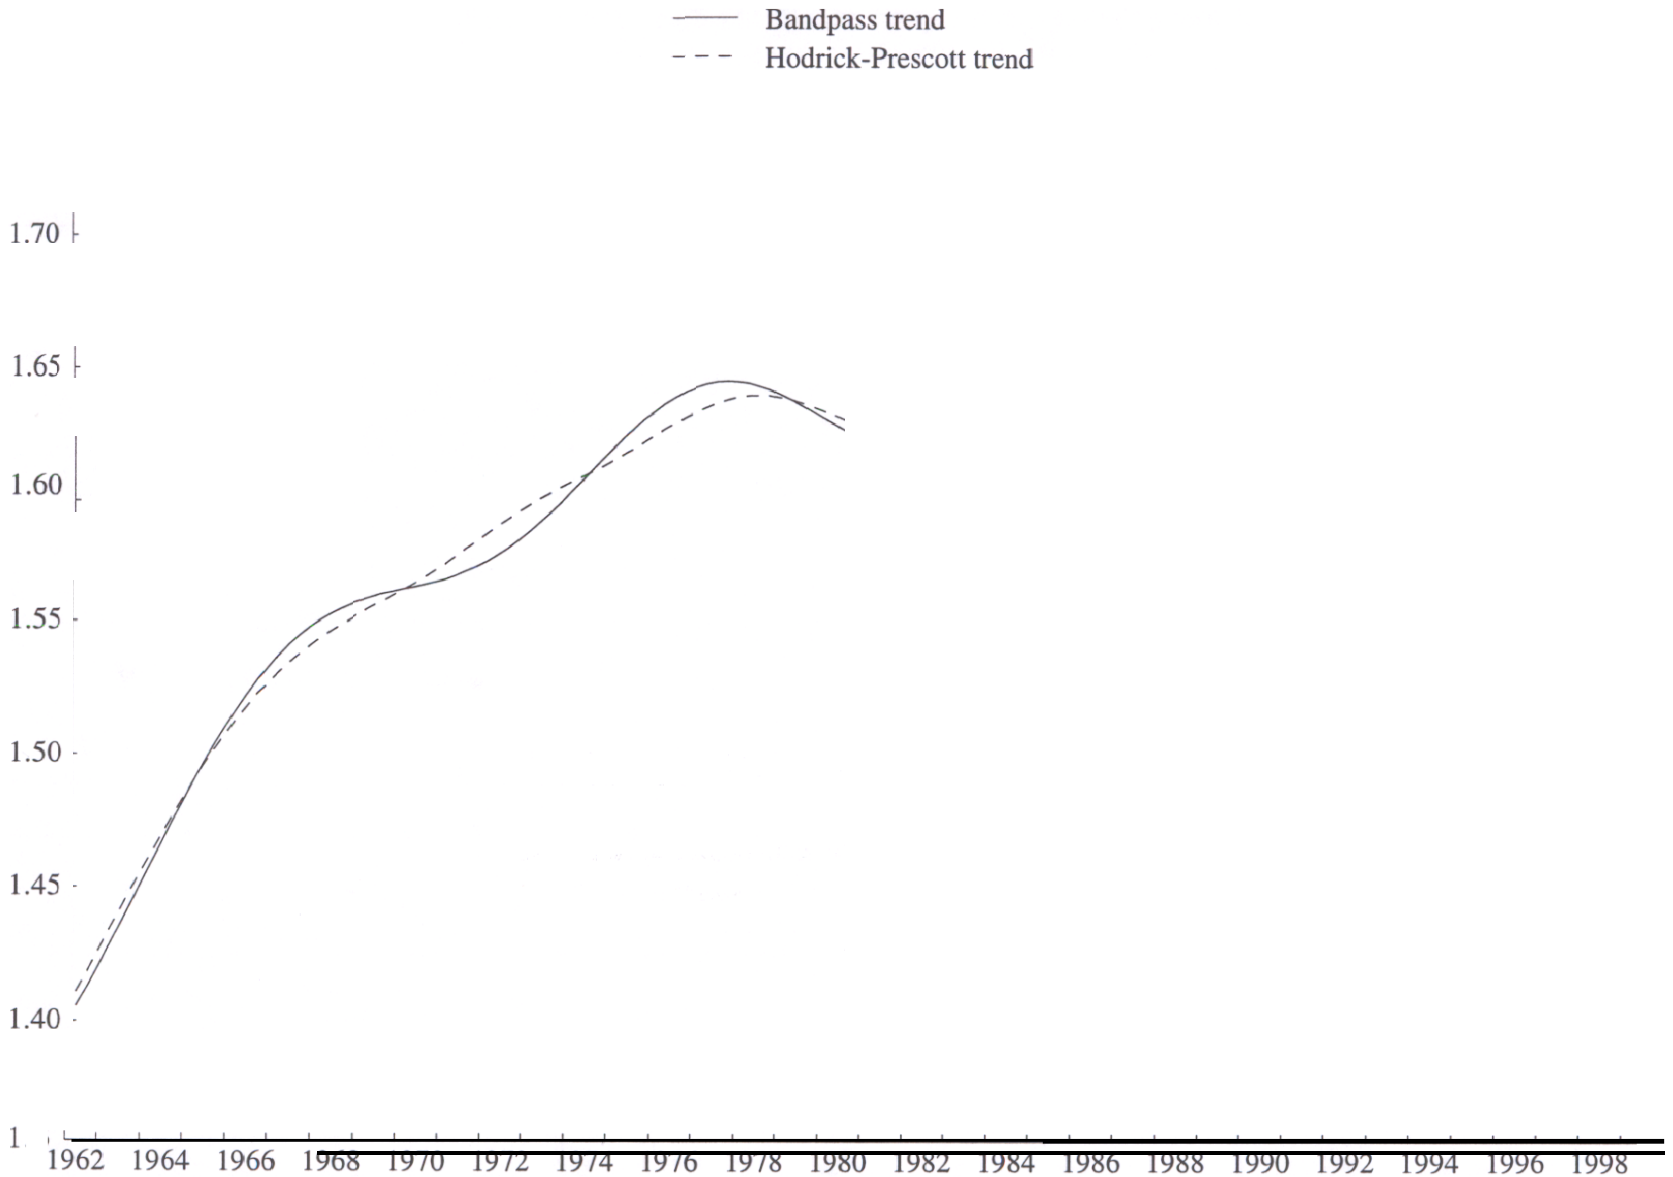

Figure 4

Actual Log TFP vs. Kalman Trend

(Separate restrictions imposed on I(1) and I(2) variances)

