Center for Research on Economic and Social Theory
and
Department of Economics
Working Paper Series

The Causes of Business Cycles and the Cyclicality of Real Wages

Charles A. Fleishman

August, 1993
Number 93-21

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109-1220

The Causes of Business Cycles and the Cyclicality of Real Wages

Charles A. Fleischman
University of Michigan

August 14, 1993

Special thanks to Matthew Shapiro, Robert Barsky, Susanto Basu, Charlie Brown, Carol Evans, Jonathan Parker, and the participants in the National Bureau of Economic Research's High Frequency Economic Fluctuations Summer Institute for extremely helpful suggestions and conversations, and the Benard L. Maas Research Foundation for financial support. Any errors are the responsibility of the author.

Abstract

This paper estimates the cyclicality of real wages using a VAR approach. Long-run restrictions on the behavior of aggregate hours and output identify labor supply, technology, oil price, and aggregate demand shocks. It is shown that real wages are procyclical in response to technology and oil price shocks but are countercyclical in response to aggregate demand shocks. The evidence is consistent with models where nominal wages are stickier than nominal prices. The results point out the importance of looking at the cyclicality of real wages in response to the shocks that drive business cycles, rather than at the simple correlation between real wages and output or the unemployment rate.

I. INTRODUCTION

There has been a great deal of recent interest in understanding the cyclicality of real wages. Bils (1985) and Solon, Barsky, and Parker (SBP, 1992) found real wages to be significantly procyclical, overturning many years of macroeconomic wisdom. Neftci (1978) and Sargent (1978) had earlier found real wages to be countercyclical, while Geary and Kennan (1982) found real wages to be acyclical.

A model's ability to explain procyclical movements in real wages has become a benchmark by which macroeconomists judge business cycle theories. Keynesian models with sticky nominal wages have been criticized and dismissed because they predict countercyclical real wages. McCallum (1986) summarizes the criticisms of models with sticky nominal wages. He writes: "If wage stickiness alone were responsible for the real effects of monetary actions, with product prices adjusting flexibly, then we should observe countercyclical movements in the real wage. That we do not has recently been reconfirmed by Bils (1985).[1]"

Real Business Cycle (RBC) models (Kydland and Prescott, 1990; and others) explain procyclical real wages within their explicitly market clearing framework. Technology shocks move output and real wages, which are equal to the marginal product of labor, in the same direction. These models, however, have been criticized for, among other things, predicting real wage movements that are much more procyclical than those observed in aggregate data. SBP showed that countercyclical composition bias in aggregate real wage series causes studies using aggregate data to understate the true procyclicality of real wages. Despite their consistency with the stylized fact of real wage cyclicality, RBC models are still controversial because they allow no channel through which nominal variables can affect real output (King, Plosser, Stock, and Watson, 1991).

Recent New Keynesian models that attribute movements in real output to the effects of aggregate demand shocks are also consistent with procyclical real wages. There are several varieties of models that have been classified under the broad New Keynesian heading. Models

[1]McCallum (1986), p. 428. An exception to this is the knife-edged case in Taylor (1979) where prices are a constant markup over nominal wages, and real wages are acyclical.

of imperfect competition with nominal prices stickier than nominal wages (Mankiw, 1985; and many others) predict procyclical movements in real wages in response to aggregate demand shocks. In direct response to the findings of procyclical real wages, Rotemberg and Woodford (1991) offer an alternative model that does not rely on nominal rigidities. They explain the procyclical response of real wages to aggregate demand shocks as the result of imperfect competition and countercyclical markups of prices over marginal costs, of which nominal wages are a significant component.

The empirical studies and theoretical models cited above constrain real wages to exhibit the same cyclicality, independent of the causes of the fluctuations in output. Blanchard and Watson (1986) identify several types of shocks that cause business cycles. They also show that the cyclical covariance of real wages varies across post-war business cycles, suggesting that real wage cyclicality may differ depending on the causes of business cycles.

The Fischer (1977) model of long-term nominal wage contracts allows both aggregate supply and demand shocks to affect output. In this model, a positive aggregate demand shock (monetary shock) causes real wages to decline, allowing firms to hire more labor and increase output. The Fischer model has been criticized for predicting countercyclical real wages. The model, however, also predicts procyclical movements in real wages in response to aggregate supply shocks. Bils and SBP studied periods with several important supply shocks, including two OPEC oil price shocks and the onset of the "productivity slowdown." This might explain the overall procyclical pattern of real wages found in these two studies, even if long-term nominal wage contracts are important.

This paper examines the importance of studying the cyclicality of real wages in response to the various shocks that affect the economy. It builds upon an existing literature that has identified aggregate supply and demand shocks within vector autoregression models of the economy. Blanchard and Quah (BQ, 1989) and Shapiro and Watson (SW, 1988) identify aggregate supply shocks as those that can have permanent effects on real output, and aggregate demand shocks as those that can have only transitory effects on real output. This methodology is derived from the neoclassical synthesis, which posits that output movements can be divided into their long-run growth components and short-run cyclical components. By allowing the factors that move the economy in the long run to contribute to variations in output at all frequencies, BQ and SW take a more agnostic view of business cycles.

SW build transitory dynamics resulting from aggregate demand shocks onto a simple real business cycle model. I extend the SW model to explicitly incorporate real wage dynamics. The principal aim of this paper is to include the real wage in the model without adding an additional structural shock. SW identify five shocks that explain all of the variation in output at all horizons. Any additional shock must have no effect on output (or the other variables in the SW model) at any horizon. This shock must also have only transitory effects on the real wage. Following SW and many earlier balanced growth models, I use a constant returns to scale Cobb-Douglas production function for long-run aggregate output. With the Cobb-Douglas production function, there is nothing that has permanent effects on real wages but not on output. In the long run, output is determined by the level of labor supply, the level of technology, and real oil prices, which, in turn, are determined by the accumulation of past labor supply shocks, technology shocks, and oil price shocks. I will show that in the long run, the real wage is determined by the level of technology and real oil prices.

