Matthew D. Shapiro and Mark W. Watson
YALE UNIVERSITY/NBER, NORTHWESTERN UNIVERSITY/NBER

Sources of Business Cycle Fluctuations

1. *Introduction*

What is the source of business cycle fluctuations? Most theories take the answer to this question to be axiomatic. The essence of Keynesian theories is that in the short-run the willingness of agents to absorb the output of the economy determines the quantity of output produced. On the other hand, classical and new classical theories do not allow the possibility that output can deviate from capacity, except for very short intervals. In these theories, prices and rates of return adjust so that a change in aggregate demand does not cause output to change. Here, we attempt to quantify the sources of economic fluctuations by making minimal and plausible identifying restrictions that do not depend on a theory of the business cycle.

Standard textbook treatments of macroeconomic fluctuations separate the high frequency, business cycle fluctuations from the low frequency, growth fluctuations. This dichotomy lies at the heart of most Keynesian and rational expectations models.[1] In these models, shocks to aggregate demand temporarily move the economy away from some "full-employment" "potential," or "natural" level of output. The natural level of output is determined by the capital stock, the labor force, and technology in long-run equilibrium.[2] These supply-side factors are assumed to be independent of the business cycle phenomenon. This dichotomy, which is

1. Textbook treatments of Keynesian economics treat business cycles as fluctuations around a long-term deterministic trend. Sophisticated Keynesian macroeconometric models, such as the Fair model, incorporate a production function that determines output in the long-run. Rational expectations with misperceptions models of the cycle (Lucas, 1973) also have monetary impulses that move output temporarily from a trend level.
2. In Milton Friedman's (1968) words, the natural rate is "ground out by the Walrasian system of general equilibrium equations" (p. 8) even if unexpected monetary disturbances move output in the short-run (p. 9).

central to the neoclassical synthesis, superimposes business cycles as short-run disequilibrium phenomena on an economy in long-run equilibrium.

This business cycle/growth dichotomy has been vitiated by new research on two fronts. First, research on the time series properties of main economic aggregates indicates that output can be characterized as following an integrated process.[3] Extracting the long-run trend from data generated by integrated process cannot be accomplished by simple regression detrending methods. Auxiliary assumptions concerning the covariation of the trend and cyclical components of the data are necessary. Once covariation of the trend and cyclical components is allowed, the rationale for detrending loses much of its appeal.

Second, some recent theories of macroeconomic fluctuations attribute all of the variability in output to real factors.[4] These real business cycle theories account for fluctuations at all frequencies by the same shock. There is, then, no meaningful dichotomy between the short-run and the long-run.

In this paper, we take seriously the message of these challenges to the neoclassical synthesis; shocks that move the economy at business cycles frequencies may also affect the economy in the long-run. Indeed, we use economic theory about the long-run impacts of different shocks to identify our model. Yet, we do not take this challenge to its extreme. Specifically, we do not maintain that all fluctuations in output are attributable to growth shocks. To the contrary, we view fluctuations as arising from a mixture of shocks; our goal is to disentangle these shocks.

The key identifying restriction underlying our empirical work is the simple, but powerful assumption made by Blanchard and Quah (1988), which we state as:

The level of output is determined in the long-run by supply shocks, such as shocks to technology and labor supply.

This identifying assumption does not exclude the possibility that these shocks also account for the high frequency movements in output as they would, for example, in a real business cycle model. It also does not exclude the possibility that short-run fluctuations are largely explained by aggregate demand shocks, such as shocks to the money supply or velocity, or by shocks to fiscal policy or animal spirits. It only excludes the possibility that the aggregate demand shocks permanently affect the level of output. The assumption allows the data to choose a description closer to the Keynesian

3. See Nelson and Plosser (1982) and Campbell and Mankiw (1987a).
4. See Kydland and Prescott (1982), Long and Plosser (1983), and Prescott (1986).

view, in which fluctuations are predominantly transitory, or fit a description closer to the real business cycle view, in which fluctuations are largely the result of permanent shocks.[5]

In the next section of the paper, we sketch the economic model that guides our empirical analysis. In Section 3, we give the precise econometric specification. We present our findings in Section 4 and offer concluding remarks in Section 5.

2. *Model*

Our econometric specification is motivated by a model in which the long-run properties of real variables are determined by a simple neoclassical growth model. In this model, long-run movements in output can be attributed entirely to exogenous changes in labor input and technological progress. In the short-run, output may deviate from its long-run steady state value. These deviations may arise from shocks to the permanent levels of labor input and technology, which lead to a transition from one steady state to another, or they may be caused by aggregate demand disturbances. Hence, movements in output arise from three sources: labor supply disturbances, technological disturbances, and aggregate demand disturbances. The first two of these—the supply shocks—have a permanent effect on the level of output; the third has only a temporary effect.

Interest rates and the rate of inflation are also included in the empirical model. All three sources of shocks are allowed to have both long-run and short-run effects on the level of inflation and the level of the nominal interest rate, but not on the real interest rate.

Two identifying assumptions allow us to separate these three sources of shocks from a dynamic reduced form, which includes labor input, output, inflation, and nominal interest rates. The first was alluded to above, aggregate demand disturbances have no long-run effect on output. This assumption allows us to determine the historical influence of aggregate demand and aggregate supply on the variables in the model. The second identifying restriction is that the long-run level of labor supply is exogenous. This assumption allows us to divide the aggregate supply effect into

5. Blanchard and Quah (1988) use this assumption in a bivariate model of output and unemployment. They assume that output is integrated, but that unemployment is stationary, and that supply shocks are responsible for the stochastic growth component of output. Other researchers have relied on the distinction between permanent and transitory shocks for identification. Campbell and Mankiw (1987b) identify long-run movements in output as the part of output orthogonal to unemployment changes. King, *et al.* (1987) identify the long-run movement in output as the common long-run component in output, consumption, and investment. Blanchard (1986) analyses a model where the identifying assumption is long-run homogeneity of demand schedules.

the components arising from labor input and from technology. In the long-run, labor supply is influenced neither by aggregate demand nor by the level of technology. We could relax this assumption to allow permanent real wage growth to affect labor supply, but doing so would only affect the decomposition of the permanent supply component into labor supply and technology. The decomposition between supply and demand would not be affected.

In standard models of long-term growth, the shocks to technology and labor supply together with capital accumulation determine the level of output in the long-run. Suppose that labor supply and technology evolve according to

$$h_t^* = \delta_h + h_{t-1}^* + \Theta_h(L)v_t \quad (2.1)$$

and

$$\varepsilon_t^* = \delta_\varepsilon + \varepsilon_{t-1}^* + \Theta_\varepsilon(L)e_t \quad (2.2)$$

where ε_t^* and h_t^* are the log levels of technology and labor supply and where v_t and e_t are serially and mutually uncorrelated shocks. The lag polynomials $\Theta_h(L)$ and $\Theta_\varepsilon(L)$ are assumed to have absolutely summable coefficients and roots outside the unit circle. That is, the dynamics described by the polynomials are transitory.[6]

We define the long-run log level of output y_t^* as

$$y_t^* = \alpha h_t^* + (1 - \alpha)\, k_t^* + \varepsilon_t^* \quad (2.3)$$

where k_t^* is the long-run level of capital. That is, we assume that the production function is Cobb-Douglas in the long-run. Yet, as shown below, we allow output to deviate in the short-run from this relationship.

We now introduce our first restriction from economic theory by assuming that the steady state capital-output ratio is a constant

$$k_t^* = y_t^* + \eta \quad (2.4)$$

where η is the constant log capital-output ratio. The Solow-Swan[7] growth model would generate a constant η, which is a function of δ_h, δ_ε, and the

6. Unless otherwise stated, all of the lag polynomials that we use in this paper will have these properties. Thus, they will always give rise to transitory dynamics. Where necessary we will invert them.

7. Solow (1956) and Swan (1956).

economy's saving and depreciation rates. Substituting (2.4) into (2.3) and rearranging yields

$$y_t^* = h_t^* + (1/\alpha)\,\varepsilon_t^* \quad (2.5)$$

where the constant $\eta(1 - \alpha)/\alpha$ is suppressed.

If we were willing to identify y_t^* and h_t^* with the actual log levels of output and labor, the equations above would define a real business cycle model with a much simpler propagation mechanism for the shocks than, for example, Kydland and Prescott's (1982). We close our model, however, by adding aggregate demand disturbances that allow output and inputs to deviate temporarily from their long-run levels.

To allow output and labor to move independently of the labor and productivity shocks in the short-run, we introduce two aggregate demand shocks, denoted by ν_t^1 and ν_t^2. These can be thought of as goods market (IS) and money market (LM) shocks. They are assumed to be serially uncorrelated and uncorrelated with the growth shocks. We cannot disentangle these shocks. Reasonable specifications of the goods and money market do not restrict just one of these shocks to affect the price level in the long-run. Both labor input, h_t, and output, y_t, can deviate temporarily from their long-run values because of these aggregate demand shocks, or because of transitory adjustments to permanent labor and or technology shocks.[8] Namely,

$$h_t = h_t^* + \Xi_h(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]' \quad (2.6)$$

and

$$y_t = y_t^* + \Xi_y(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]'. \quad (2.7)$$

The dependence of h_t on all of the shocks in the model allows flexible responses of labor to aggregate demand and real wages. Equation (2.6) allows labor supply to be elastic in the short-run. Indeed, in the short-run, workers can be off their labor supply schedules. Output and hours can deviate from their long-run levels as they would in a wide range of models, such as the inflation-augmented Phillips curve, the Lucas supply model, or the Fischer-Taylor contract model. Moreover, equations (2.6) and (2.7) break the tight link between output and inputs so that "off the production

8. Tobin's (1955) dynamic aggregative model is the first to superimpose a business cycle model on a neoclassical growth model. It features wage inflexibility as the source of cyclical fluctuations.

function" behavior or labor hoarding can be captured in the estimates. We only assume that the production function holds in the long-run (equation 2.3).

Differencing (2.6) and (2.7) and applying (2.1) and (2.2) yields

$$\Delta h_t = \Theta_h(L)v_t + (1 - L)\Xi_h(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]' \quad (2.8)$$

and

$$\Delta y_t = \Theta_y(L)v_t + \alpha^{-1}\Theta_\varepsilon(L)\varepsilon_t + (1 - L)\Xi_y(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]' \quad (2.9)$$

which are two of the reduced form equations that we estimate.[9]

To complete the model we add equations describing the inflation rate and the nominal interest rate. The inflation reduced form is

$$\Delta\pi_t = \Xi_\pi(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]', \quad (2.10)$$

which implies that the rate of inflation is integrated, its first difference is stationary, and that all of the shocks can have a long-run effect on the level of inflation.