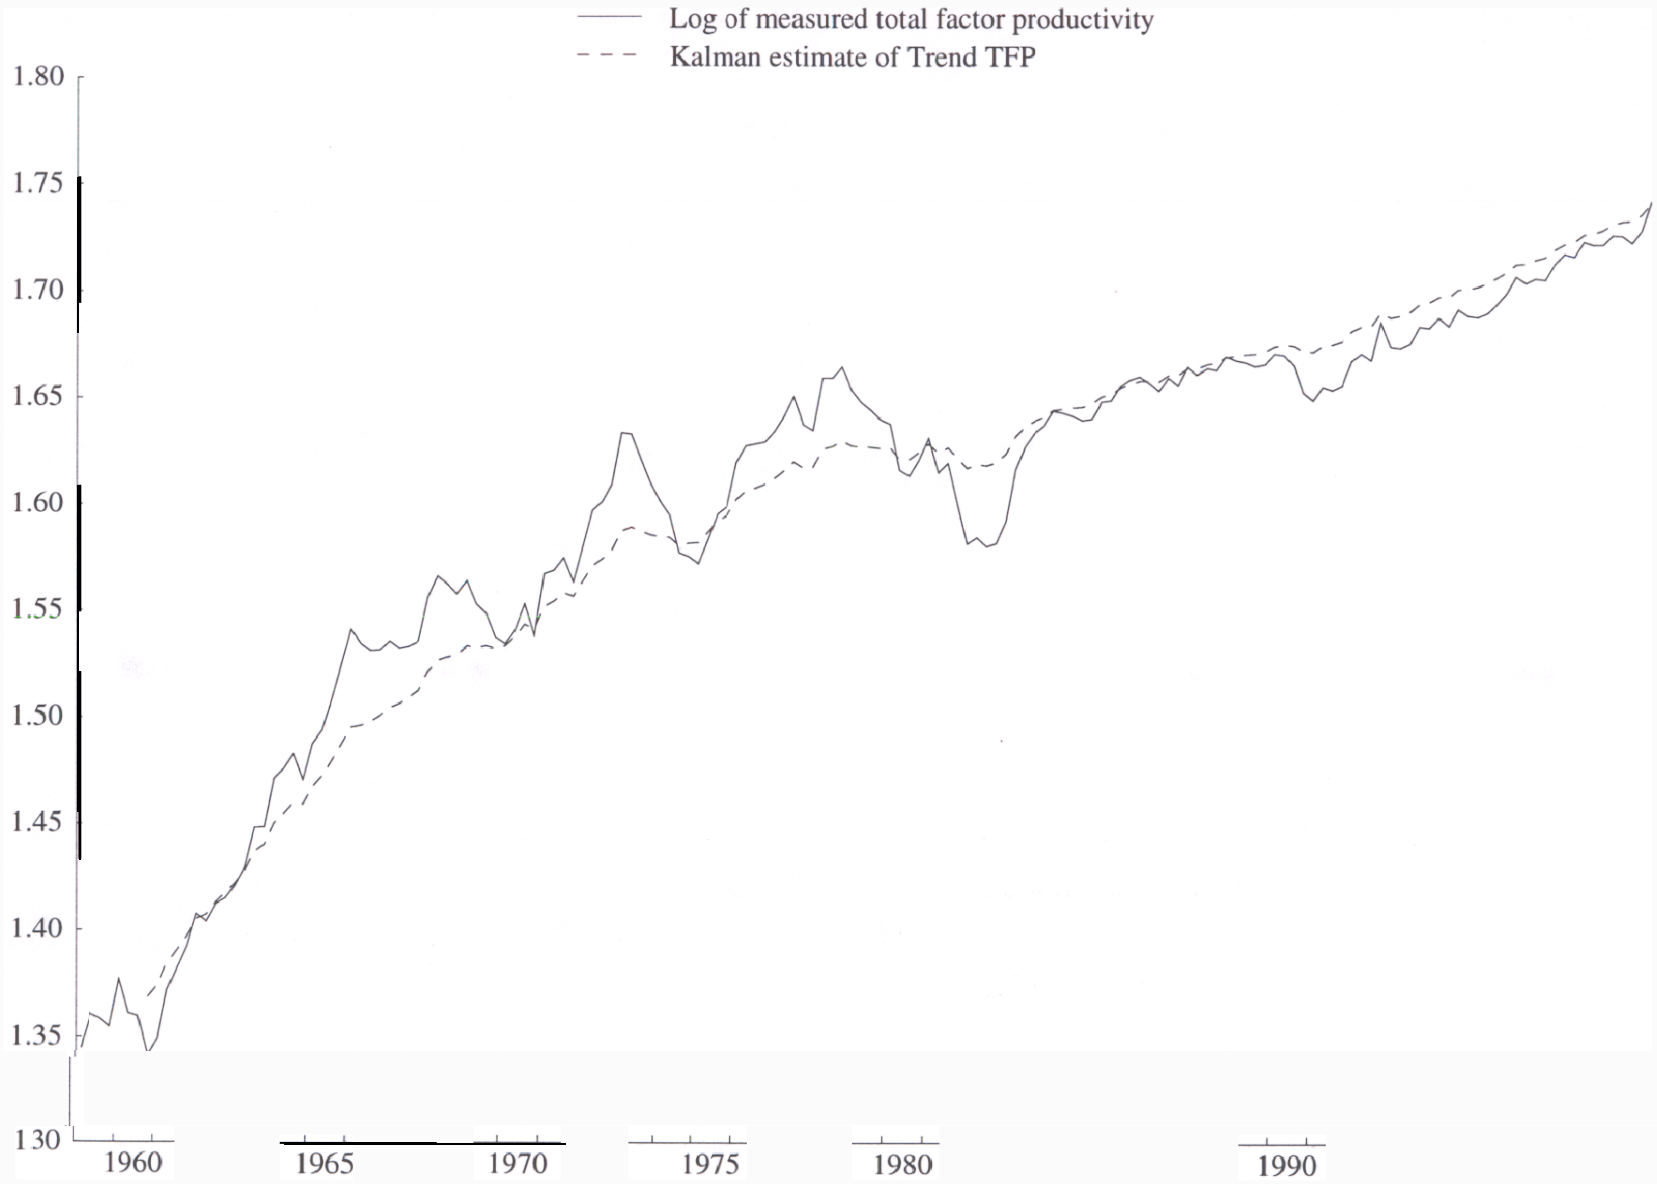

Figure 5

Actual Log TFP vs. Estimated Trend from Kalman Filter with 1973 & 1995 Dummies

Figure 6a

Two Measures of Log Trend TFP

Kalman with 1973 & 1995 Dummies vs. H-P Filter (1600)