I use this variant of the SW model to study the cyclicality of real wages in response to supply and demand shocks. I examine the impulse response functions of output, real wages, and average labor productivity in response to each of the three supply shocks. In addition, I use the forecast errors in output and real wages due to each of the shocks to construct two measures of their cyclical covariance. The first is the simple correlations between output and real wages in response to the shocks. The second is the output elasticity of real wages in response to the shocks. This second measure allows comparison with other studies of aggregate real wage cyclicality.

The results are generally consistent with Keynesian models of aggregate supply with nominal wages stickier than nominal prices, and with simple models of labor supply and demand. Labor supply shocks cause roughly acyclical movements in real wages in the short run. Technology shocks and oil price shocks result in very procyclical real wages. Aggregate demand shocks move real wages countercyclically.

These findings are consistent with other studies that have examined the cyclicality of real wages in response to different shocks affecting the economy. Liederman (1982) shows that Barro-style unanticipated movements in money supply result in countercyclical real wages.

Cogley (1993), using a VAR methodology to test the neutrality of money in the short and medium run, finds that nominal wages are stickier than nominal prices, resulting in countercyclical movements in real wages in response to monetary shocks. Blanchard (1989), using a traditional VAR with contemporaneous rather than long-run restrictions, finds slightly countercyclical real wages in response to aggregate demand shocks. Sumner and Silver (1989) find that real wages move either procyclically or countercyclically, depending on the choice of sample period. They find that in years in which there is a positive correlation between change in the inflation rate and change in the unemployment rate (countercyclical inflation), which they interpret as years dominated by aggregate supply shocks, the real wage is procyclical. In contrast, they find that in years in which the change in the inflation rate and the unemployment rate are negatively correlated, the real wage is countercyclical.

Following this introduction, Section II introduces the model I use to study the cyclicality of real wages. Section III describes the data and addresses data related issues that determine the proper specification of the model. Section IV explains the estimation strategy. Section V discusses the results, and Section VI concludes.

II. THE MODEL

This paper develops a close variant of the model presented in Shapiro and Watson (1988). First, I describe the basic five equation SW model. Next, I explain how I append a real wage equation to the basic model by including a non-allocative measurement error shock.

The Basic Shapiro and Watson Model

SW used their model to understand the causes of business cycles. Their most important identifying assumption is the decomposition of output into permanent and transitory components. The SW model is based on a simple balanced growth model of the economy. In the long run, output is determined by the size of the labor force (labor supply) and the state of technology. The long-run capital-to-output ratio is assumed constant. In the short run, variables are allowed to differ from their long-run levels.

In the SW model, long-run labor supply and technology evolve according to:

$$(1) \qquad h_t^* = \delta_h + h_{t-1}^* + \theta_h(L)v_t$$

and

$$(2) \qquad \epsilon_t^* = \delta_\epsilon + \epsilon_{t-1}^* + \theta_\epsilon(L)e_t$$

where h^*_t is long-run labor supply and ε^*_t is the long-run level of technology.[2] v_t is the labor supply innovation and e_t is the innovation in technology. The innovations are mutually and serially uncorrelated. $\Theta_h(L)$ and $\Theta_\epsilon(L)$ have strictly transitory dynamics, allowing the long-run levels of technology and labor supply to exhibit more intricate dynamics than a simple random walk. Long-run labor supply is exogenous. The position of the perfectly inelastic labor supply curve is determined solely by the accumulation of previous labor supply shocks, such as growth in population or changes in labor force participation rates.

The long-run level of output is determined according to the Cobb-Douglas production function:

$$(3) \qquad y_t^* = \alpha h_t^* + (1-\alpha)k_t^* + \epsilon_t^*$$

Applying the implication of balanced growth models that the capital-to-output ratio is constant ($y^*_t = \eta + k^*_t$), and suppressing constants, the production function can be rewritten as:

$$(4) \qquad y_t^* = h_t^* + (\frac{1}{\alpha})\epsilon_t^*$$

Equation (4) shows the restrictions that only labor supply and technology shocks can have permanent effects on output. Given the production function (3) and the constancy of the capital-to-output ratio, I can solve for the labor demand curve:

$$(5) \qquad w_t^* = (\alpha - 1)h_t^* + (1 - \alpha)y_t^* + \epsilon_t^*$$

[2]All variables are natural logarithms, and all constants are subsequently suppressed.

Using (4), I can eliminate output from (5) to find:

$$(6) \qquad w_t^* = (\frac{1}{\alpha})\epsilon_t^* \qquad = y_t^* - h_t^*$$

In the long run, the real wage will be a function only of the accumulation of technology shocks and will be equal to the marginal product of labor ($y^*_t - h^*_t$), holding the long-run capital-to-output ratio constant.

Along a balanced growth path with constant returns to scale technology, the capital stock evolves according to:

$$(7) \qquad \Delta k_t^* = \Delta h_t^* + (\frac{1}{\alpha})\Delta\epsilon_t^*$$

A labor supply shock has a long-run effect on the labor force but no effect on the level of technology. In response to a labor supply shock, the change in the capital stock is equal to the change in the labor force, leaving the capital-to-labor ratio unchanged. The long-run marginal product of labor, and hence the long-run real wage, is a function of the capital-to-labor ratio. With the capital-to-labor ratio unchanged in the long run, a labor supply shock can have no permanent effect on the real wage. Conversely, the long-run real wage responds to a technology shock because the long-run capital stock adjusts in response to a technology shock but the long-run labor force does not. The change in the long-run capital-to-labor ratio causes a permanent change in the real wage.

Equations (1), (2), (3) and (6) define a very simple real business cycle model that attributes all movements in hours, output, and wages to aggregate supply shocks and permits no deviations from the long-run values. The SW model allows for strictly transitory movements in output and hours by including two aggregate demand shocks, which are loosely interpreted as IS and LM shocks. One additional variable must be included in the model for each shock. In order to identify two demand shocks, Shapiro and Watson included the ex poste real interest rate and the inflation rate.

SW place no restrictions on the short-run dynamics of any of the variables in the system.