Equations (2.3) and (2.4) imply that the long-run real interest rate is constant. Shocks to the system can have only short-run effects on the real rate, so the real rate is stationary. Given the definition of the real interest rate as the difference between the nominal interest rate and the expected inflation rate, the restriction on the real rate implies a restriction on the joint behavior of the nominal interest rate i_t and the inflation rate. Specifically, the nominal interest rate and the inflation rate are cointegrated, leading to the reduced form

$$i_t - \pi_t = \Xi_i(L)\,[v_t\; e_t\; \nu_t^1\; \nu_t^2]. \quad (2.11)$$

Summarizing, the model can be written as

9. Here, and for the remainder of the paper, constant terms are suppressed. They are included in the estimated equations.

$$\begin{bmatrix} \Delta h_t \\ \Delta y_t \\ \Delta \pi_t \\ i_t - \pi_t \end{bmatrix} = A(L) \begin{bmatrix} v_t \\ e_t \\ \nu_t^1 \\ \nu_t^2 \end{bmatrix}. \quad (2.12)$$

The matrix polynomial A(L) is a function of the polynomials $\Xi_h(L)$, $\Xi_y(L)$, $\Xi_\pi(L)$, $\Xi_i(L)$, $\Theta_h(l)$, and $\Theta_y(L)$ appearing in (2.8) through (2.11). Our identifying restrictions can be written in terms of the long-run multipliers, that is, the elements of $A(1)$. Setting the lag operator L equal to one in (2.8) and (2.9) shows that the long-run multiplier from ν_t^1 and ν_t^2 to h_t and y_t are zero, and that the long-run multiplier from e_t to h_t is zero. Consequently, the matrix of long-run multipliers $A(1)$ is lower block triangular, i.e.

$$A(1) = \begin{bmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix}. \quad (2.13)$$

Because we place no restrictions on a_{34} the identification scheme we employ cannot be used to separate the two aggregate demand shocks.[10] We report only their joint impact in our empirical analysis.

The model summarized in (2.12) and (2.13) might reasonably characterize aggregate hours, output, inflation, and interest rates, were it not for the large oil shocks that occurred in the 1970s and 1980s. We introduce exogenous oil price changes into our model. Below, we support this specification for oil shocks. We also assume that oil price changes have no long-run effect on labor supply, which is consistent with our assumption that there are no wealth effects in labor supply. Oil prices are allowed to have a permanent effect on all of the other variables in the model. Denoting the change in real oil prices by

$$\Delta o_t = \xi_t \quad (2.14)$$

the model becomes

10. With conventional exclusion restrictions, which we adjure in this paper, one could identify these shocks.

$$\begin{bmatrix} \Delta h_t \\ \Delta o_t \\ \Delta y_t \\ \Delta \pi_t \\ i_t - \pi_t \end{bmatrix} = C(L) \begin{bmatrix} v_t \\ \xi_t \\ e_t \\ \nu_t^1 \\ \nu_t^2 \end{bmatrix} \quad (2.15)$$

where $C(1)$ retains the lower block triangular structure of $A(1)$.

We estimate equations (2.15) and discuss the results in Section 4. Before proceeding to that discussion, we give details in the next section of the econometric method and specification.

3. *Econometric Method and Specification*

In this section, we present the precise form of our estimated equations and discuss how we impose the identifying restrictions introduced in the last section. These restrictions are a combination of covariance restrictions and restrictions on long-run multipliers. There are several equivalent methods for imposing these identifying restrictions. We discuss a simple instrumental variables approach.

We assume that the $C(L)$ in equation (2.15) is invertible, so that it can be written as

$$D(L)\, X_t = \omega_t \quad (3.1)$$

where $D(L) = C(L)^{-1}$, X_t is the 5×1 vector $(\Delta h_t, \Delta o_t, \Delta y_t, \Delta \pi_t, i_t - \pi_t)'$ and ω_t is the vector of disturbances $(v_t, \xi_t, e_t, \nu_t^1, \nu_t^2)'$. Following the assumptions made in Section 2, we assume that the roots of $|D(z)|$ are outside the unit circle and that ω_t is vector white noise. Our goal in the empirical analysis is to use the observed data to estimate the disturbances ω_t and the moving average polynomial $C(L)$. To do so, we appeal to identifying assumptions derived from the model in Section 2. The classical approach to the identification problem is to impose exclusion restrictions in the equations so that "endogenous" variables have no effect on "exogenous" variables, and specific exogenous variables affect some, but not all of the endogenous variables. Criticisms of these restrictions are well known. In rational expectations models, restrictions across the coefficients in $D(L)$ and covariance restrictions on the matrix of structural disturbances are used to identify the model. These restrictions typically impose tight constraints on the dynamics of the model.

In "structural" VAR approaches (Bernanke (1986), Blanchard and Watson (1986), or Sims (1986)), the dynamics of the model are left unconstrained and identification is achieved by imposing constraints on contemporaneous relations of the data through $D(0)$ and the covariance matrix of ω_t. These restrictions are similar to the classical exclusion restrictions and are often difficult to justify on *a priori* grounds.

An alternative identification scheme is used by Blanchard and Quah (1988).[11] They constrain $D(1)$, the long-run multipliers, as well as the covariance matrix of ω_t, to identify the model. We use this approach in our empirical analysis. In particular, we use the block lower triangular structure of $D(1)$ (inherited from $C(1)$), together with the assumption that the supply shocks v_t, e_t, and ξ_t are mutually uncorrelated with each other and uncorrelated with the demand disturbances, to identify the supply disturbances, the impulse response functions of these disturbances, and a linear combination of the demand disturbances. To this end, we write the first equation of (3.1), the equation for Δh_t as

$$\Delta h_t = \sum_{j=1}^{p} \beta_{hh,j}\Delta h_{t-j} + \sum_{j=0}^{p} \beta_{ho,j}\Delta o_{t-j} + \sum_{j=0}^{p} \beta_{hy,j}\Delta y_{t-j}$$

$$+ \sum_{j=0}^{p} \beta_{h\pi,j}\Delta \pi_{t-j} + \sum_{j=0}^{p} \beta_{hi,j}(i_{t-j} - \pi_{t-j}) + v_t. \quad (3.2)$$

Since $D(1)$ is the lower triangular, the long-run multipliers from Δo_t, Δy_t, $\Delta \pi_t$, and $i_t - \pi_t$ to Δh_t are zero, so the coefficients of their lags each sum to zero. Imposing these constraints yields

$$\Delta h_t = \sum_{j=1}^{p} \beta_{hh,j}\Delta h_{t-j} + \sum_{j=0}^{p-1} \gamma_{ho,j}\Delta^2 o_{t-j} + \sum_{j=0}^{p-1} \gamma_{hy,j}\Delta^2 y_{t-j}$$

$$+ \sum_{j=0}^{p-1} \gamma_{h\pi,j}\Delta^2 \pi_{t-j} + \sum_{j=0}^{p-1} \gamma_{hi,j}(\Delta i_{t-j} - \Delta \pi_{t-j}) + v_t \quad (3.3)$$

so that only differences of Δo_t, Δy_t, $\Delta \pi_t$, and $i_t - \pi_t$ enter the equation. Clearly, equation (3.3) cannot be estimated by ordinary least squares, since it includes contemporaneous values of some of the regressors, which are correlated with v_t. We estimate the equation by instrumental variables using lags one through p of Δh_t, Δy_t, $\Delta \pi_t$, $i_t - \pi_t$, and lags zero through p

11. A closely related identification procedure is employed in King, *et al.* (1987).

of Δo_t as instruments. The current value of Δo_t can be used because it is exogenous.

Similarly, the equation for Δy_t is specified as

$$\Delta y_t = \sum_{j=1}^{p} \beta_{yh,j}\Delta h_{t-j} + \sum_{j=0}^{p} \beta_{yo,j}\Delta o_{t-j} + \sum_{j=1}^{p} \beta_{yy,j}\Delta y_{t-j} + \sum_{j=0}^{p-1} \gamma_{y\pi,j}\Delta^2\pi_{t-j} + \sum_{j=0}^{p-1} \gamma_{yi,j}(\Delta i_{t-j} - \Delta\pi_{t-j}) + \beta_{yv}v_t + e_t \quad (3.4)$$

where the differences of $\Delta\pi_t$ and $i_t - \pi_t$ are included in the equation to impose the constraint that the long-run multipliers from π_t and $i_t - \pi_t$ to Δy_t are zero. Equation (3.4) can be estimated using the same set of instruments as (3.3) plus, v_t, the estimated residual for (3.3). Recall that v_t is uncorrelated with e_t. The instrumental variables procedure makes their sample analogues uncorrelated by construction.

The equations estimated for π_t and $i_t - \pi_t$ are reduced forms. They are

$$\Delta\pi_t = \sum_{j=1}^{p} \beta_{\pi h,j}\Delta h_{t-j} + \sum_{j=0}^{p} \beta_{\pi o,j}\Delta o_{t-j} + \sum_{j=1}^{p} \beta_{\pi y,j}\Delta y_{t-j} + \sum_{j=1}^{p} \beta_{\pi\pi,j}\Delta\pi_{t-j} + \sum_{j=1}^{p} \beta_{\pi i,j}(i_{t-j} - \pi_{t-j}) + \beta_{\pi v}v_t + \beta_{\pi e}e_t + a_t^1 \quad (3.5)$$

and

$$i_t - \pi_t = \sum_{j=1}^{p} \beta_{ih,j}\Delta h_{t-j} + \sum_{j=0}^{p} \beta_{io,j}\Delta o_{t-j} + \sum_{j=1}^{p} \beta_{iy,j}\Delta y_{t-j} + \sum_{j=1}^{p} \beta_{i\pi,j}\Delta\pi_{t-j} + \sum_{j=1}^{p} \beta_{ii,j}(i_{t-j} - \pi_{t-j}) + \beta_{iv}v_t + \beta_{ie}e_t + a_t^2. \quad (3.6)$$

The error terms a_t^1 and a_t^2 are linear combinations of the structural aggregate demand shocks v_t^1 and v_t^2. Since these disturbances are uncorrelated with the regressors, equations (3.5) and (3.6) can be estimated by ordinary least squares. We include the estimated v_t and e_t in equations (3.5) and (3.6) as regressors and instruments; the estimated a_t^1 and a_t^2 are by construction uncorrelated with those estimated supply shocks.[12]

Finally, oil prices are exogenous, so they are simply specified as

$$\Delta o_t = \xi_t. \quad (3.7)$$

All equations include constant terms. The results from estimating (3.3) through (3.7) are the subject of the next section.[13]

4. *Results*

4.1. DATA

The variables considered in our model are total hours worked (h_t), output (y_t), inflation (π_t), the nominal interest rate (i_t), and real oil prices (o_t). The Appendix gives the details of the sources of the data. Estimates reported in this paper are based on quarterly U.S. data from 1951:1. Data before 1951 are used as initial conditions in autoregressions. The end of the sample period is discussed below. The data for labor hours, output, and price are for the nonfarm private economy, excluding housing. We choose output for the nonfarm, non-housing private sector, rather than the whole economy because there are serious conceptual difficulties in relating the output to the inputs of housing, government, and farms. Housing and government are imputed in the national accounts. Farmers are largely

12. In the RATS packages, the equations can be estimated without including the disturbances and then transformed via the standard Cholesky decomposition. This decomposition picks out a different linear combination of the aggregate demand shocks, but since only their joint effect is identified, this difference is inessential.
13. Blanchard and Quah (1988) use a different technique to estimate models subject to these long-run Wold causal orderings. They estimate the unrestricted vector autoregression for X_t and then transform the system by post-multiplying the VAR by a matrix that imposes the necessary restrictions on the long-run multipliers and the residual covariance matrix. There is unique matrix that simultaneously diagonalizes the VAR innovation covariance matrix and triangularizes the matrix of long-run multipliers. When the only constraints on the system are a lower triangular matrix of long-run multipliers and a diagonal innovation covariance matrix, the model is just-identified, and this procedure can be thought of as "indirect least squares." The instrumental variable approach that we outline can be thought of as two stage least squares. When the model is just-identified, these two estimation methods produce identical estimators and are equivalent to the FIML estimator. The model that we estimate is overidentified. In particular, oil prices are assumed to be strictly exogenous, and this imposes overidentifying restrictions. These overidentifying restrictions are easy to impose in our instrumental variable approach, but are much more difficult to impose in the indirect least squares approach.

self-employed, so measures of their hours of work are unreliable. Moreover, studying the nonfarm business sector allows us to abstract from the major changes in aggregate labor productivity caused by workers leaving farms.[14]

4.2. DATA ANALYSIS

Our modelling and estimation strategy depends critically on the correct differencing of our time series. In Table 1 we present a variety of unit root test statistics that underlie our choice of specification. In the top panel we present the familiar Dickey-Fuller *t*-statistics, which test for a root of unity, versus a root less than unity. In the next column we present the largest estimated root from a sixth order autoregression, denoted by $\hat{\rho}$. In the hours, output, and productivity regressions we included a time trend in the autoregression to eliminate deterministic drift in these series. The *t*-statistics for hours, output, labor productivity, inflation, and interest rates are far less extreme than the 10 percent critical values. The estimated values of $\hat{\rho}$ are less than unity, but under the null hypothesis of a unit root, these estimates have a substantial negative bias. As pointed out in Schwert (1987), this bias is particularly severe when the first differences of the data have a large moving average component. Such moving average components might explain the small value of $\hat{\rho}$ for inflation.