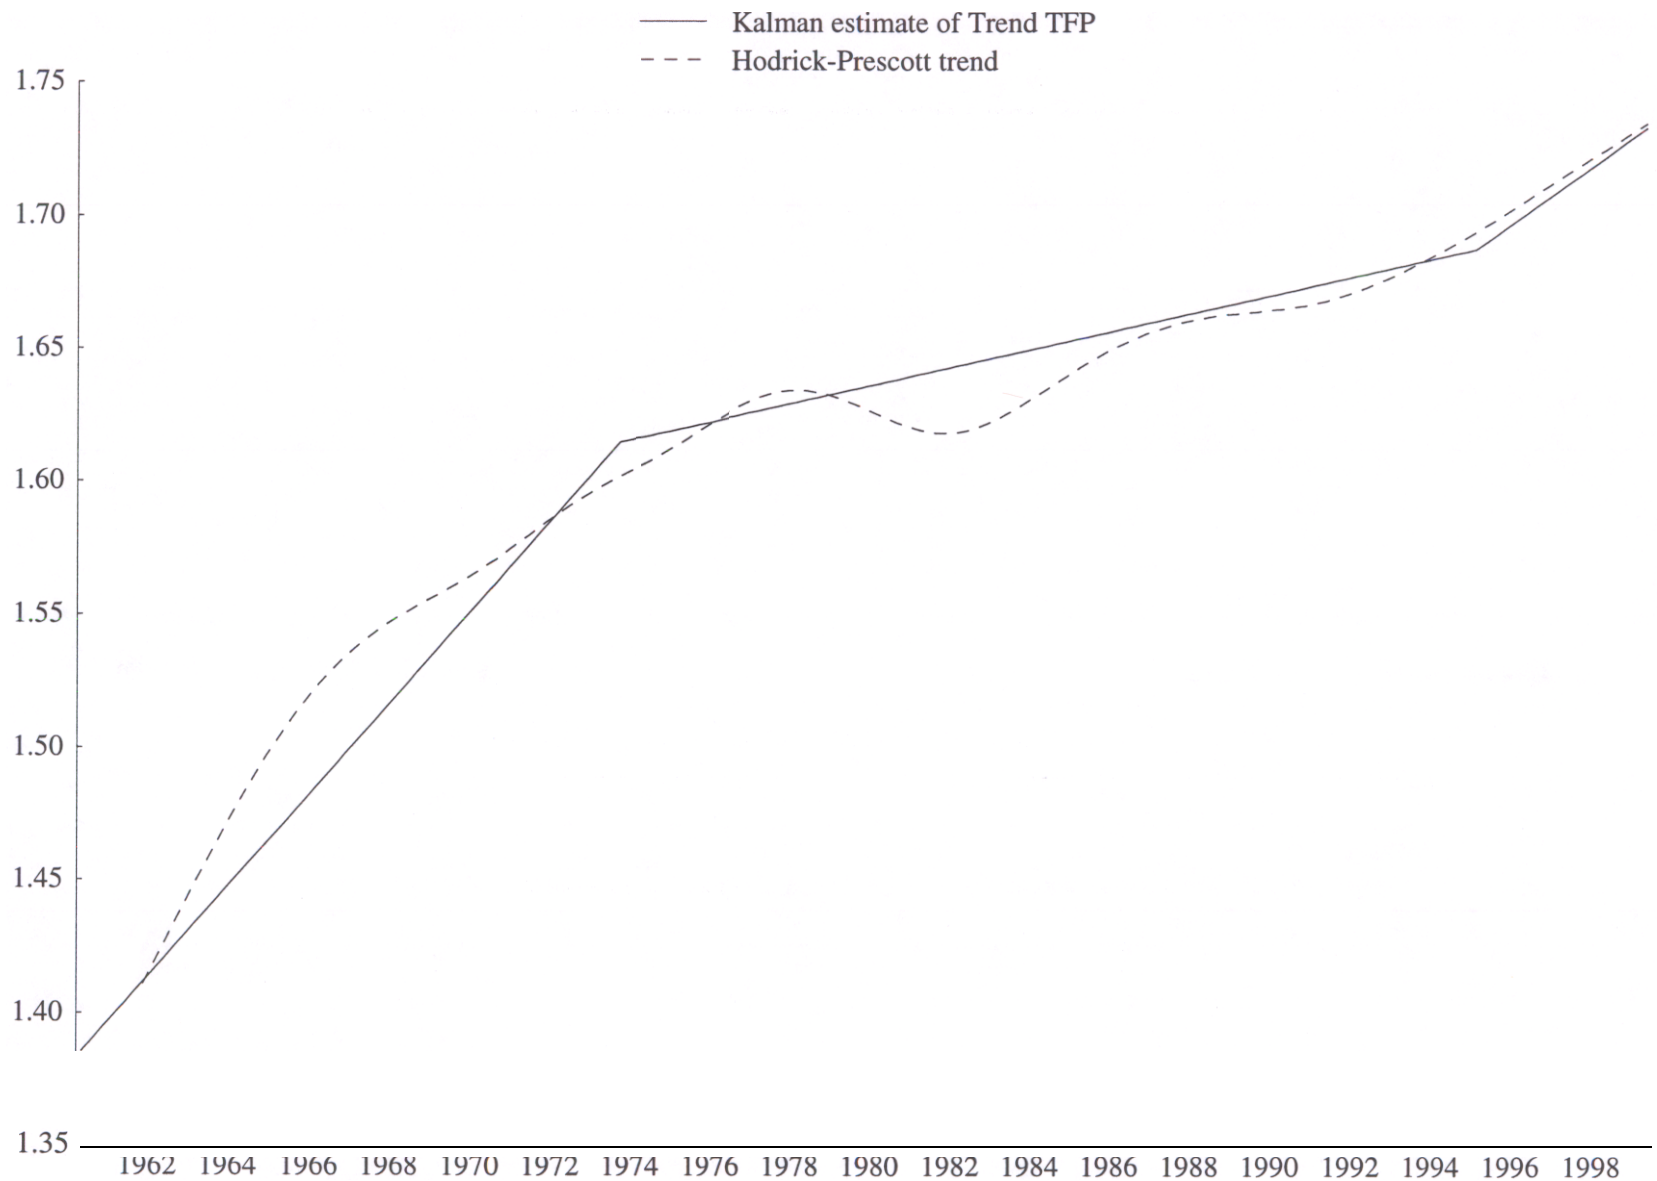

re (

o Measu f g Tre d TFP

alman with 1973 um vs. H-P Fi r (6400)