Movements in actual hours and output will be a combination of the long-run dynamics from (1) and (2) and strictly transitory dynamics. Transitory movements in hours and output can be explained by the effects of aggregate demand shocks, which are neutral in the long run, and by off-the-labor supply curve and off-the-production function behavior that can result from the adjustment to either aggregate supply or aggregate demand shocks.

Equation (1) shows that hours contain a unit root. Together (1) and (2) imply that output also contains a unit root that can be attributed to both labor supply and technology. In order to estimate the model, actual hours and output must be differenced. The first-differenced equations are:

$$(8) \qquad \Delta h_t = \theta_h(v_t) + (1 - L)\Xi_h(L)(v_t, e_t, d_t^1, d_t^2)$$

and

$$(9) \qquad \Delta y_t = \theta_h(v_t) + \frac{1}{\alpha}\theta_e(e_t) + (1 - L)\Xi_y(L)(v_t, e_t, d_t^1, d_t^2)$$

where d^1_t and d^2_t are the two aggregate demand shocks. These shocks are serially uncorrelated and uncorrelated with the aggregate supply shocks. Equations (8) and (9) include the restrictions that aggregate demand shocks have only transitory effects on hours and output, and show that the effects of the permanent shocks can have strictly transitory components, as well.

Equations (8) and (9) also embody the cross-equation restriction implied by (4): Labor supply shocks have the same long-run impacts on hours and output. From (8), all long-run movements in labor supply will be due to $\Theta_h(L)v_t$. Similarly, (9) shows that long-run movements in output not due to movements in technology (captured by $\Theta_e(L)e_t$) are due to $\Theta_h(L)v_t$. Equation (6) shows that this also implies that labor supply shocks have no long-run effect on the real wage, because they have no long-run effect on the output-to-labor ratio, which is equal to the long-run capital-to-labor ratio.

Appending Real Wage Dynamics to the Basic Shapiro and Watson Model

To study the dynamic properties of the real wage, I include an additional equation for the

real wage and an additional shock. Were I to include an equation for the real wage without including an additional shock, I would be making the strong overidentifying assumption that all of the variance in the real wage is explained by the shocks already included in the model. Clearly, this overidentifying restriction would be rejected by the data.[3]

The additional real wage shock must satisfy several conditions. First, it can have no effect on output or any of the other variables included in the SW model at any horizons. SW identify five structural shocks that span the space of shocks that can affect the variables included in the basic model. Second, in the long run, the real wage is determined by the level of technology. The additional shock can have only transitory effects on the real wage.

The real wage shock will represent the non-allocative movements (or non-movements) in real wages that do not match movements in productivity and that are not explained by the five structural shocks. The presence of long-term labor contracts, both formal and implicit, implies that there are non-allocative movements (or non-movements) in real wages. These contracts, however, probably cannot explain all of the movements in real wages not explained by the structural shocks. Other possible interpretations for the real wage shock, like changes in the rates of unionization or in the markup ratio, would imply an affect on employment and, hence, output.

In the absence of economically meaningful candidates for this type of shock, I follow Altug (1989) in adding a "measurement error" shock to the real wage equation. Altug confronts a similar problem in estimating the dynamics of a real business cycle model driven by a single technology shock. Recognizing that the time series of the actual variables will differ somewhat from the dynamics implied by the model, Altug includes a vector of measurement errors, one for each variable. Here, following Altug, I allow the real wage to differ from the real wage implied by the effects of the five structural shocks by including a "measurement error" shock. The actual real wage will behave according to:

[3]Indirectly, I will perform this test and show that the overidentifying restriction is rejected. I estimate the unrestricted model and find that approximately 40 percent of the one-quarter ahead forecast error in real wages cannot be explained by the five structural shocks. Smaller, though statistically significant, fractions of the forecast error in real wages at longer horizons are also unexplained by the structural shocks.

$$(10) \quad w_t = w_t^* + \Gamma_w(L)(v_t, e_t, d_t^1, d_t^2) + \Lambda_w(L)(m_t)$$

where m_t is the measurement error shock. These shocks are serially uncorrelated and uncorrelated with the structural shocks.

Equation (6) shows that there is a cointegrating relationship between the real wage and the output-to-labor ratio, which is a linear transformation of the marginal product of labor. Equations (6), (8), and (9) also show that labor supply and aggregate demand shocks have only transitory effects on the real wage and average labor productivity. To identify the "measurement error" shock, I study the difference between the real wage and the average productivity of labor:

$$(11) \quad w_t - (y_t - h_t) = \Gamma_w(v_t, e_t, d_t^1, d_t^2) + \Lambda_w(L)(m_t)$$

Equations (8), (9), and (11) describe the dynamics of the model specified by equations (1) through (3) by imposing the long-run restrictions on hours, output and wages. The dynamics of the other two variables in the system, the inflation rate and the real interest rate, are not restricted by the model.

The Role of Oil Prices

Even after including the dynamics of the inflation rate and the real interest rate, the model described above is still incomplete. During the 1970s and 1980s, there were several oil shocks. Of particular note are the first OPEC shock in 1973-74, the second OPEC shock in 1979 following the fall of the Shah of Iran, and the collapse of oil prices during 1985-86. In addition, it appears that the process for real (or nominal) oil prices has changed from one characterized by infrequent though generally large discrete changes in nominal oil prices that persisted through the early 1980s, to one more closely resembling the volatility of other commodity price series since the early 1980s.

Hamilton (1983) shows that oil price shocks were an important cause of recessions even

in the pre-OPEC era.[4] SW add real oil prices to the model and treat them as exogenous. This assumption is equivalent to stating that the level of aggregate demand in the U.S. does not affect real oil prices. Because the real oil price is exogenous, SW write the oil price shock as the change in the real price of oil.

I can now complete the description of the restrictions of the SW model. SW treat oil shocks as supply shocks that have permanent effects on output, but no permanent effects on hours.[5] With the capital-to-output ratio assumed constant in the long run, this implies that oil shocks have permanent effects on real wages because of their effects on the capital-to-labor ratio.[6] An increase in real oil prices will lower the long-run capital-to-labor ratio, causing a permanent decline in real wages.