Unit root tests cannot be performed on the unobserved *ex ante* real interest rate; we present results for the *ex post* real rate. Since the null hypothesis of a unit root in the *ex ante* real rate implies a unit root in the *ex post* real rate, little is lost in this substitution. The results for the *ex post* real interest rate $i_t - \pi_{t+1}$ are qualitatively different from the results for the other variables. The Dickey-Fuller *t*-statistic is much closer to the 10 percent critical value (its *p*-value is approximately 12 percent), and the estimated value of ρ is only 0.81. Thus, there is stronger evidence supporting the hypothesis that the real rate is stationary than there is supporting the hypothesis that the other variables are stationary.

In the bottom panel of the table we present the multivariate unit root tests developed in Stock and Watson (1987). The first statistic, $Q_f(4,3)$, tests the null hypothesis of 4 versus 3 unit roots among the four variables h_t, y_t, π_t, and i_t. The null of four unit roots is strongly rejected: the *p*-value of the test is 0.3 percent. The data, therefore, appear to be cointegrated. The next statistic, $Q_f(3,2)$ tests for 3 versus 2 unit roots in the four variable system. Here the data are consistent with the null of 3 unit roots: the *p*-value for the test is 85 percent. Thus, there appears to be only one cointegrating relationship among the data.

14. See Denison, 1974, pp. 62–4.

In summary, these results suggest that h_t, y_t, $y_t - h_t$, π_t, and i_t each contain a unit root, that there is one cointegrating relationship, and that the stationary linear combination of the data is $i_t - \pi_t$, implying a stationary real interest rate. Recall that stationarity of the real interest rate is one of the restrictions imposed on the data by our neoclassical model of long-term growth.

Unit root tests never provide sharp discrimination between the unit root hypothesis and the hypothesis that the data are stationary, but highly serially correlated. It is possible, especially in the case of inflation, that we are making a type two error by falsely accepting the null of a unit root, or in the case of the real rate, making a type one error by falsely rejecting the null. The univariate results for the nominal interest rate suggest that either inflation or the real rate has a root very close to unity. If the large root is less than one, then an expectations theory of the term structure suggests that interest rates should become more stationary (that is, have smaller $AR(1)$ coefficients) as the term increases. But, interest rates do not get more stationary as the term increases. The values of $\hat{\rho}$ for 6-month, 1-year, 5-year, 10-year, and 20-year nominal Federal interest rates vary between 0.96 and 0.98. The conclusion is that either inflation or the real rate has a unit root. Our data analysis, together with our priors, leads us to accept the unit root in inflation and reject the unit root in the real rate.

Table 1 UNIT ROOT DESCRIPTIVE STATISTICS

A. Univariate Unit Root Tests		
Series	Dickey-Fuller t-statistic	$\hat{\rho}$
h_t	−2.71	.93
y_t	−2.47	.93
$y_t - h_t$	−0.99	.98
π_t	−1.96	.85
i_t	−1.84	.96
$i_t - \pi_{t+1}$	−2.48	.81
B. Multivariate Unit Root Tests Four Variable System: h_t, y_t, π_t, i_t		
4 vs. 3 Unit Roots	$Q_f(4,3) = -62.84$	p-value = 0.3 percent
3 vs. 2 Unit Roots	$Q_f(3,2) = -13.83$	p-value = 84.8 percent

Note: The Dickey-Fuller t-statistics are calculated from a regression, including six lags of the differenced data. The regressions for h_t, y_t, and $y_t - h_t$ included a constant and time trend. The regressions for π_t, i_t, and $i_t - \pi_{t+1}$ included a constant. The 10% critical values for h_t, y_t, and $y_t - h_t$ are −3.12. The 10% critical values for π_t, i_t, and $i_t - \pi_{t+1}$ are −2.57. $\hat{\rho}$ is largest autoregressive root in the sixth order autoregression used to calculate the Dickey-Fuller t-statistics. The multivariate tests (Q_f) are described in Stock and Watson (1987). They are calculated using linearly detrended data with a VAR(6) correction.

Finally, before proceeding to the results, we offer support for our specification of exogenous oil price changes. Oil prices are, in principle, endogenous. On average, real oil prices should increase by the real rate of interest, with innovations in the price reflecting shocks to demand and supply. Yet, over our sample period oil price changes are dominated by four exogenous events: the Yom Kippur War in 1973, the fall of the Shah in 1979, price decontrol in 1981, and the 1986 "collapse" of OPEC. That these events dominate the data is obvious from Figure 1, which plots the percentage change in real oil prices over the sample period.

4.3. RESULTS FOR BASIC MODEL

We estimated the model in equations (3.3) through (3.7) using six lags of the data together with a constant. Initially, we carried out the analysis using data through 1987:2, but it quickly became obvious that this led to a

Figure 1 OIL PRICE CHANGES

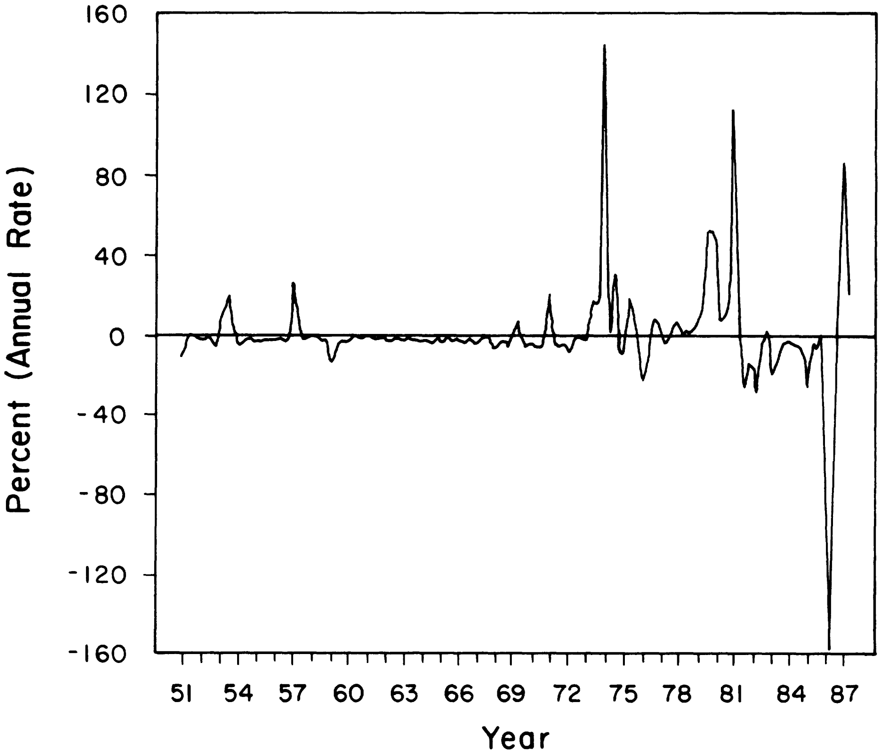

possible serious misspecification for the role of oil prices. The largest oil shock during the sample period occurred during the 1986 collapse of OPEC: during 1986 oil prices fell 50 percent. This dramatic decrease in prices coincided with sluggish growth. Averaging this period of positive covariation between oil price changes and output growth together with the 1974–1975 and 1979–1981 periods of negative covariation misses the possibility that the dynamic response of the variables in the model is different for oil price decreases than it is for oil price increases. The most straightforward way to allow for this asymmetric response is to interact the lags of oil prices over the 1986–1987 period with a dummy variable. Since we allow six lags of oil prices in our model, full interaction of the lags with the dummy variable over the 1986:1–1987:2 period results in a perfect fit over that period. Consequently, we present results for the estimated model using data through 1985:4.

The results for the estimated model are summarized in Figure 2 and Table 2. The graphs give the response of the logs of labor, output, the price level, inflation, and the nominal and *ex ante* real interest rates to shocks in labor supply, oil and technology.[15] The impulse responses are normalized as follows: the labor supply shock has a unit long-run impact on hours, the oil shock represents a 1 percent increase in oil prices, and the technology shock has a long-run impact of 1.6 on output. The long-run elasticity of output with respect to technology is $1/\alpha$ (see equation 2.5). Since the share of labor averages approximately 0.625, our impulse response functions trace out the effect of a 1 percent long-run increase in technology.

Since our identification procedure does not enable us to untangle the two aggregate demand shocks, we do not report the aggregate demand impulse response functions. Any impulse response functions that we reported would depend on arbitrary normalizations that would make interpretation difficult.