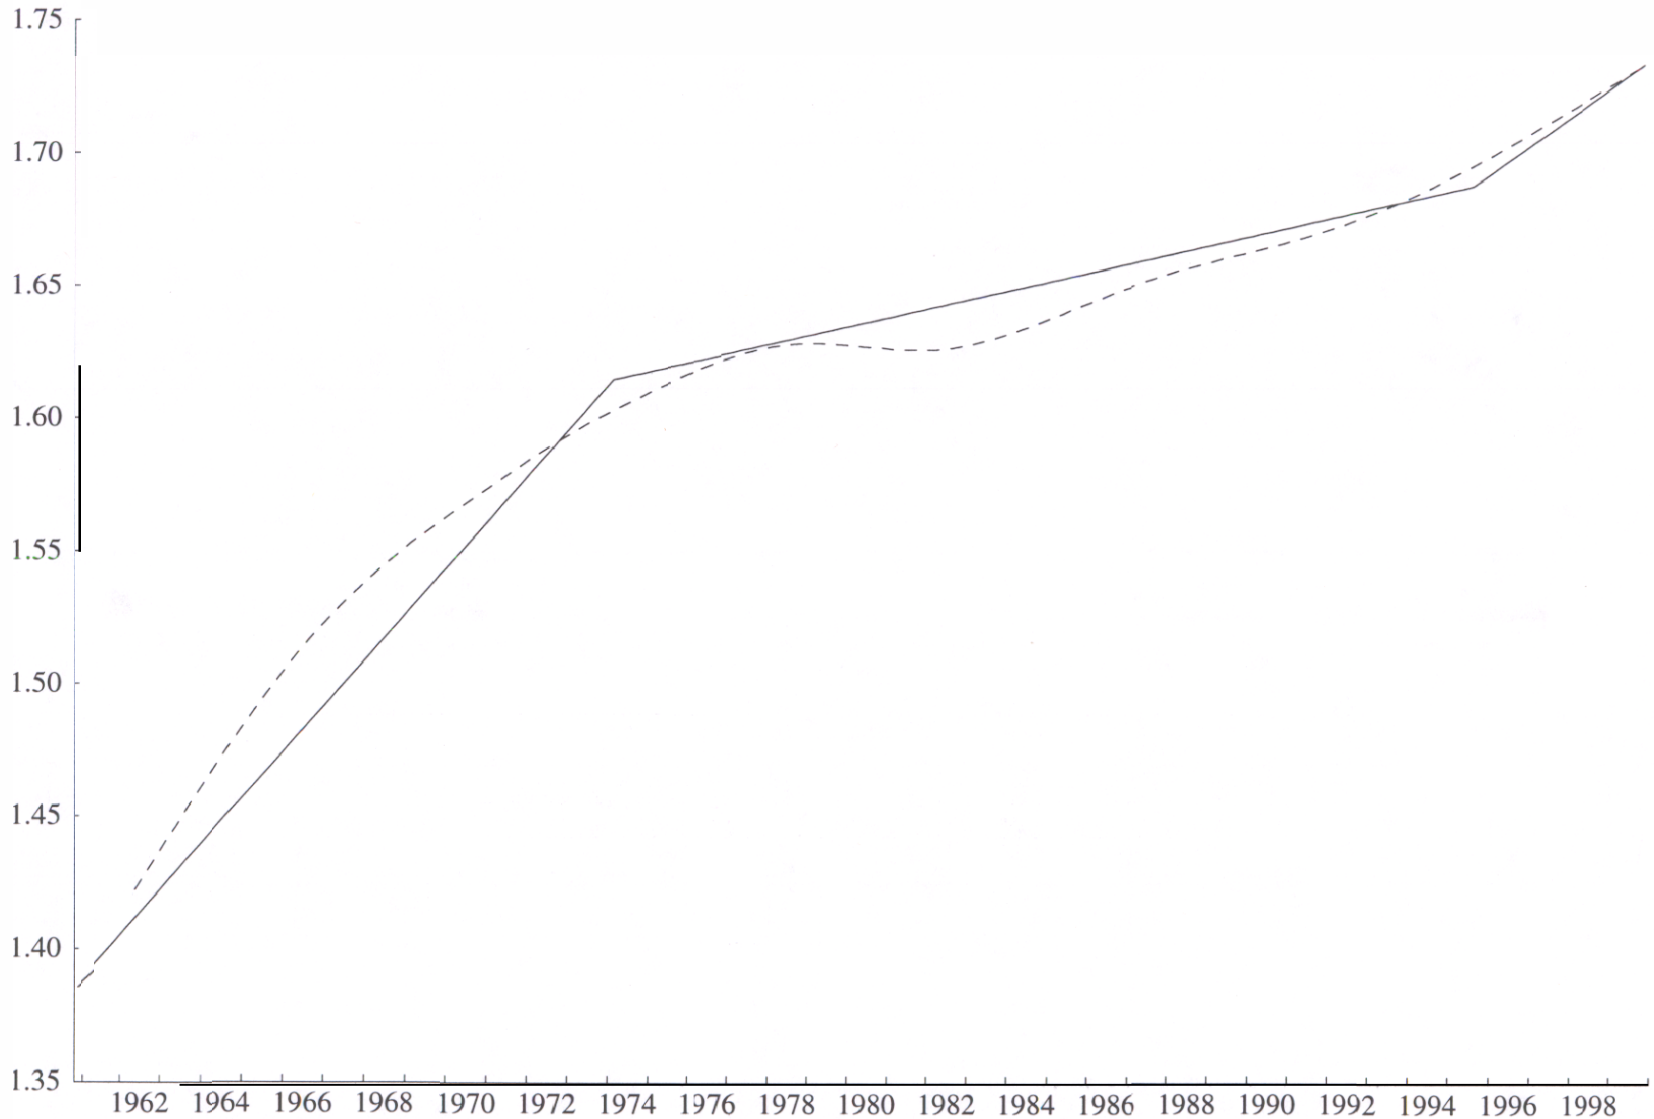

Figure 7

Total Factor Productivity vs. Markov-switching trend

(with two-stage model of the business cycle, no rapid recovery)

Figure 8

Total Factor Productivity vs. Markov-switching trend

(with 'plucking' model of the business cycle)