III. DATA

This section discusses the data used in this paper. First, it describes the data series I use. Next, it discusses the time series properties of these data series. I focus on unit root tests, as

---

[4]Mork (1989) extends Hamilton's analysis and finds that oil price shocks have asymmetric effects on output. He finds that the large oil price decrease in 1986 did not cause a significant increase in output.

[5]I allow real oil prices to have permanent effects on output thus treating real oil prices as I do other aggregate supply shocks. Arguments can be made that real oil price shocks are really aggregate demand shocks, affecting output through wealth effects and not through the production function. If an increase in oil prices lowers the value of existing capital by making it more expensive to operate, owners of capital will be worse off and will adjust their lifetime consumption plans. Thus, even if the effects of oil price shocks operate through aggregate demand channels, their effect will be permanent. Because of the separation of shocks into those with permanent and strictly transitory effects on output, it is appropriate to treat oil price shocks as aggregate supply shocks in this model.

[6]Real oil price shocks should not have permanent effects on output when output is measured as value added, as implied by the production function (3) and the choice of data (real gross domestic product). If a value-added production function exists then the long-run capital-to-output ratio cannot be affected by real oil prices and real oil prices can have no permanent effect on output or real wages. By leaving the long-run effects of real oil prices on output and real wages unrestricted, I allow the data to speak to this issue.

well as prior beliefs, in determining the appropriate degree of differencing of the data. Finally, it analyzes the properties of the real oil price series that can affect the choice of sample period.

Data Sources

All data series are quarterly and seasonally adjusted, unless otherwise indicated. The choice of data series was dictated by the goal of most accurately describing the dynamics of the domestic U.S. economy.

Output is measured by the gross domestic product (GDP) of the non-farm, non-housing private sector.[7] Nominal GDP is deflated by the appropriate GDP deflator. The hours data are hours of all persons in the non-farm business sector from the Bureau of Labor Statistics Productivity and Costs Document. The hours data includes the housing sector, but this should have little effect on the results, as the housing sector employs relatively few people. Nominal oil prices are measured as the producer price index for crude oil and are not seasonally adjusted. The nominal interest rate is the market rate on U.S. three month treasury bills.

The wage series is constructed as follows: I begin with total nominal wages and salaries paid in the private sector (excluding government and government enterprises) from the National Income and Product Accounts. I deflate this by the GDP implicit price deflator for the non-farm, non-housing private sector. Finally, I divide by hours of all persons in the non-farm business sector to construct a measure of real hourly wages. I chose this wage series over the more frequently used series of wages of production and non-supervisory employees in the manufacturing sector because my series provides broader coverage and is more comparable with the output and hours series I use.[8] This wage series does not exclude farm and housing sector workers. Again, because these sectors contribute little to the aggregate wage bill, the results should not be sensitive to their inclusion.

[7]See Shapiro and Watson (1988) for a discussion of the problems associated with using farm and housing sector data.

[8]I have examined the sensitivity of the results to the choice of wage series, and found them to be quantitatively very similar.

Time Series Properties

Estimation of a VAR system like that described above must take into account the time-series properties of the data. In particular, I must consider two factors: The left hand side variables for each equation must be stationary, and all cointegrating relationships must be taken into account.

As discussed above, equation (1) implies that h_t is an integrated process. Equation (4) also implies that y_t has two permanent components, one due to growth in the labor force (v_t) and the other due to growth in technology (e_t). Labor supply shocks can have no long-run effects on the output-to-labor ratio. I take this cointegrating-like relationship into account by substituting $(y_t - h_t)$ for y_t when estimating the model. $(y_t - h_t)$ still has a permanent component due to growth in technology. The results of augmented Dickey-Fuller tests using six lags shown in Table 1 indicate that $(y_t - h_t)$ still contains a unit root. This series must be first differenced.

There is no similar theoretical guidance for determining the appropriate differencing for the inflation rate and real interest rate series. The augmented Dickey-Fuller t-statistic for the presence of a unit root in the inflation rate is -2.64. This is very close to the 10 percent critical value of -2.57 for the test including a constant. The hypothesis of a unit root can not be rejected at the 5 percent significance level, and can barely be rejected at the 10 percent level.

I chose to difference the inflation rate. This allows aggregate supply and demand shocks to have permanent effects on the inflation rate. Barsky (1987) shows that inflation has become much more persistent in the post-war period. Shapiro (1993) analyzes the times series properties of the inflation rate measured by change in the Consumer Price Index (CPI). Shapiro finds that the results of unit root tests are sensitive to the choice of sample period. The longer the period is extended past the Volcker disinflation of the early 1980s, the more mean-reverting inflation looks. One argument for this is that Volcker was successful in using monetary policy to eliminate much of the inflation that persisted through the 1970s. Given the low inflation rates during the 1980s, I find this somewhat persuasive evidence that Volcker's disinflation has had permanent, or at least very persistent, effects. Similarly, the high inflation of the 1970s resulting in part from the two OPEC oil shocks was quite persistent. These experiences suggest that inflation is quite persistent, and that aggregate supply and demand shocks should be allowed to have permanent or very persistent effects on its level.

In choosing the appropriate differencing for a series it is important to consider the context in which it is being used (Blough 1991). For the purposes here, it is more appropriate to allow the aggregate demand and supply shocks to have permanent effects on the level of inflation and let the data determine their importance. I examine the sensitivities of the main findings of this paper to treating inflation as a stationary process and find that they are robust to this modification.

The augmented Dickey-Fuller t-statistic for the real interest rate is -3.08. The hypothesis of a unit root can be rejected at better than the five percent significance level (critical value of -2.89) but not at the 10 percent level. The Solow model predicts a constant real interest rate at all steady states, assuming population growth is constant. The stationarity of the change in hours is equivalent to saying that population growth rates are mean reverting. Because of the inconclusive results from the Dickey-Fuller test and the view that real interest rates are constant in the long run, I treat the real interest rate as stationary.

Equation (6) clearly indicates that real wages are integrated. Equation (6) also shows that there should be a cointegrating relationship between w_t and (y_t-h_t) with a cointegrating vector (1, -1). The augmented Dickey-Fuller t-statistic for the $w_t-(y_t-h_t)$ is -3.20. This is slightly larger than the 10 percent critical value of -3.08 for the test including a constant and a time trend. Although the hypothesis of a unit root cannot be rejected at the 10 percent significance level, I chose to treat this series as trend stationary.