A 1 percent shock in long-run labor supply has a 0.4 percent impact effect on hours. After five to six quarters, hours reach 80 percent of their long-run level. The labor shock increases output by 0.6 percent in the long-run. Recall that we expect a unit long-run elasticity of output with respect to the labor supply shock. We cannot reject the null that the elasticity is one.[16]

Oil price increases lead to reductions in hours and in output. The output response reaches a trough after six quarters when a 1 percent oil price

15. The *ex ante* real interest rate is computed using the expected inflation rates implied by the model.
16. The t-statistic for this null hypothesis equals 1.7.

increase leads to a decline in output of 0.1 percent. The point estimate of the long-run elasticity of output with respect to oil prices is −0.07. Oil prices have a small, positive long-run effect on inflation. A 1 percent

Figure 2 IMPULSE RESPONSE FUNCTIONS: STOCHASTIC TREND IN HOURS

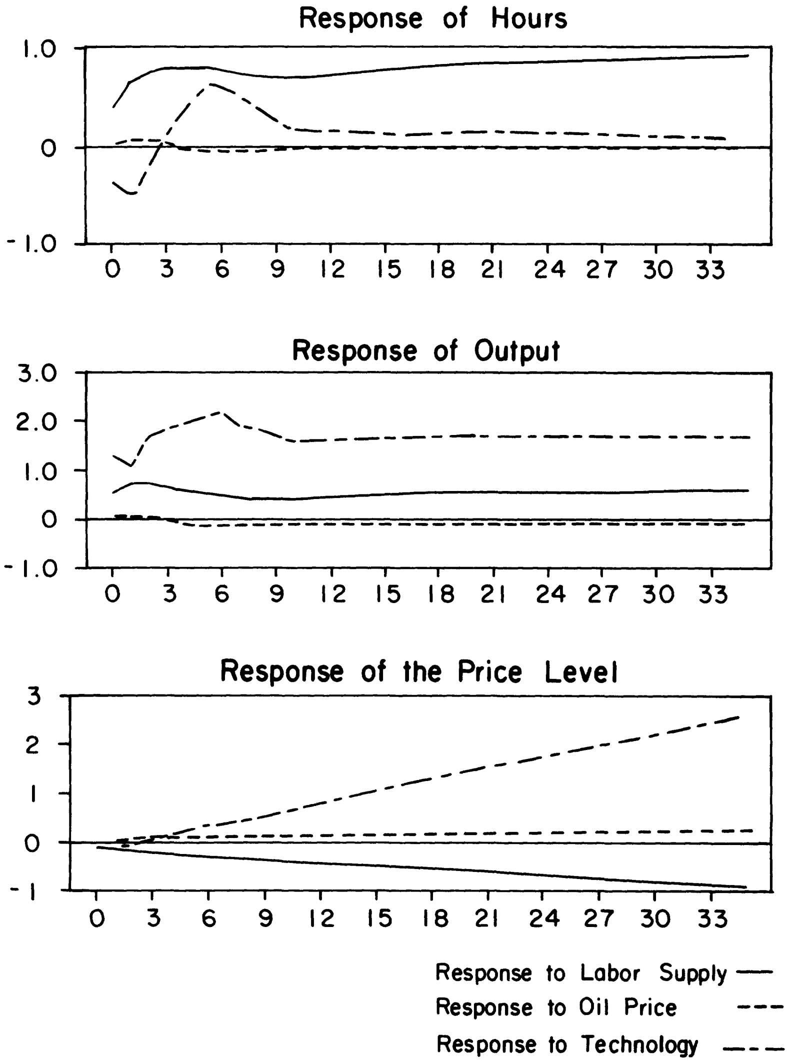

increase in oil prices leads to an increase in the price level of roughly 0.09 percent after two years.

Increases in technology have little effect on hours. Their effect on output is immediate; the impact effect of output is 80 percent of the long-run effect.

Figure 2 (Continued)

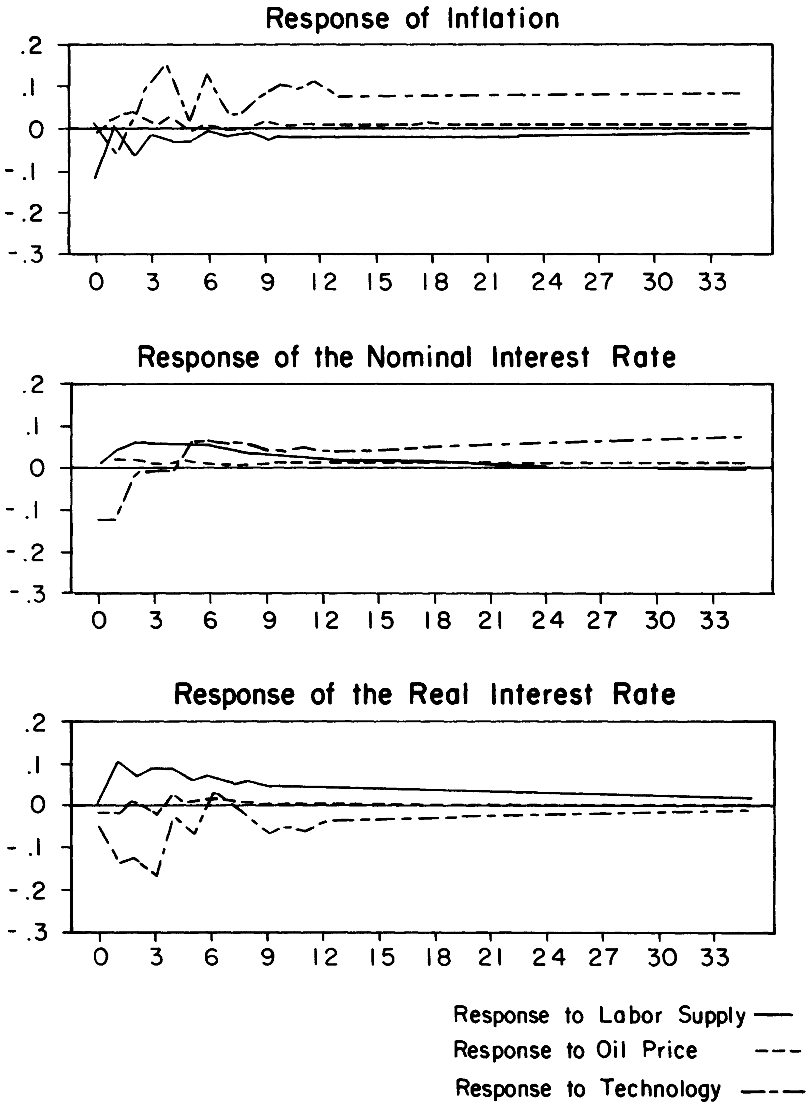

Table 2 DECOMPOSITIONS OF VARIANCE: STOCHASTIC TREND IN HOURS

	Fraction of hours explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	58.9 (20.8)	0.3 (2.0)	4.3 (10.1)	36.5 (19.5)
4	57.7 (21.0)	0.8 (2.6)	1.4 (7.2)	40.1 (19.0)
8	64.7 (19.4)	2.1 (3.0)	2.1 (6.3)	31.0 (15.8)
12	68.7 (17.9)	2.3 (3.1)	1.8 (5.9)	27.2 (14.0)
20	76.8 (14.6)	1.7 (2.3)	1.2 (5.2)	20.3 (11.1)
36	86.2 (10.0)	1.0 (1.4)	0.7 (4.3)	12.1 (7.2)
∞	100.0	0.0	0.0	0.0

	Fraction of output explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	45.9 (17.7)	0.8 (2.7)	25.2 (15.1)	28.1 (17.6)
4	48.4 (18.4)	1.1 (2.0)	22.2 (14.5)	28.3 (17.1)
8	40.1 (16.3)	8.0 (6.1)	32.3 (13.7)	19.5 (13.1)
12	38.1 (15.4)	9.9 (7.0)	35.3 (13.6)	16.7 (11.2)
20	40.3 (15.2)	10.6 (7.3)	36.6 (13.9)	12.5 (8.7)
36	45.2 (15.7)	10.4 (7.5)	36.5 (14.4)	7.8 (5.9)
∞	61.7 (20.8)	7.5 (9.4)	31.9 (15.9)	0.0

	Fraction of price level explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	10.3 (15.2)	0.4 (2.3)	0.0 (14.2)	89.2 (19.1)
4	6.3 (12.3)	5.6 (5.7)	0.1 (12.9)	88.1 (16.9)
8	5.3 (11.7)	6.2 (6.9)	0.5 (12.0)	88.0 (16.7)
12	4.3 (11.7)	4.7 (6.8)	0.8 (12.2)	90.2 (16.9)
20	3.6 (12.0)	3.6 (7.1)	1.4 (12.5)	91.5 (17.3)
36	3.0 (12.7)	2.9 (7.4)	1.8 (12.7)	92.3 (18.0)
∞	1.6 (15.3)	2.4 (9.8)	2.5 (11.1)	93.5 (19.1)

Note: See text for details of these computations. Standard errors are in parentheses.

Table 2 contains the variance decompositions for a variety of forecast horizons.[17] The table presents the fraction of the forecast errors variance for each of the variables that is attributed to each of the shocks. Since we can

17. The standard errors reported in Tables 2 and 3 were calculated using Monte Carlo simulation. The simulations were carried out using draws from the normal distribution for the innovations in hours, output, price, and the interest rate. The historical sample path of oil prices was used in all of the simulations. Three hundred Monte Carlo draws were carried out.

Table 2 DECOMPOSITIONS OF VARIANCE:
STOCHASTIC TREND IN HOURS (CONTINUED)

	Fraction of inflation explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	10.3 (15.2)	0.4 (2.3)	0.0 (14.2)	89.2 (19.1)
4	7.6 (10.7)	9.8 (5.9)	0.6 (10.7)	82.0 (14.5)
8	6.3 (9.3)	10.4 (5.4)	1.5 (9.2)	81.8 (12.4)
12	5.3 (9.3)	8.9 (5.5)	1.7 (9.4)	84.1 (12.6)
20	4.6 (10.0)	7.0 (5.7)	2.0 (10.1)	86.4 (13.5)
36	3.7 (11.5)	5.3 (6.2)	2.2 (11.0)	88.8 (15.0)
∞	1.6 (15.3)	2.4 (9.8)	2.5 (11.1)	93.5 (19.1)

	Fraction of real interest rate explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	10.3 (15.2)	0.4 (2.3)	0.0 (14.2)	89.2 (19.1)
4	14.9 (14.0)	3.2 (3.7)	1.5 (12.0)	80.4 (16.7)
8	21.3 (13.1)	7.3 (4.8)	2.3 (10.0)	69.1 (14.6)
12	22.5 (13.6)	6.6 (4.7)	2.2 (10.0)	68.6 (14.6)
20	24.0 (14.3)	6.7 (4.8)	2.3 (10.2)	67.1 (14.9)
36	24.9 (15.0)	6.5 (4.9)	2.3 (10.5)	66.3 (15.3)
∞	34.1 (17.6)	10.2 (8.5)	3.6 (11.1)	52.1 (19.1)

	Fraction of nominal interest rate explained by shock to			
Quarter	Labor supply	Oil	Technology	Aggregate demand
1	0.7 (5.5)	8.5 (6.5)	7.8 (10.9)	83.0 (12.1)
4	11.7 (10.5)	14.3 (7.8)	3.4 (7.0)	70.6 (12.2)
8	13.8 (12.4)	11.5 (8.2)	2.4 (6.4)	72.3 (13.2)
12	12.0 (12.3)	11.3 (8.8)	2.2 (6.7)	74.4 (13.7)
20	8.4 (11.9)	10.5 (9.1)	1.9 (7.3)	79.2 (14.4)
36	4.8 (12.2)	8.2 (8.9)	1.9 (8.7)	85.1 (15.5)
∞	1.6 (15.3)	2.4 (9.8)	2.5 (11.1)	93.5 (19.1)

Note: See text for details of these computations. Standard errors are in parentheses.

observe a linear combination of the aggregate demand shocks, we report the variance explained by aggregate demand. Our identifying restrictions imply that 100 percent of the variance of hours is explained by the labor supply shock at the infinite horizon, and that 100 percent of the variance of output is explained by shocks to labor supply, oil, and technology at the infinite horizon. At shorter horizons, aggregate demand is allowed to have an impact on these variables. The results in Table 2 suggest that this impact is substantial. Approximately 40 percent of the variability in hours and 30 percent of the variability in output over the one year horizon is attributed to aggregate demand. Shocks to technology account for roughly 20 percent of the variability in output, but explain little of the variation in hours during the first year. As the horizon increases from 4 to 8 to 20 quarters, the variability in output attributed to aggregate demand falls from 28 percent to 20 percent to 12 percent; the variability attributed to technology increases from 22 percent to 32 percent to 37 percent. Oil prices explain only a small fraction of the variability in output.