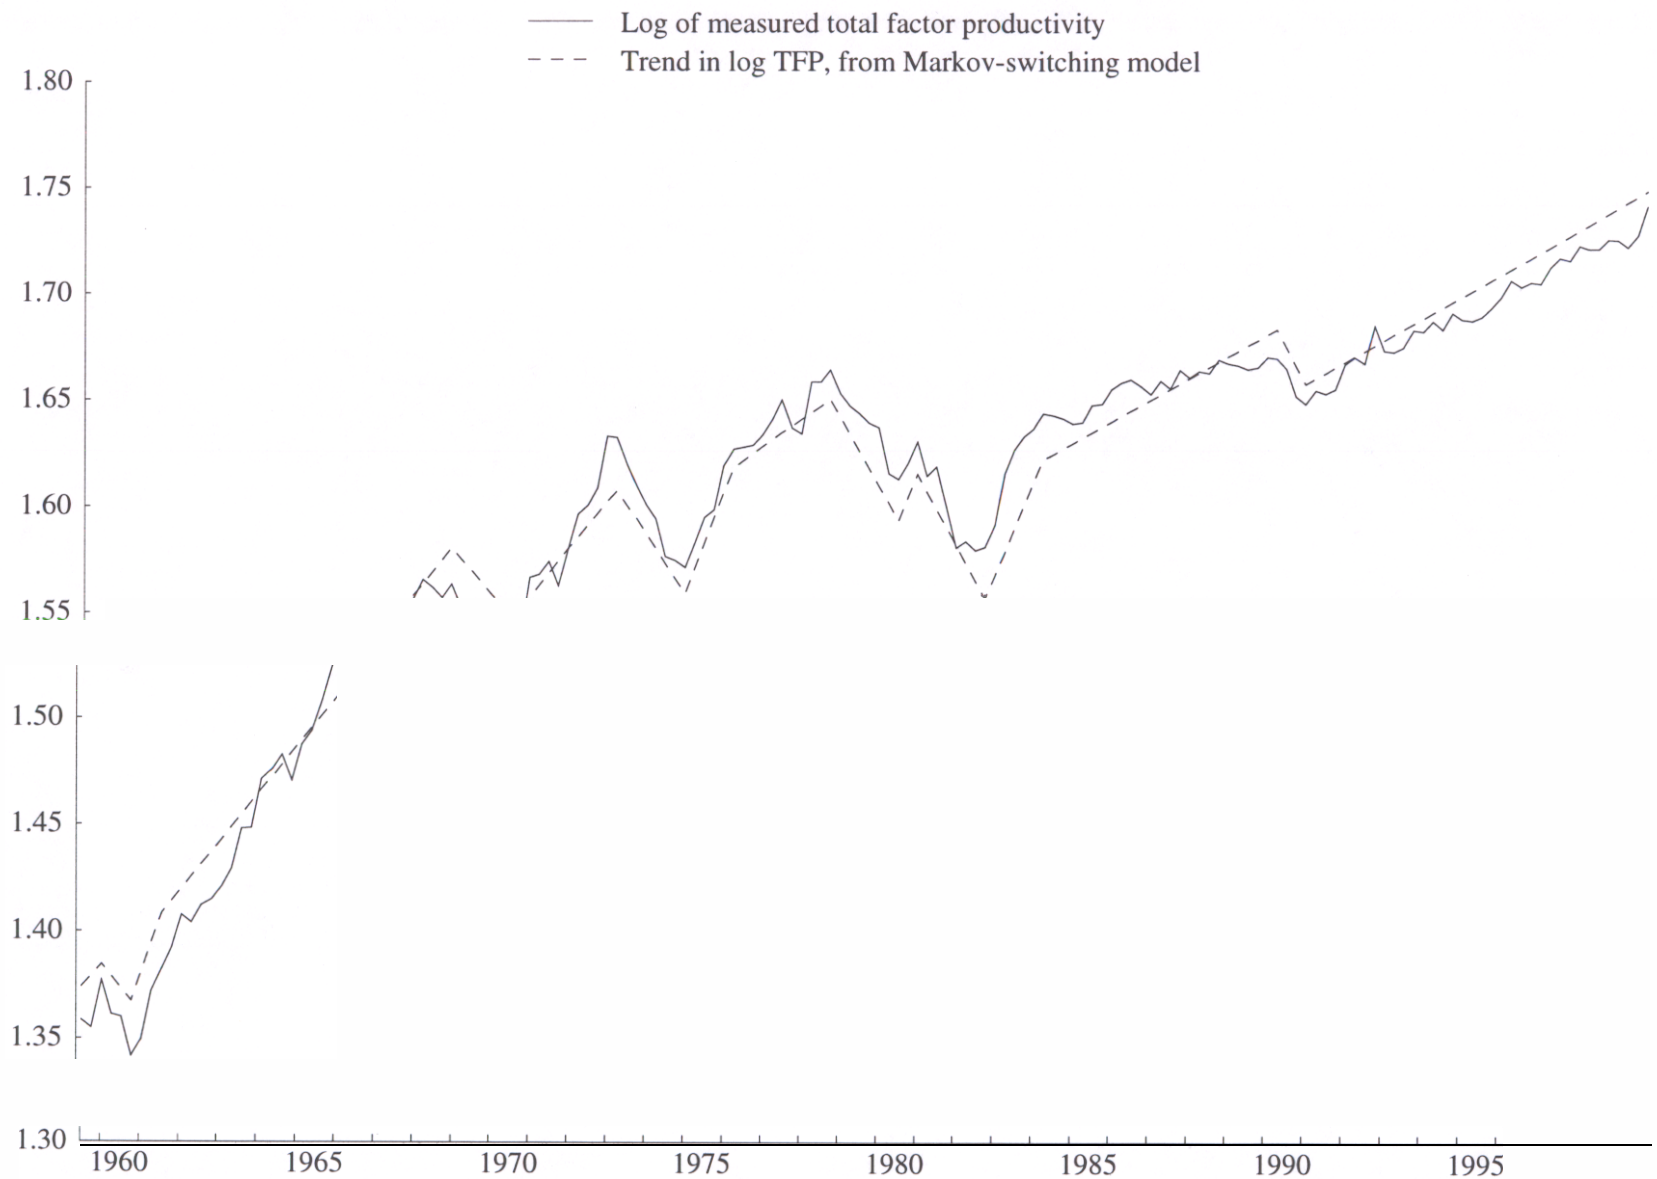

Figure 8b

Deterministic Trend (split in 1973 & 1995) vs. Markov-switching trend

(with 'plucking' model of the business cycle)