Permanent movements in the difference between the real wage and marginal product of labor are difficult to explain. Even within models with sticky wages and long-term implicit labor contracts, real wages adjust to clear the labor market in the long run. Further, the "measurement error" described above, which allows long-term implicit contracts to be a possible source of this difference, is inconsistent with permanent movements in this series due to the presence of a stochastic trend. The series must be detrended for it to be stationary. I estimate the trend to be -.26 percent per quarter (at an annualized rate) with a standard error of .014. The presence of a deterministic trend is also difficult to explain, though the lack of permanent effects of the measurement error on real wages is more consistent with the measurement error story. I find that measurement errors have very persistent effects on real wages, suggesting a possible specification error. The negative trend may be due to the inclusion of the shrinking wage bill of farm workers

in the numerator of the constructed real wage series and the exclusion of the hours of these same workers from the denominator. I examine the sensitivity of the results to the choice of trend-stationarity. Because the measurement error is exogenous to the rest of the model, the choice of difference stationarity vs. trend stationarity affects only the response of the real wage to the five structural shocks. The results are generally robust to this change in specification, although the response of wages to labor supply shocks does differ somewhat.

Sample Period Selection

I estimate the model for two sample periods: 1951:1 to 1985:4 and 1951:1 to 1992:2. The first period is identical to that used by Shapiro and Watson (1988). This sample period has several advantages over the longer period. First, during the 1980s the process for real oil prices seems to have changed from one dominated by long periods of relative constancy interrupted by large fairly discrete increases, to a highly variable process more closely resembling the volatility of other commodity prices sometime during the 1980s. Figure 1 shows real oil prices from 1951:1 to 1992:2. I examine the correlograms of the real oil price series, allowing for breaks in the process between 1979 and 1988. There is no conclusive break point. For example, prior to 1983:1, the first five autocorrelations are 0.96, 0.92, 0.88, 0.82, and 0.76 and from 1983:1 to 1992:2 they are 0.90, 0.76, 0.69, 0.61, and 0.54. Similarly, the autocorrelations prior to 1986:1 are 0.98, 0.96, 0.94, 0.91, and 0.88 and beginning in 1986:1 are 0.37, -0.10, -0.09, -0.14, and -0.22.

In addition to the ambiguous choice of a breaking point, a second issue is that the only large real oil price decline occurred during 1986. Mork (1989) find that oil price shocks have asymmetric effects on output. The large decrease in the price of oil did not cause a substantial increase in output. Including this in the sample period might produce extremely misleading results. Because of this problem, Shapiro and Watson chose to limit their sample period to ending in 1985:4.

I chose to estimate the model for both sample periods. Despite the problems with the longer period, limiting the sample excludes more than six years of available data. I find that results from the longer period are substantially different from those from the shorter period. These results are also difficult to interpret because they appear to be dominated by the fall in oil prices following the collapse of OPEC in 1986. In an attempt to correct for the possibility that this large fall in oil prices affected the economy differently from prior real oil price shocks, I include a variable in the analysis that effectively dummies out the change in oil prices during the first two quarters of 1986. In addition, to deal with the change in the persistence of the oil price series, I include a variable that is equal to zero prior to 1983:1 and is equal to the change in real oil prices beginning in 1983:1. I chose 1983:1 as the breaking point after examining the correlograms described above. After I make these modifications, the results are still quite different from those from the shorter period, again with the differences dominated by the negative oil shocks and increased volatility of oil prices since 1985. I have not examined the sensitivity of the results to changes in the breaking point.

IV. ESTIMATION STRATEGY

Following Shapiro and Watson, I write the model in its Wold moving average form as:

$$(12)\qquad \begin{bmatrix} \Delta h_t \\ \Delta o_t \\ \Delta(y_t - h_t) \\ \Delta \pi_t \\ r_t \\ w_t - (y_t - h_t) \end{bmatrix} = A(L) \begin{bmatrix} v_t \\ \omega_t \\ e_t \\ d^1_t \\ d^2_t \\ m_t \end{bmatrix}$$

The identification of the model is summarized by the matrix of long-run multipliers:

$$(13)\qquad A(1) = \begin{bmatrix} a_{11} & 0 & 0 & 0 & 0 & 0 \\ 0 & a_{22} & 0 & 0 & 0 & 0 \\ 0 & a_{32} & a_{33} & 0 & 0 & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} & 0 & 0 \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} & 0 \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} & a_{66} \end{bmatrix}$$

This paper estimates the model described by (12) and (13). It uses an instrumental variables approach similar to that described by Shapiro and Watson. The first five equations will be estimated as in SW, except that I impose the additional overidentifying equation that labor supply shocks have no permanent effects on real wages or the marginal product of labor. SW test but do not impose this restriction. They find that the hypothesis that labor supply shocks have the same permanent effects on hours and output cannot be rejected. Imposing the restriction that these responses are the same in the long run is equivalent to imposing the restriction that labor supply shocks have no permanent effects on real wages, because they have no permanent effects on the capital-to-labor ratio.

The real wage equation will be estimated as a supplementary equation, though care will be taken to make the wage shock orthogonal to the five structural shocks.

Equation (12) can be summarized as $X_t = A(L)s_t$, where s_t is the vector of structural shocks. $E(s_t s'_t) = \Omega$, a diagonal matrix, and $E(s_t s'_{t-1}) = 0$. The first describes the mutual orthogonality of the shocks; the second forces the shocks to be serially uncorrelated. The vector of structural shocks, s_t, can be directly estimated.

A(L) and s_t will be estimated by inverting the estimable, finite autoregressive form $A(L)^{-1}X(t) = s_t$. Given that A(1) from (13) is lower triangular, it can easily be shown that $A(1)^{-1}$ is also lower triangular. This transformation allows me to impose the long-run restrictions embodied in A(1) on the sum of the estimated autoregressive parameters. This will be demonstrated below.