Our results are quite close to results found by other researchers who use different measures of output and different specifications. King, Plosser, Stock, and Watson (1987) find that about 30 to 40 percent of the 8 quarter ahead variability in per capita GNP can be attributed to transitory factors (corresponding to aggregate demand in our framework, but unspecified in theirs). Blanchard and Quah (1988) use data on real GNP and unemployment. They attribute from 10 to 40 percent of the 8 quarter ahead variability of GNP to the temporary, aggregate demand shock, depending on the detrending procedure for the unemployment rate. All of these results attribute less than half, but still a substantial fraction of the variance of output to shocks that have a temporary effect on the level of output.

Aggregate demand is the main determinant of the variability in prices, inflation, and the nominal and real interest rate. It explains approximately 90 percent of the variability in prices and inflation, from 70 to 90 percent of the variability in nominal interest rates, and roughly the same percentage of the *ex ante* real rate.

The variance decompositions show the importance of the shocks in explaining the *average* variability in output. Of equal importance is the role that these shocks played in specific historical episodes. Our procedure produces estimates of the quarter-to-quarter shocks. Because these are serially uncorrelated, they are difficult to interpret. In Figure 3, we plot the 8 quarter ahead forecast error in output and its components. These are simply an eight-period weighted average of estimated shocks, where the weights are given by the impulse response functions. Again, the parameter estimates are based on data through 1985:4.

A striking feature of the graph is the post-sample 1986–1987 period. Using the estimated model through 1985:4, the oil price decline during 1986 provides a dramatic stimulus to growth. Actual output growth was sluggish, so the positive stimulus from oil is countered by large negative contributions from labor and technology. Given the sharp drops in output following the oil shocks of the 1970's, as well as during the pre-OPEC era (Hamilton, 1983), our model predicts a strong increase in output following the big decline in oil prices in 1986. Given that the boom did not occur during that period of time, our procedure offsets the positive effect of the oil price decrease on output with negative shocks to the other components that permanently affect output. As noted above, one would like to accommodate this episode by allowing for an asymmetric response of output to oil prices. Given that oil prices fell by a large amount only once in the sample, it is not possible to estimate such an asymmetric response.

The results through 1985:4 are not complicated by the asymmetric response to oil prices. Throughout the early 1980, oil prices were important factors that negatively affected output. Declines in aggregate demand coincide nearly perfectly with the output "double dip"; the decline is particularly severe during 1981–1982. Labor supply is essentially neutral until the very end of the 1981–1982 recession when it turns down sharply, while technology does not play a role in the recessions of the 1980s.

Oil and labor supply are the major factors in the 1974–1975 recession. Aggregate demand does not play a role during this period, although it fluctuates noticeably in the years immediately following it.

The largest negative impact of technology occurs during the 1970 recession and corresponds to the beginning of the productivity slowdown. Note also that there is a lower frequency contribution of technology to the forecast errors in output that corresponds to the extremely strong performance of measured productivity growth in the 1960s, and its subsequent slowdown in the 1970s.[18]

In addition to its roles in the recessions of the 1980s, aggregate demand appears to have played the major role in the recessions of 1957–1958 and 1960.

Finally, at the beginning of the sample, there is a large movement in the labor supply variable related to the Korean demobilization. This anomaly remains in the results even if an exogenous variable accounting for military employment is included in the system.[19]

4.4. ROLE OF PERMANENT LABOR SHOCKS IN OUTPUT FLUCTUATIONS

A striking feature of our results is the large role that permanent labor supply shifts play in the variability of output at all frequencies. Labor supply explains 40 percent of the 8 quarter ahead variability in output (Table 2). Moreover, permanent shifts in labor input are the first or second most important factor in the recessions of 1954, 1958, and 1975 (Figure 3). Why do these results arise and should they be regarded as surprising?

18. Note that negative values for the contribution of technology in Figure 4 usually do not correspond to declines in the level of technology because it has a positive drift.
19. This variable is the ratio of military employment to the civilian labor force. Its movements, which closely match those of the ratio of Federal purchases to private output, are dominated by the Korean war and to a lesser extent by the Vietnam War.

Figure 3 COMPONENTS OF FORECAST ERROR FOR OUTPUT: STOCHASTIC TREND IN HOURS

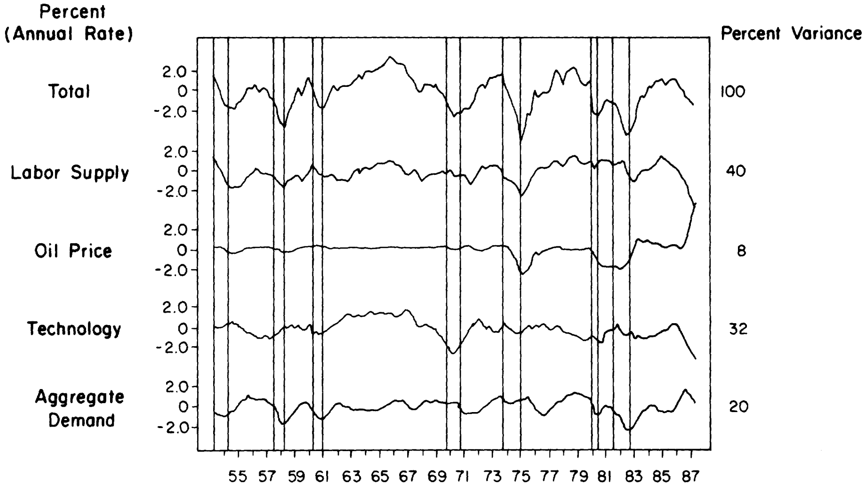

Economists have long attributed about half of long-run changes in the level of output to exogenous changes in labor input.[20] This decomposition of variance at very long horizons is almost entirely noncontroversial.[21] Now consider why, in our estimates, the shock to labor should be important at all frequencies. Labor supply shocks are important because we allow them to have a stochastic, rather than a deterministic trend and because the stochastic trend is estimated to have a large variance. Our findings are based on a simple, standard, and widely accepted model of long-term growth on which business cycle dynamics are superimposed. Because we find our specification so plausible, we are reluctant to dismiss it. Yet, because the important role of the permanent labor shock is inconsistent with our prior beliefs, we investigate alternative specifications.

4.4.1. Measure of Labor Input We measure labor input as total hours worked in the sector. Given that a production function is at the heart of our growth model, using hours worked as the labor variable is appropriate. For questions of low frequency movements in labor input, smoother variables, such as labor force or population are perhaps just as appropriate. In the notation of the model of Section 2, labor force or population could be used

20. See many careful studies by Denison (1974, for example) and others.
21. It has been challenged recently by Romer (1987).

in an equation for h_t^* with actual hours worked fluctuating in a stationary manner about h_t^*. In such a formulation, the labor supply shock would be the structural error in the labor force equation. The residual stationary deviations of hours from labor force would be attributed to aggregate demand.

This solution, attractive as it may seem, fails because the deviation of hours worked from labor force is not stationary. The first graph in Figure 4 shows the deviation of hours worked from labor force (in logarithms). This

Figure 4 HOURS AND LABOR FORCE

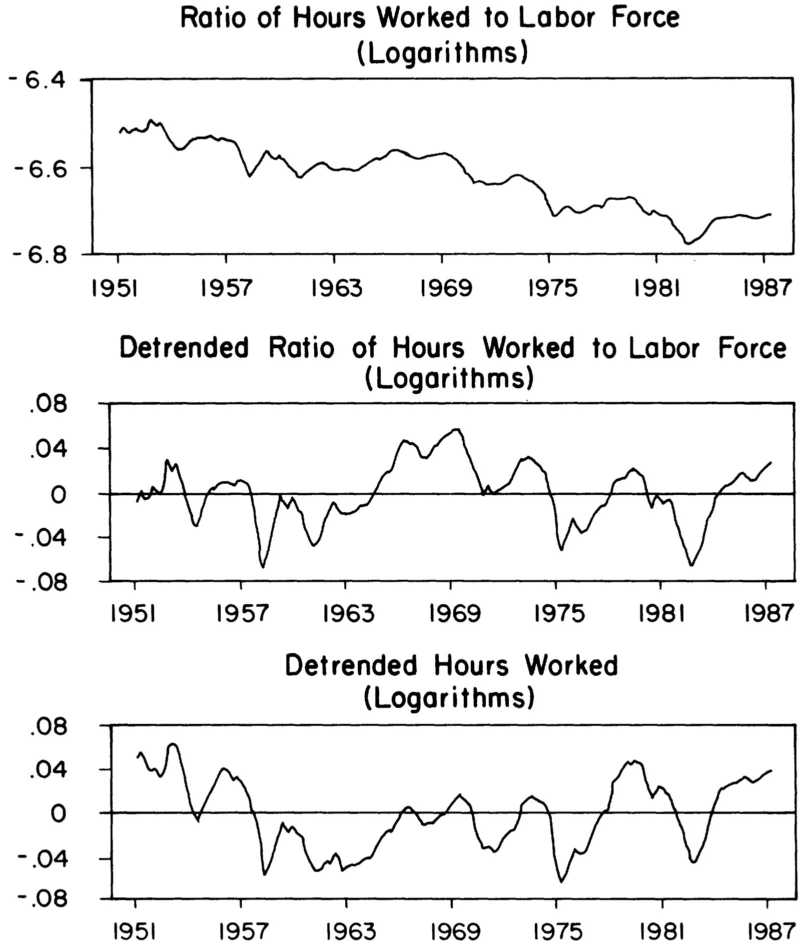

deviation clearly contains a trend.[22] The trend arises from a convolution of the decline in the average work week, the increase in female participation in the labor force, and the recent increase in part-time work. If we treated this trend as stochastic, it would play a role nearly identical to labor supply shock in the estimates just discussed. Alternatively, we could treat it as deterministic and abstract from issues of weekly hours and participation in the calculations.[23]

The fluctuations of the detrended deviation of the logs of hours and labor force are very similar to those of detrended log hours. These series are graphed in the second two panels of Figure 4. Because the series are so similar, the model we are about to discuss—one with trend-stationary hours but ignoring the labor force data—is very similar to the trend-stationary labor supply model which includes the labor force data.

4.4.2. Trend-Stationary Labor Supply As discussed above, we find that the permanent labor stock is important at all frequencies because labor appears to have a stochastic trend with large estimated variance. Harvey (1985), Watson (1986), and Clark (1987) point out that the sum of a stochastic trend (a random walk) and an independent, highly serially-correlated stationary process have an ARIMA representation with long-run properties that are poorly approximated by low order autoregressions. A low order autoregression could attribute some of the cyclical variability in the series to the stochastic trend. Therefore, the large stochastic trend in hours that we find may arise from a confusion between trend and stationary components.

To check for misspecification of this form we have carried out a variety of experiments, including doubling the lag length on all variables in the model, and doubling the lag length of the variable in the hours equation.[24] Qualitatively, the results are unchanged. Labor supply remains an important determinant of the business cycle variability in output. Including many lags of output in the hours equation should provide ample opportunity for removing the cyclical movements from its disturbance.

22. The *detrended* deviation of log hours from log labor force is highly serially correlated. Indeed, one marginally cannot reject the null hypothesis that it has a unit root. The deviation has a $\hat{\rho}$ of 0.90 and a Dickey-Fuller t-statistics of -3.4. See the notes to Table 1 for details of these computations.
23. Blanchard and Quah (1988) face a nearly identical problem in dealing with the trend in the unemployment rate. Their results are sensitive to whether or not unemployment is detrended.
24. This is only a partial response to the criticism, since we have estimated unconstrained autoregressive models. Proponents of unobserved component models would estimate parsimonious constrained ARIMA models. See below for a further discussion of this econometric issue.