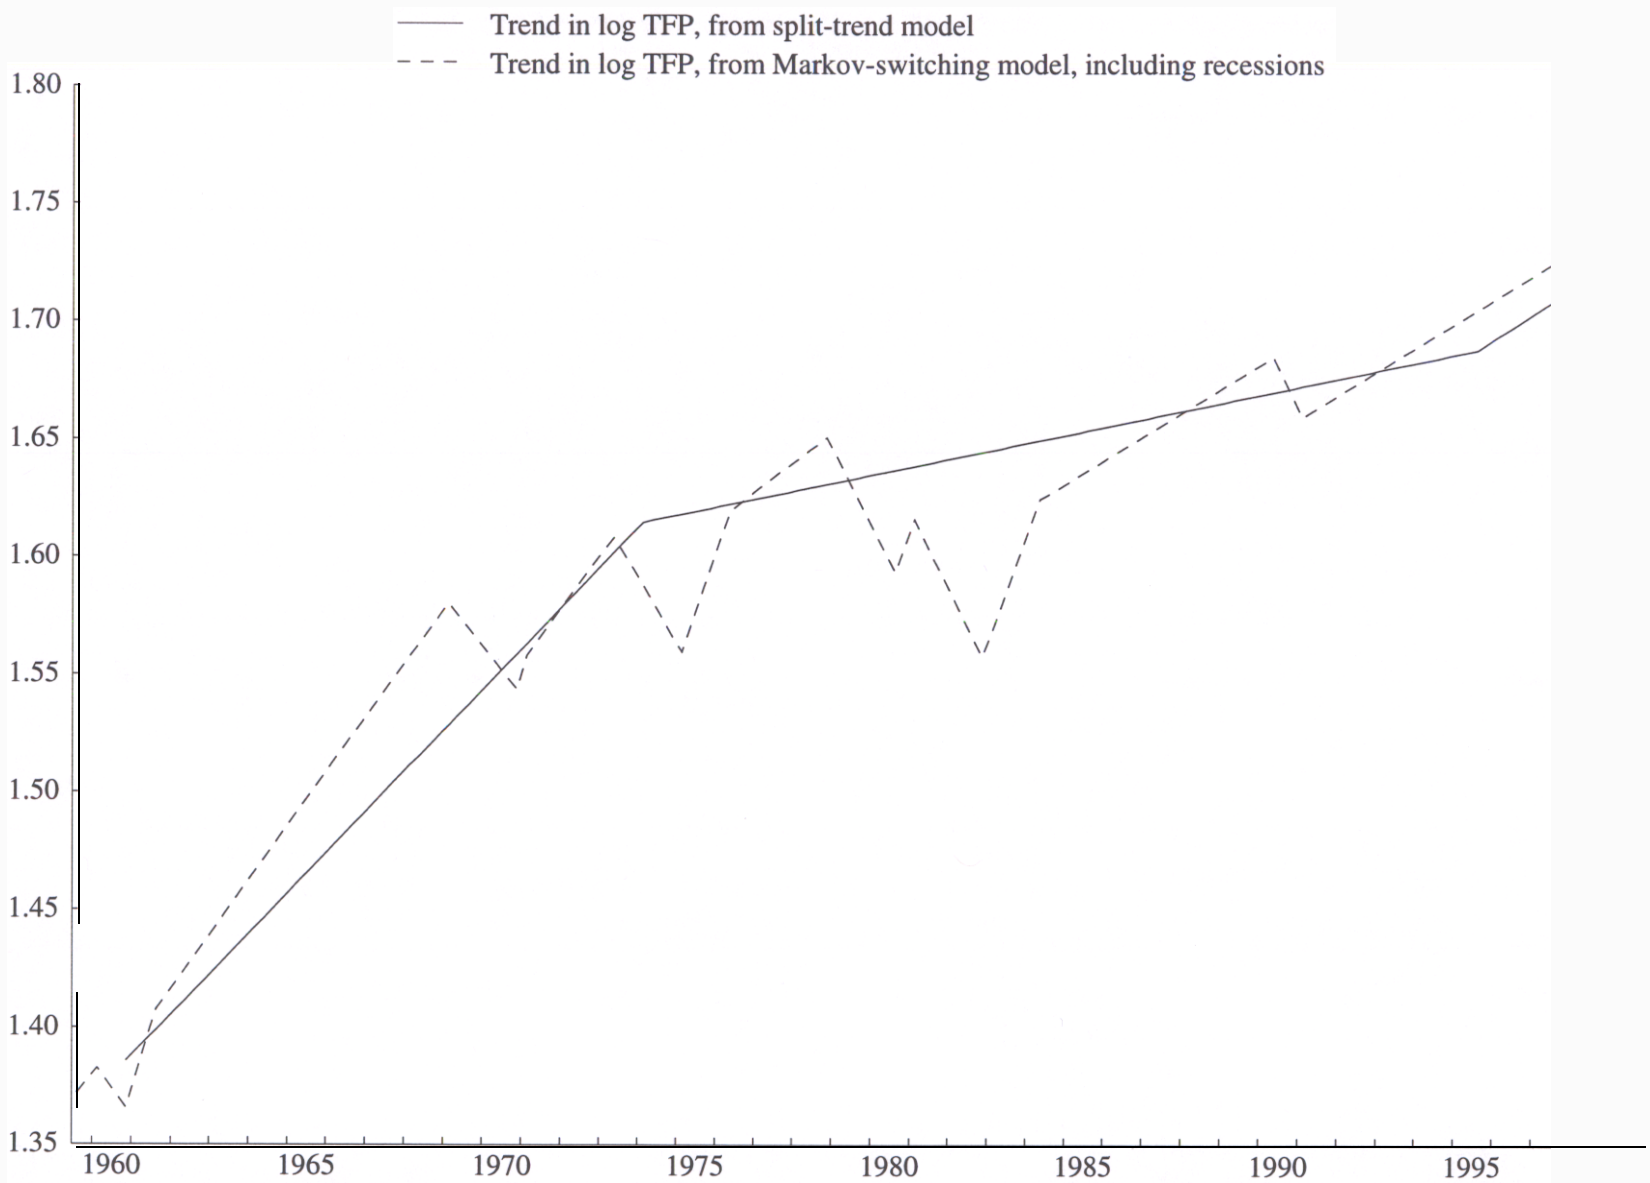

Figure 9

Markov Switching Model

(with 'plucking' model of the business cycle)

Conditional probability of being in fast trend growth regime (2.3%) vs. slow trend growth regime (1.0%)