As described above, I estimate the autoregressive form of the model. The hours equation will be estimated as:

$$(15) \qquad \Delta h_t = \sum_{j=1}^{p} \beta_{hh,j} \Delta h_{t-j} + \sum_{j=0}^{p} \beta_{ho,j} \Delta o_{t-j} + \sum_{j=0}^{p} \beta_{hy,j} \Delta (y-h)_{t-j} + \sum_{j=0}^{p} \beta_{h\pi,j} \Delta \pi_{t-j} + \sum_{j=0}^{p} \beta_{hr,j} r_{t-j} + v_t$$

Equation (14) does not impose the long-run restrictions from $A(1)^{-1}$. The restrictions require that there be no permanent effects on hours from any of the shocks other than the labor-supply shocks. I can impose these restrictions by restricting the sets of coefficients on changes in oil prices, changes in output, changes in the rate of inflation, and the real interest rate to each sum to zero. I implement this restriction by differencing the right hand side variables whose associated shocks must have no permanent effects. I can rewrite equation (14) as[9]:

$$(15) \qquad \Delta h_t = \sum_{j=1}^{p} \beta_{hh,j} \Delta h_{t-j} + \sum_{j=0}^{p-1} \gamma_{ho,j} \Delta^2 o_{t-j} + \sum_{j=0}^{p-1} \gamma_{hy,j} \Delta^2 (y-h)_{t-j} + \sum_{j=0}^{p-1} \gamma_{h\pi,j} \Delta^2 \pi_{t-j} + \sum_{j=0}^{p-1} \gamma_{hr,j} \Delta r_{t-j} + v_t$$

Similarly, I write the output equation as

[9]The following is an example of how double differencing works: Suppose there is a unit sized shock to oil prices in period t, and no other shocks. Using equation (21), $\Delta h_t = \gamma_{ho,0}$, $\Delta h_{t+1} = -\gamma_{ho,0} + \gamma_{ho,1}$, $\Delta h_{t+2} = -\gamma_{ho,1} + \gamma_{ho,2}$, $\Delta h_{t+3} = -\gamma_{ho,2} + \gamma_{ho,3}$, $\Delta h_{t+4} = -\gamma_{ho,3} + \gamma_{ho,4}$, $\Delta h_{t+5} = -\gamma_{ho,4} + \gamma_{ho,5}$, $\Delta h_{t+6} = -\gamma_{ho,5}$. This sums to zero. If the net change in Δh_t is equal to zero, then there is no long-run effect of oil price shocks on the level of h_t.

$$(16) \qquad \Delta (y-h)_t = \sum_{j=0}^{p-1} \gamma_{yh,j} \Delta^2 h_{t-j} + \sum_{j=0}^{p} \beta_{yo,j} \Delta o_{t-j} + \sum_{j=1}^{p} \beta_{yy,j} \Delta (y-h)_{t-j} + \sum_{j=0}^{p-1} \gamma_{y\pi,j} \Delta^2 \pi_{t-j} + \sum_{j=0}^{p-1} \gamma_{yr,j} \Delta r_{t-j} + e_t$$

the inflation rate equation as

$$(17) \qquad \Delta \pi_t = \sum_{j=0}^{p} \beta_{\pi h,j} \Delta h_{t-j} + \sum_{j=0}^{p} \beta_{\pi o,j} \Delta o_{t-j} + \sum_{j=0}^{p} \beta_{\pi y,j} \Delta (y-h)_{t-j} + \sum_{j=1}^{p} \beta_{\pi\pi,j} \Delta \pi_{t-j} + \sum_{j=1}^{p} \beta_{\pi r,j} r_{t-j} + d_t^1$$

and the real interest rate equation as

$$(18) \qquad r_t = \sum_{j=0}^{p} \beta_{rh,j} \Delta h_{t-j} + \sum_{j=0}^{p} \beta_{ro,j} \Delta o_{t-j} + \sum_{j=0}^{p} \beta_{ry,j} \Delta (y-h)_{t-j} + \sum_{j=0}^{p} \beta_{r\pi,j} \Delta \pi_{t-j} + \sum_{j=1}^{p} \beta_{rr,j} r_{t-j} + d_t^2$$

The output equation imposes the restrictions that aggregate demand shocks have no permanent effects on output, and that labor supply shocks have the same long-run impact on output as they do on hours. I impose the restriction that labor supply shocks have no long-run effect on the output-to-labor ratio by double differencing the change in hours.

Unless there are restrictions on the contemporaneous effects of the labor supply shock on output, the change in the inflation rate, or the real interest rate, the contemporaneous values of these variables will be correlated with the residual in the hours equation. Thus equation (15) can not be estimated by OLS. Similarly, equations (16) through (18) cannot be estimated by OLS because of the contemporaneous correlation of the structural shocks (e_t, d^1_t, d^2_t) with the included contemporaneous regressors (except the contemporaneous change in real oil prices, which is exogenous). Following Shapiro and Watson, I will use an instrumental variables procedure that uses lags 1 through p of Δh_t, Δy_t, $\Delta \pi_t$, and r_t, and lags 0 through p of Δo_t as instruments. Because equations (15) through (18) can be ordered as in the text and estimated in sequence, the residuals from the prior equations are also valid instruments and are used. For example, the residual from the hours equation is used as an instrument in the output equation. Including the residuals from the prior equations as instruments orthogonalizes the estimated residuals, making it unnecessary to use the Choleski decomposition to estimate the structural shocks.

However, the two demand shocks estimated by (17) and (18) will not be separately identified. Although the estimation procedure imposes an ordering on the two equations that causes there to be no contemporaneous effects of d^2_t on the change in the inflation rate, I find this contemporaneous restriction implausible. For this reason, following Shapiro and Watson, I avoid analysis of the separate effects of the two demand shocks on output and the real wage.