The most extreme case of this misspecification occurs when hours contain no stochastic trend component and are characterized as stationary deviations from a deterministic trend. Differencing hours would introduce a unit moving average root into the model, which could not be inverted to yield an autoregressive representation. In this case, our models with 6 lags and our model with much longer lags would both be misspecified. It is unlikely, however, that they would give the same qualitative results. Even if the long lags could not eliminate the stochastic trend, they could make its variance small.

The estimates based on the differenced-stationary specification for labor are valid even if labor supply is trend stationary, but only if the estimation procedure allows for unit moving average roots. We do not undertake the difficult task of estimating a loosely parameterized vector ARMA model. Yet, it is instructive to consider the univariate ARMA process for hours to check for the presence of unit MA roots. Campbell and Mankiw (1987a) discuss the difficulties in estimating processes where a unit MA root might cancel an over-differenced dependent variable. For aggregate GNP, their results indicate that it is difficult to distinguish the trend-stationary AR(2) model from the ARIMA(1,1,1). For our log hours variable, the trend-stationary AR(2) estimates are (with constant and trend suppressed):

$$h_t = \underset{(.07)}{1.54}\, h_{t-1} - \underset{(.07)}{0.61}\, h_{t-2} + v_t, \quad \text{S.E.E.} = 0.757, \quad Q(36) = 26.5.$$

The ARIMA(1,1,1) estimates are (with the constant suppressed)

$$\Delta h_t = \underset{(.10)}{0.38}\, \Delta h_{t-1} + v_t + \underset{(.11)}{0.39}\, v_{t-1}, \quad \text{S.E.E.} = 0.776, \quad Q(36) = 29.3$$

Here S.E.E. is the standard error of estimate and $Q(36)$ is the Box-Pierce test.[25] Note that in the univariate setting there is no evidence that the moving average root is near unity. Were there a unit moving average root in the hours equation of the vector system, there would also be one in the univariate equation. Although the univariate test is not as powerful as a multivariate test, and we have explored only a limited number of ARIMA models, the univariate estimates do suggest that excluding MA components from the VAR estimates is not a serious problem. Hence, we believe that a unit moving average root is not a major source of misspecification.

25. The estimates of the ARIMA model are exact maximum likelihood and are computed using a computer program kindly provided by John Campbell.

Notwithstanding these findings, one can still argue that the estimate of $\hat{\rho}$ reported in Table 1 for hours is 0.93 which, if it was a precise estimate, would suggest that hours exhibit persistent, but stationary deviations about a linear trend. The estimate is not precise. A value of $\hat{\rho}$ equal to 0.93 is roughly the median value one would expect to find if the true value of ρ was 1. That is, there is significant downward bias in $\hat{\rho}$ when the true value of ρ is close to one. On the other hand, despite the bias in the estimate of $\hat{\rho}$, one also cannot reject the hypothesis that hours are borderline-stationary.

Prior knowledge is needed to resolve the problem. One possible prior is that the true underlying trend in hours comes from population growth whose trend is very smooth and is likely to be well-approximated by a deterministic function of time. An alternative prior is that the stochastic growth component in hours is trivially small compared to its stationary component. Both priors suggest that deviations of hours from a deterministic trend are, for all practical purposes, stationary.

Therefore, we present estimates consistent with this prior by estimating a model where labor is stationary around a deterministic trend. We view the estimates with detrended labor as an extreme but instructive case. They show the consequence of a prior that the stochastic trend in labor has low variance by taking the extreme position that the variance is zero. The trend-stationary model is a special case of our basic model with stochastic labor, but with the variance of the long-run component in labor set to zero. An econometric difficulty (estimating a loosely parameterized vector ARMA model) necessitates estimating the trend-stationary model as a separate, special case. In principle, it is nested by the stochastic trend model. If we estimate the stochastic labor model with labor differenced (and, in fact, the process is trend stationary) the estimated process will have a unit-moving average root, which should undifference the labor model. Yet, because we do not have explicit moving average components in our estimation, this undifferencing cannot take place in practice.

Specifically, the model with trend-stationary labor is as follows: Hours are assumed to be stationary around a deterministic trend. Output is still integrated, since we maintain the assumption that productivity is integrated.[26] Since detrended hours are now stationary, there are now three transitory shocks in the model. We now associate these shocks with aggregate demand. Oil prices and technology permanently affect the level of output. A summary of the results for this model can be found in Figures 5 and 6 and in Table 3.

26. From Table 1, the estimated $\hat{\rho}$ for average productivity is 0.98. Hence, there is less doubt about the non-stationarity of output or output per hours than for hours.

In Figure 5 we present the impulse response functions. The responses to changes in oil prices are much the same as they were in the model with stochastic labor supply growth. The responses to shocks in technology are different. Hours now fall sharply in response to shock to technology and output increases very slowly.

Table 3 presents the variance decompositions. Oil explains roughly the same fraction of output as it does in the model with differenced-stationary hours. The contributions from aggregate demand and technology are substantially different. In this model, aggregate demand explains 90 percent of output over the first year, and 80 percent at the 8 quarter horizon. Indeed even though we constrain aggregate demand to have no long-run effect on output, it still accounts for roughly 35 percent of the variability of output at the 8-year horizon. This result is a consequence of labelling shocks in the hours equation as aggregate demand rather than as labor supply. Recall that these shocks are very persistent.

The historical 8 quarter decomposition, shown in Figure 6, tells much the same story as the variance decompositions. Aggregate demand is now more important, oil retains its importance for the 1974 and 1980–1982 periods, and technology is somewhat less important.

The two sets of estimates tell markedly different stories about the sources of economics fluctuations in the postwar United States. Unfortunately, the data do not clearly support one model or the other. It is necessary to refer to priors when considering the likely role of permanent labor supply responses. While the models give very different answers to the question of the relative importance of transitory/permanent shocks, much of these differences can be attributed to the allocation of the shock to hours. That is, our results suggest that permanent components other than labor supply—productivity and oil prices—have been less important than is suggested by others. Productivity is somewhat more important at business cycle frequencies in the model with stochastic labor supply growth, but even there it explains only one-third of the 8 quarter variation in output.

4.5 SOLOW RESIDUAL

We would also like to incorporate explicitly a measure of technology, such as the Solow (1957) residual, into the estimation. It might seem consistent with our modelling strategy to assume that the long-run changes in the Solow residual measure long-run changes in technology. But a difficulty arises in using the Solow residual because it is inherently measured as a rate of change. If this measure contains errors due either to data or specification problems, these errors will accumulate in the measures of the level of technology. Hence, the accumulated Solow residual will contain a permanent component that is attributable to measurement error in addition to the permanent component that represents technological progress. Such difficulties could arise from measurement issues alone. Specifically, transitory measurement error in capital accumulation leads to a permanent error in the accumulated Solow residual. Additionally, if measured input flows

Figure 5 IMPULSE RESPONSE FUNCTIONS:
DETERMINISTIC TREND IN HOURS

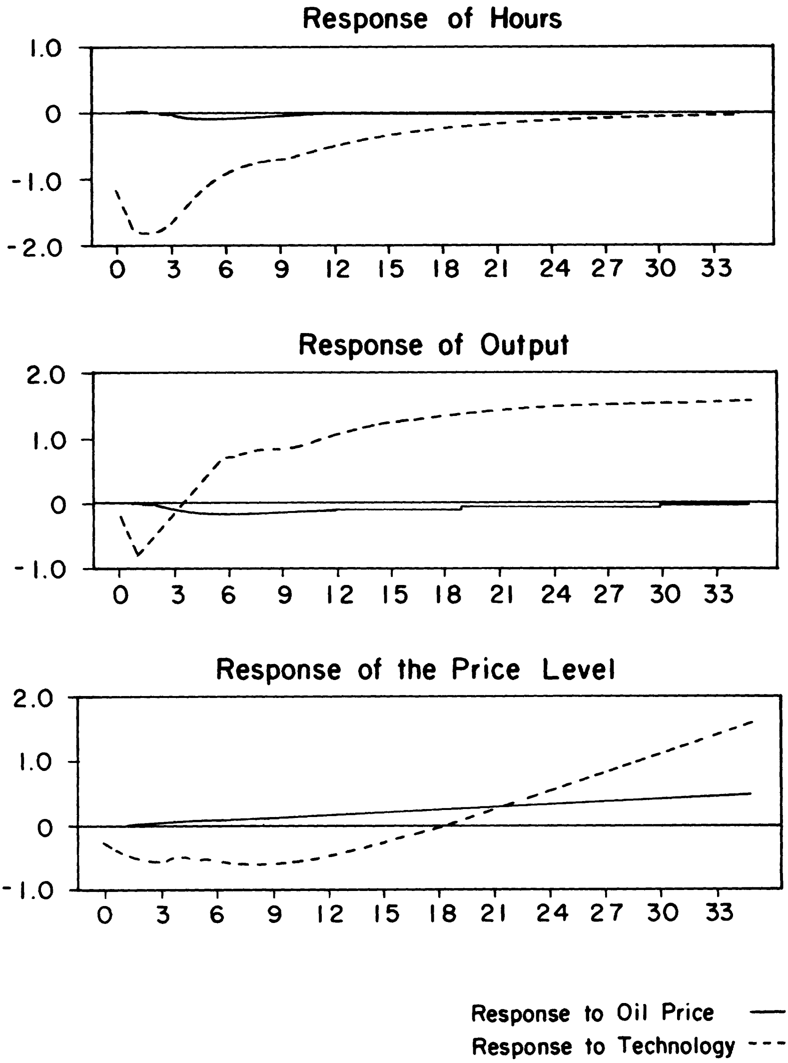

are not always equal to input services (labor is hoarded) then the accumulated Solow residual will have a permanent component similar to that arising from measurement error. Similarly, Hall (1988) shows that the measured Solow residual contains a business cycle component if the assumption of perfect competition is incorrect.