Estimation of the real wage equation is different. First, recall that neither the real wage "measurement error" shock or the change in real wages entered into equations (14) to (18). I treat the "measurement error" shock as exogenous to the system, and allow it to affect only the relationship between real wages and the average productivity of labor, which is equal to -- excluding constants -- the marginal product of labor in the long run after the capital stock has adjusted. Equation (6) shows that in the long run, real wages are equal to the marginal product of labor. In the short run I allow off-the-labor supply curve behavior by consumers, labor hoarding behavior by firms, and long-term labor contracts where the wage is not allocative in the short run. Thus the labor supply shocks, technology shocks, oil shocks, and demand shocks can have transitory effects on the difference between the real wage and marginal product. The main observation being made is that these five shocks cannot explain all of the variance in the real wage. The dynamics of the real wage minus the marginal product will help us identify the "measurement error" shock, as well as the cyclicality of real wages in response to supply and demand shocks.

I will estimate the following equation for the real wage:

$$(19) \qquad (w-(y-h))_t = \sum_{j=0}^{p} \beta_{wh,j} \Delta h_{t-j} + \sum_{j=0}^{p} \beta_{wo,j} \Delta o_{t-j} + \sum_{j=0}^{p} \beta_{wy,j} \Delta (y-h)_{t-j} + \sum_{j=0}^{p} \beta_{w\pi,j} \Delta \pi_{t-j} + \sum_{j=0}^{p} \beta_{wr,j} r_{t-j} + \sum_{j=1}^{p} \beta_{ww,j} (w-(y-h))_{t-j} + m_t$$

Equation (19) can be estimated by OLS because m_t is by assumption orthogonal to all of the included regressors. By including the contemporaneous values of the left-hand side variables from equations (15) through (18), as well as the full set of regressors and instruments from those equations, as regressors in equation (19), I ensure that the structural residuals and the oil price shock are orthogonal to "measurement error" shock.

V. Results

I estimate the model described by equations (15) through (19) using six lags. Each equation includes a constant. The oil price equation is not estimated, but the series of changes in real oil prices is demeaned. The model is estimated over both sample periods, but the discussion will focus on the shorter sample.

I present the results in three ways. First, I examine the impulse response functions of output, average labor productivity, and real wages in response to each of the three aggregate supply shocks.[10] Second, I analyze the decomposition of forecast error variance for output and real wages to assess the importance of each of the shocks for explaining output and real wage movements. Third, I use the 4-, 8-, and 16-quarter ahead forecast errors in output and the real wage to construct two measures of cyclical covariance. I examine both the simple correlations of the forecast errors in real wages and output and the elasticities of real wages with respect to output in response to each of the shocks.

[10]It is not possible to separately identify the impulse response functions for the aggregate demand shocks because the shocks themselves are not separately identified. Any analysis of the impulse response functions for the two demand shocks would be misleading.

Impulse Response Functions

I use the impulse response functions and k-quarter ahead forecast errors to examine the cyclicality of real wages in response to each of the shocks: If output and wages move in the same direction in response to a shock, I describe this as a procyclical movement in the real wage; conversely, if wages and output move in opposite directions then the real wage is countercyclical.

Figure 2 (panels a through i) shows the impulse response functions of output, the real wage, and average labor productivity to the three supply shocks. Each figure is plotted with one-standard-error bands around the estimated impulse response function. The standard errors were calculated by bootstrapping from the vectors of structural shocks and the measurement errors 500 times, simulating the paths of all variables using the bootstrapped shocks, and re-estimating the model. For all simulations, the actual series of real oil prices is used.

Figures 2a through 2c show the response to a unit sized labor supply shock. As shown earlier, this shock will have a unit sized long-run effect on hours and output, and no permanent effect on real wages or the average productivity of labor. The effect of labor supply shocks on real wages is not precisely estimated. In the short run, this effect is nearly zero. After two years, the effect is negative but a one standard error band around the estimated response includes a wide range of positive and negative values. Given the imprecision of the estimates, it is difficult to conclude that real wages are either procyclical or countercyclical in response to labor supply shocks.

The point estimates of the impulse response functions indicate that real wages adjust more slowly than average labor productivity to a labor supply shock. Labor supply shocks predominantly represent changes in the number of new entrants to the labor force. The slower adjustment of real wages may indicate that the slope of the age-wage relationship is shallower than the slope of the age-productivity relationship following the entry of younger workers into the labor force. If firms can construct compensation schemes where wages grow more slowly than productivity, this will discourage quits and reduce turnover costs.

After the first quarter, labor productivity falls. This may be partially explained by slow adjustment of the capital stock in response to an increase in the labor force. This explanation, however, is difficult to reconcile with the forecastability of the labor supply. Even though the hours series has a large stochastic component, most changes in the labor force are due to chages

in the growth of the working age population that are forecastable many years in advance, giving firms ample time to adjust the capital stock to accomodate new workers.[11] However, factors that affect the labor force participation rate, like permanent changes in tax policy, increases in the availability of disability insurance, and changes in female labor force participation, may be less forecastable. The capital stock may be slow to adjust to these labor supply shocks. Alternatively, the decrease in average labor productivity may reflect changes in the composition of the labor force toward less experienced workers. Compositional changes could be due to either changes in the retirement rates of older, more productive workers or changes in the entrance rate of younger, less experienced workers.

Figures 2d though 2f show the responses of output, real wages, and average labor productivity to oil price shocks. Initially, all three variables increase in response to an increase in real oil prices. This initial increase is small relative to the large decreases in each of the variables that trough after about one year. Real wages and output move closely together, indicating that real wages are procyclical in response to oil price shocks. Rotemberg and Woodford (1992) show that with value-added data this result is a puzzle unless there is imperfect competition and countercyclical markups. They argue that a constant or procyclical markup is not able to explain the large decline in output and real wages in response to oil price shocks. However, as I show below, the responses of output and real wages to aggregate demand shocks are more consistent with sticky wage models than with models with countercyclical markups.

The procyclicality of real wages in the short run in response to oil price shocks is consistent with the more traditional notion of oil price shocks as aggregate supply shocks. If an increase in oil prices shifts in the aggregate supply curve along a stable aggregate demand curve, output should fall and prices should increase. If nominal wages are stickier than nominal prices, the increase in the price level will result in decreases in the real wage. Alternatively, an oil price shock can be characterized as a shock to the labor demand curve. In a market clearing model, a shift in the labor demand curve along a stable labor supply curve moves hours and wages in

[11] If the capital stock adjusts in anticipation of an increase in the labor supply, output should increase prior to the realization of the labor supply shock, as well. The model estimated here does not deal with this possibility.

the same direction.