Figure 5 (Continued)

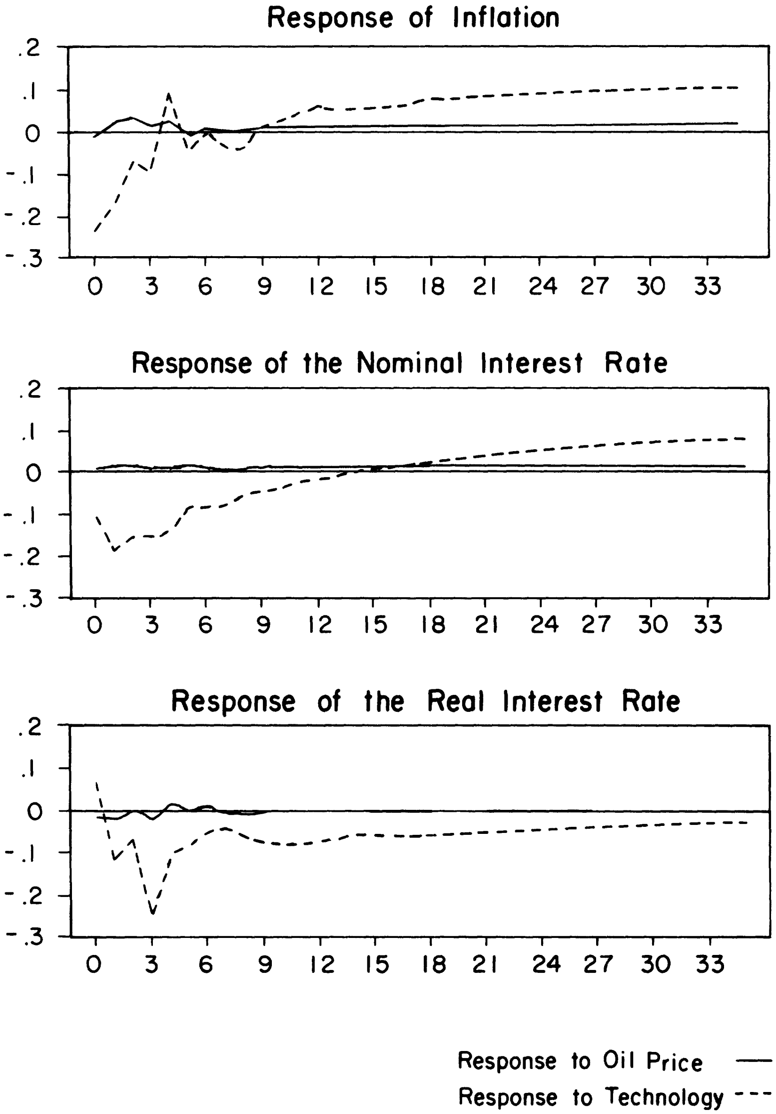

Figure 6 COMPONENTS OF FORECAST ERROR FOR OUTPUT: DETERMINISTIC TREND IN HOURS

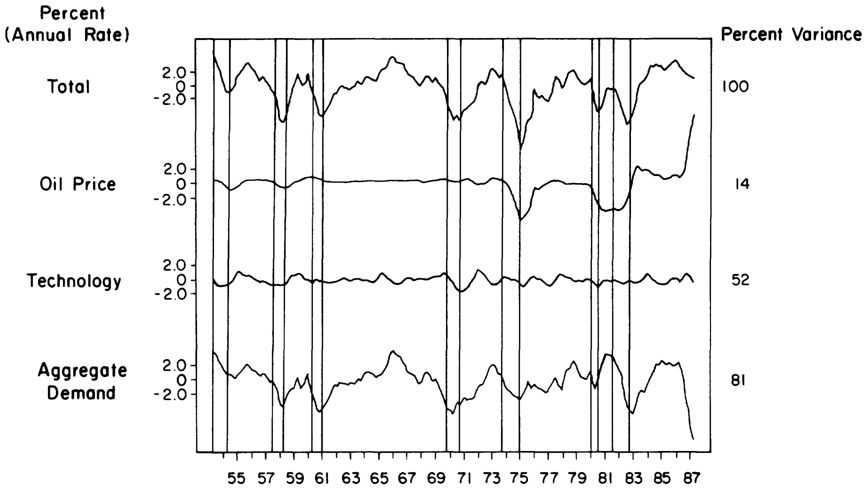

Despite these difficulties in explicitly incorporating the Solow residual into the model, it is interesting to see how our estimated technological shock relates to this widely-studied measure of technological progress. In the previous paragraph we suggest that the relationship at high frequencies is likely to be weak. Yet, if the measurement errors are fairly small, one might expect to find a relationship in the long-run between the Solow residual and our technological shock. We compute the fraction of variance at frequency zero of the Solow residual accounted for by our estimated shocks to technology.[27] In brief, we find that our technological shocks are closely related to the Solow residual at low frequencies. For our basic model with differenced-stationary hours, the technological shock accounts for 62 percent of the variation of the Solow residual in the long-run; for the model with trend-stationary labor, that figure is 75 percent.[28] Therefore, we

27. The definition of and data for the Solow residual are discussed in detail in the Appendix. The variance decompositions are computed based on a regression of the Solow residual on current and six lagged values of the five shocks, plus a constant six lags of the Solow residual itself.

28. In the hours differenced-stationary model, labor shocks account for 6%, oil price shock for 8%, aggregate demand for 16%, and the residual for 8% of the long-run variation in the Solow residual. In the hours trend-stationary model, the decompositions are 6% for labor, 4% for oil prices, 9% for aggregate demand, and 6% for the residual.

conclude that our estimated technological shock corresponds closely to more familiar estimates of technological progress.[29]

4.6. FURTHER CONSIDERATIONS

We conclude this section with a discussion of a few minor empirical issues, and some general comments about the identifying assumptions that we use. First, consider an empirical observation about the long-run output-capital ratio, which is assumed to be constant in our equation (2.4). With our data, the ratio wanders between 1.04 at the beginning of the sample to 0.85 at the end of the sample. Its sample path looks more like a random walk than stationary oscillations around a constant mean. We are skeptical that building a variable output-capital ratio into our model would be fruitful. The mean and variance of its drift is small relative to the other drifts in the model, so we believe that ignoring it does not substantially affect our results.

An important limitation of our model is that aggregate demand disturbances are synonymous with transitory disturbances. Purely transitory aggregate supply and technological disturbances will be misclassified as aggregate demand disturbances. If aggregate demand disturbances have a long-run impact on capacity, they will be misclassified as labor supply and technological disturbances. We would be reluctant to apply this technique to European countries that appear to display hysteresis in unemployment (Blanchard and Summers, 1986). For postwar U.S. data, there is a stronger case for stationarity of the unemployment rate.[30]

We now turn to the limitations of the technique. For many VAR exercises, the degree of differencing and cointegration of the data is not a crucial issue. The researcher can estimate the model in levels and let the VAR estimate unit roots if it chooses. Inference issues can be subtle, but many of the usual inference procedures are asymptotically valid even in the presence of unit roots and cointegration. Identification procedures, such as ours, that rely on the long-run multipliers depend critically on the location of unit roots. So, for example, we have already seen how the results can change when the assumption that hours are difference stationary is changed. In addition, our assumption that inflation contains a unit root is not innocuous. We have estimated a modification of our five variable system replacing $\pi_t = (1 - L)p_t$ with $(1 - \lambda L)p_t$ where λ is estimated by maximum likelihood. We find that values of λ greater than 0.9 provide local

29. See Shapiro (1987) for further discussion and evidence that the Solow residual is a good measure of technological innovations despite the potential presence of cyclical errors.

30. Unemployment is the only series for which Nelson and Plosser (1982) reject the null hypothesis of a unit root.

maxima of the likelihood function and that results similar to those reported in the paper follow from this model. There is another local maxima of the likelihood function of comparable size near $\lambda = 0$. Those estimates yield results somewhat different from those reported in the paper. We believe those results are unreliable. They are based on autoregressive models with roots near unity, and consequently the long-run multipliers, upon which our identification rests, are close to being undefined.

Table 3 DECOMPOSITIONS OF VARIANCE: DETERMINISTIC TREND IN HOURS

	Fraction of hours explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	0.0 (1.9)	61.1 (24.9)	38.9 (24.9)
4	0.2 (2.0)	51.4 (23.1)	48.3 (22.9)
8	6.4 (6.3)	41.5 (19.8)	52.1 (19.7)
12	8.4 (7.7)	40.2 (19.2)	51.4 (19.1)
20	8.6 (7.9)	39.9 (18.9)	51.5 (19.0)
36	8.7 (7.9)	39.7 (18.8)	51.6 (19.0)
∞	1.9 (9.4)	40.2 (29.6)	57.8 (29.4)
	Fraction of output explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	0.4 (2.5)	0.8 (12.0)	98.8 (12.2)
4	2.2 (3.5)	3.9 (11.2)	93.9 (11.7)
8	14.2 (8.6)	5.2 (10.0)	80.7 (11.7)
12	16.4 (9.3)	10.4 (11.2)	73.2 (12.3)
20	15.2 (8.7)	26.4 (13.1)	58.5 (13.0)
36	11.0 (7.7)	51.7 (12.5)	37.3 (11.3)
∞	2.0 (8.9)	98.0 (8.9)	0.0
	Fraction of price level explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	0.3 (2.3)	4.9 (16.9)	94.7 (16.9)
4	6.6 (5.6)	6.0 (16.1)	87.4 (16.4)
8	8.3 (7.4)	3.3 (14.0)	88.4 (15.1)
12	7.5 (8.1)	2.1 (13.5)	90.5 (15.0)
20	7.9 (9.6)	0.6 (12.8)	91.5 (15.0)
36	9.5 (11.4)	0.6 (13.1)	89.9 (15.7)
∞	13.9 (14.8)	6.6 (16.0)	79.5 (19.8)

Note: See text for details of these computations. Standard errors are in parentheses. Aggregate demand includes the shock to hours worked.

Finally, VARs do not eliminate omitted variable bias. It is critical in all structural VAR exercises that the VAR forecast errors span the space of structural disturbances. Except in unusual circumstances, the number of variables in the VAR must be at least as large as the number of structural disturbances driving the variables. Hence, the statistical model must be based on an underlying economic model that takes into account the major shocks impinging on the aggregate economy.

Table 3 DECOMPOSITIONS OF VARIANCE: DETERMINISTIC TREND IN HOURS (CONTINUED)

	Fraction of inflation explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	0.3 (2.3)	4.9 (16.9)	94.7 (16.9)
4	11.0 (5.5)	4.8 (12.7)	84.2 (13.1)
8	11.8 (5.6)	4.1 (10.6)	84.1 (11.2)
12	11.1 (6.2)	3.4 (10.2)	85.5 (11.1)
20	10.8 (7.3)	3.2 (10.1)	86.0 (11.5)
36	11.5 (9.1)	4.0 (11.2)	84.5 (13.2)
∞	13.9 (14.8)	6.6 (16.0)	79.5 (19.8)

	Fraction of real interest rate explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	0.3 (2.3)	4.9 (16.9)	94.7 (16.9)
4	3.6 (3.9)	4.6 (12.5)	91.9 (12.3)
8	6.5 (4.3)	7.4 (10.3)	86.0 (10.2)
12	6.0 (4.2)	7.4 (10.3)	86.6 (10.2)
20	5.5 (4.3)	7.9 (10.7)	86.6 (10.5)
36	5.2 (4.4)	8.3 (11.3)	86.4 (10.9)
∞	15.0 (4.3)	17.0 (8.0)	68.0 (8.1)

	Fraction of nominal interest rate explained by shock to		
Quarter	Oil	Technology	Aggregate demand
1	8.5 (6.1)	8.1 (10.9)	83.4 (11.7)
4	13.4 (7.6)	15.6 (12.1)	70.9 (12.5)
8	10.5 (8.1)	12.0 (11.6)	77.4 (12.3)
12	11.0 (9.0)	9.6 (10.8)	79.4 (11.9)
20	13.1 (10.0)	6.3 (9.2)	80.6 (11.7)
36	14.4 (11.1)	4.7 (9.8)	80.9 (13.3)
∞	13.9 (14.8)	6.6 (16.0)	79.5 (19.8)

Note: See text for details of these computations. Standard errors are in parentheses. Aggregate demand includes the shock to hours worked.

5. *Conclusions*

We now summarize the main results from our model in which labor supply is allowed to have a stochastic trend. Aggregate demand accounts for between 20 and 30 percent of the variation in output at business cycle horizons. Moreover, it is an important factor in most episodes labelled as recessions in the NBER chronology. Technological change accounts for roughly one-third of output variation. Adverse technological shocks are not an important factor in recessions except for the recession in 1970, which roughly coincides with the beginning of the productivity slowdown. Favorable technological shocks play an important role in explaining strong growth in the 1960s. Additionally, our estimated technological shocks and the observed Solow residual are highly correlated at low frequencies.

Oil price shocks are a key factor in explaining the recessions that followed the two OPEC crises, but are unimportant on average.

The estimates imply that permanent shifts in labor input play a large role in explaining output fluctuations at all frequencies. It is not surprising to find that changes in labor are important in explaining low frequency movements in output. Our estimate that labor supply changes account for one-half the long-run changes in the level of output corresponds closely with the findings of growth accounting research.

Our finding that the permanent shocks in labor account for at least 40 percent of output variation at all horizons is, however, quite surprising. Yet, this finding follows from a simple and widely accepted growth model, together with our specification for the stochastic process followed by hours. We find that changes in hours have a permanent component and that changes in output do not account for much of the cyclical variability of hours. Hence, permanent, autonomous shocks to hours will play an important role at business cycle frequencies.

In order to accommodate the prior belief of many economists—which we share—that changes in labor supply are fairly smooth, we estimate an alternative model where hours worked are stationary about a deterministic trend. Detrending hours is an extreme solution because it implies there is no stochastic component to the trend in labor supply. Our basic model with stochastic trend could have told us that variance of the trend is small. Indeed, had we allowed for unit moving average roots in the estimates, the trend-stationary case is nested in the basic model with stochastic trend. We do not find a unit moving average root in the univariate ARIMA model for hours, and so we believe that explicitly incorporating moving average components into the model would not alter the results.

Despite our belief that the model with stochastic trend in labor is the best econometric specification, we present results with trend-stationary labor because of our prior that labor supply changes smoothly and because of the econometric difficulty in distinguishing between stochastic and deterministic trends. In the model with deterministic labor, aggregate demand is very important in explaining output at business cycle frequencies and has a very persistent effect on output. This result arises because the low frequency, high variance, autonomous movements in labor input are attributed to aggregate demand rather than labor supply. Because taking out a deterministic trend is an overly stringent way of imposing the prior that labor supply shocks are smooth, these estimates provide a loose upper bound on the contribution of aggregate demand to output fluctuations.

The statistical difficulty in distinguishing between the two models should be viewed in proper perspective. The basic model with the stochastic trend in labor supply implies that the permanent components of output account for two-thirds to three-quarters of business cycle frequency variation in output. This finding is similar to those of other researchers. We are surprised that permanent movements in labor input are so important in explaining output fluctuations in the short-run. Yet, we would not want to label these shocks as aggregate demand, as is done effectively in the trend-stationary estimates. The estimated labor supply shocks are autonomous movements in labor input. The estimates take into account Okun's law by purging the estimated labor shock of movements in hours that can be explained by business cycle frequency movements in output and other variables. A theory that would attribute these shocks to aggregate demand must be able to explain why there are large movements in hours that are not explained by movements in output.

Data Appendix

This Appendix discusses the data used in the estimates.

All data are quarterly. The estimates are carried out on data from 1951:1 to 1987:2. All data are seasonally adjusted unless otherwise noted.

Output and the *price* level are measured as the 1982 dollar quantity and the deflator for total gross domestic product, less the gross domestic product of farms, the government, and the housing sector. These data are available in the National Income and Product Accounts. Given that our estimates are based on a model of long-run growth relating measured inputs to measured output, this measure is more appropriate than gross national product. First, this level of aggregation (private domestic nonfarm and nonresidential) matches hours and capital stock data. Second, this aggregation abstracts from the major imputations in the national accounts: output of owner-occupied housing is imputed based on its rental value; output of the government is imputed as its wage bill. Third, farmers are largely self-employed, so there is no meaningful hours data for them. Shocks hitting the farm sector also might be very different from shocks to the non-farm sector.

The *hours* data are hours of all persons in the nonfarm business sector. This index is published by the Bureau of Labor Statistics as part of its productivity data.

The *labor force* is defined as the civilian labor force minus agricultural and civilian government employment. These data are also published by the BLS.

The *interest rate* data are average of monthly data for three-month Treasury bills on the secondary market.

The *oil price* series is the producer price index for crude oil (PW561, not seasonally adjusted) deflated by our general price index.

Computation of a quarterly *Solow residual* is complicated by the unavailability of quarterly compensation and capital stock data. Hence, our procedure necessarily involved some interpolation.

The formula for the Solow residual is

$$\Delta\varepsilon_t^S = \Delta y_t - s_t^H \Delta h_t - (1 - s_t^H) \Delta k_t \quad \text{(A.1)}$$

where s_t^H is the share of labor compensation in nominal output and Δy_t, Δh_t, and Δk_t are growth in output, labor, and capital. The capital stock is the beginning-of-period stock. Output and labor are measured as above. The net capital stock on a constant dollar basis for nonfarm business is available on an end-of-year, not end-of-quarter basis (see August 1987 *Survey of Current Business*, for example). We calculate the quarter-to-quarter changes in the capital stock by using the quarterly gross investment series (gross private domestic nonfarm fixed investment) from the NIPA. We know the net change in the capital stock over the year from the annual capital stock data. We use this information to convert the gross flows to net flows by assuming that the ratio of gross to net investment is the same within each quarter of a given year.

The compensation for nonfarm private business employees is also only available annually (Table 6.4 of the NIPA). We add to employee compensation of proprietor's income (net of depreciation) to arrive at the annual estimate of s_t^H. The quarterly figure is then defined as a weighted average of the previous years and the current years share. The weights for the first quarter are ¾ on the previous year and ¼ on the current year; for the second quarter are ½ and ½; for the third ¼ and ¾; and 0 and 1 for the fourth quarter. This procedure approximates the standard Divisia index approximation, which is, in annual data, to take a moving average of the current and lagged year's data as an estimate of the current share.

We are grateful for extensive discussions with Olivier Blanchard and Robert Hall and for many helpful comments from Joseph Altonji, Francis Diebold, Robert Engle, Stanley Fischer, Robert Gordon, James Hamilton, Danny Quah, Christopher Sims, James Stock, and participants in the Conference and a seminar at Yale. We gratefully acknowledge financial support from NSF Grants Nos. SES-8521791 (Shapiro) and SES-8510289 and SES-8618984 (Watson).

REFERENCES

Bernanke, B. S. 1986. Alternative explanations of the money-income correlation. *Carnegie-Rochester Conference Series on Public Policy* 25(Autumn): 49–99.

Blanchard, O. J. 1986. Empirical structural evidence of wages, prices, and employment in the US. NBER Working Paper no. 2044, Cambridge, MA.

Blanchard, O. J. and L. H. Summers. 1986. Hysteresis and the European unemployment problem. *NBER Macroeconomics Annual* 1: 15–78.

Blanchard, O. J., and M. W. Watson. 1986. Are all business cycles alike? *American Business Cycle*, ed, R. J. Gordon. Chicago: University of Chicago Press.

Blanchard, O. J. and D. Quah. 1988. The Dynamic effects of aggregate demand and supply disturbances. M.I.T. Mimeo.

Campbell, J. Y. and N. G. Mankiw. 1987. Are output fluctuations transitory? *Quarterly Journal of Economics* 102: 857–80. (a)

Campbell, J. Y. and N. G. Mankiw. 1987. Permanent and transitory components in macroeconomic fluctuations. *American Economic Review Proceedings* 77: 111–17. (b)

Clark, P. K. 1987. The cyclical component of U.S. economic activity. *Quarterly Journal of Economics* 102: 797–814.

Denison, E. F. 1974. *Accounting for United States Economic Growth, 1929–1969.* Washington, D.C.: The Brookings Institution.

Friedman, M. 1968. The role of monetary policy. *American Economic Review* 58: 1–17.

Hall, R. E. 1988. The relation between price and marginal cost in U.S. industry. *Journal of Political Economy*. Forthcoming.

Hamilton, J. 1983. Oil and the macroeconomy since World War II, *Journal of Political Economy* 91: 228–48.

Harvey, A. C. 1985. Trends and cycles in macroeconomic time series. *Journal of Business and Economic Statistics* 3: 216–27.

King, R., C. Plosser, J. Stock, and M. Watson. 1987. Stochastic trends and macroeconomic fluctuations. University of Rochester. Mimeo.

Kydland, F. E., and E. C. Prescott. 1982. Time to build and aggregate fluctuations. *Econometrica* 50: 1345–1370.

Long, J. B., and C. I. Plosser. 1983. Real business cycles. *Journal of Political Economy* 91: 39–69.

Lucas, R. E. 1973. Some international evidence of output-inflation tradeoffs. *American Economic Review* 63: 326–34.

Nelson, C. R., and C. I. Plosser. 1982. Trends and random walks in macroeconomic time series: Some evidence and implications. *Journal of Monetary Economics* 10: 139–69.

Prescott, E. C. 1986. Theory ahead of business-cycle measurement. *Carnegie-Rochester Conference Series on Public Policy* 25(Autumn): 11–44.

Romer, P. M. 1987. Crazy explanations for the productivity slowdown. *NBER Macroeconomic Annual* 2: 163–202.

Schwert, G. W. 1987. Tests for unit roots: a Monte Carlo investigation. University of Rochester. Mimeo.

Shapiro, M. D. 1987. Are cyclical fluctuations in productivity due more to supply shocks or demand shocks? *American Economic Review Proceedings* 77: 118–124.
Sims, C. A. 1986. Are policy models usable for policy analysis? *Quarterly Review* Federal Reserve Bank of Minneapolis (Winter) 2–16.
Solow, R. M. 1956. A Contribution to the theory of economic growth. *Quarterly Journal of Economics* 70: 65–94.
Solow, R. M. 1957. Technical change and the aggregate production function. *Review of Economics and Statistics* 39: 312–20.
Stock, J. H., and M. W. Watson. 1987. Testing for common trends. Northwestern University. Mimeo.
Swan, T. W. 1956. Economic growth and capital accumulation. *Economic Record* 32: 334–61.
Tobin, J. 1955. A dynamic aggregative model. *Journal of Political Economy* 63: 103–15.
Watson, M. W. 1986. Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18: 49–75.

Comment

ROBERT E. HALL
Stanford University/NBER

What makes macroeconomics different from microeconomics? Macro is interested in the fundamental sources of economic fluctuations, whereas micro almost always considers movements of one actor or one market as the result of changes elsewhere in the economy. In this paper, Shapiro and Watson tackle the issue of driving forces head on. Their menu of alternative sources of fluctuations contains aggregate demand shocks, shifts of technology, movements of oil prices, and shifts of labor supply. The big surprise in the paper is that labor supply shocks are important not only in the longer-run movements of the economy, but also in the shorter-run business cycle. Neither of the major schools of macroeconomic thinking active today—real business cycles or modern Keynesianism—puts any weight on labor supply as a driving force. Taken at face value, Shapiro and Watson's results call for a major rethinking of macroeconomics.

All attempts to measure fundamental driving forces must rest on strong assumptions about identification. Following Blanchard and Quah, Shapiro and Watson use timing properties to achieve identification. Their setup requires that all long-run effects on output come either from labor supply or from oil. They claim that fairly general theoretical considerations support this identifying assumption. Certainly it can be true of the Solow growth model, where labor supply is the dominant determinant of output in the long-run. What is surprising about Shapiro and Watson's findings, however, is that shifts in labor supply are an important determinant of output in business cycle frequencies.