In the long run, output, real wages, and average labor productivity are all lower. Although this is inconsistent with equation (4), a value-added production function, it might reflect some complementarity between energy and capital usage in production. If energy becomes more expensive, firms may want to economize on its use by substituting labor input for capital and energy inputs. In the short run, capital and labor may not be very substitutable in a given production process. So, even if the long-run effect is that for a given capital stock firms would increase their demand for labor, in the short run firms will find it difficult to make this substitution and unprofitable to utilize all of their capital. Effective labor demand will decrease in the short run. In the long run, firms will use less capital in production but will be able to substitute labor for capital so the capital-to-labor ratio and the marginal product of labor would both be lower. And, since there is no long-run effect of oil prices on hours, output must be lower in the long run. The difference between the short- and medium-run substitutability of capital and labor might explain the sharp drop in output in the short run and the partial recovery just a few quarters later.

Figures 2g and 2i show the responses of output, real wages, and labor productivity to a unit sized technology shock. Over both the short run and the long run, real wages and output move together resulting in procyclical movements in real wages. The long-run effect on wages and output must be the same, and equal to $(1/\alpha)$[12] times the size of the technology shock. In the short run, an increase in technology shifts the labor demand curve out along the upward sloping labor supply curve. Hours, output, and real wages should all increase, resulting in procyclical movements in the real wage.

During the adjustment process, the increase in average labor productivity should be less than the increase in output because the capital stock should be slow to adjust to the higher level of technology. I find, however, that average productivity exceeds output over at least the first nine years. I have examined the effect of technology shocks on hours. Technology shocks can

[12]This comes from the production function. The permanent effect of a one percent increase in the level of technology is an increase of $(1/\alpha)$ percent in the capital stock, and hence a $(1/\alpha)$ percent increase in the capital to labor ratio and the real wage.

have no permanent effects on hours. In the short run, the response is unrestricted. I find that hours respond negatively to a technology shock during the adjustment process. This suggests that the income effect of technology shocks is greater than the substitution effect due to higher wages making leisure more expensive. However, the desired increase in the capital stock should mitigate against this. These effects must exactly cancel in the long run, as labor supply is unaffected by wages.

Because the short-run dynamics of the model are unrestricted, a technology shock may also cause a transitory shift in the labor supply curve. Shapiro and Watson (1988) find that technology shocks lower the real interest rate in the short-run. The reduction of the real interest rate may shift the labor supply curve to the left. When both the labor supply and labor demand curves shift, the effect on hours becomes ambiguous, while the effect on wages must be unambiguously positive. Again, because of a desired build-up in the capital stock, one might expect that interest rates should increase, shifting the labor supply curve in the opposite direction.

A third possible explanation is that most technological improvements are labor-saving in nature. In the short run, firms increase output with fewer employees and wages of those employed rise. Over the medium run, the capital stock has been built up and firms increase employment. In the long run, the capital stock and capital-to-labor ratio will both be higher.

Decomposition of Variance

The top panel of Table 2 shows the decomposition of forecast error variances for output. Initially, labor supply shocks explain 40 percent of the variance in output. At the eight year horizon, this increases to 56 percent. It is not controversial that labor supply explains more than half of the permanent movements in output. The large contribution of labor supply shocks to the forecast errors in output at business cycle frequencies may appear surprising. Quite to the contrary, this effect might be expected. Following Shapiro and Watson (1988), I allow hours to have a stochastic component. This component is estimated to be quite large. Since the dynamics of hours is nearly a random walk, factors that move hours in the long run must also be important in the short run. And, aside from the effects of technology shocks, increases in output must be closely related to increases in hours. It then follows that if labor supply shocks are important in explaining hours in the long run, they must be important in explaining output at all horizons.

Technology shocks explain little of the variance in output at short horizons. Mechanically, this comes from the negative effect of technology on hours, which reduces the effect of technology shocks on output in the short run. In the long run, technology shocks explain approximately 30 percent of the forecast errors in output. On average, oil price shocks explain little of the variance in output. As will be shown below, oil prices were important causes of the recessions of 1974-1975 and 1980-1982, but explained little of the variance in output during the rest of the period.

Aggregate demand shocks explain a large percentage of the forecast error variance in output at business cycle frequencies. I find that aggregate demand explains 43 percent of the 8-quarter ahead forecast error variance. This is substantially larger than the 20 percent found in SW. I attribute this difference to the restriction that labor supply shocks have no permanent effects on real wages. When I follow SW and do not impose the restriction that hours and output have identical long-run responses to labor supply shocks, which is equivalent to the restriction that labor supply shocks have no permanent effects on real wages, I find that aggregate demand shocks explain 23 percent of the 8-quarter ahead forecast error variance in output. This difference is also reflected in the explanatory power of technology shocks. In the unrestricted model, technology shocks explain 29 percent of the 8-quarter forecast error variance, compared with the nine percent found with the restricted model. I also find that aggregate demand shocks are more persistent in the restricted model. They explain 20 percent of the 32-quarter ahead forecast error in output in the restricted model compared with nine percent in the unrestricted model.

Table 2 also presents the decompositions of the forecast error variances for real wages. At the shortest horizons, the measurement error shock accounts for approximately 40 percent of the forecast error variance.[13] The contribution of the measurement error is quite persistent. After 32 quarters, the measurement error shock still explains nearly 26 percent of the variance in real wages.

[13]The measurement error explains 41.7 percent of the 1-quarter ahead forecast error variance in real wages with a standard error of 6.3 percent. I find this persuasive evidence that the five structural shocks do not explain all of the variance in real wages, and that the addition of the measurement error shock is appropriate in this framework.

Labor supply shocks are relatively unimportant at all horizons, explaining no more than four percent of the forecast error variance in real wages at any horizon. Technology shocks are most important in explaining the forecast error variance of real wages. At business cycle frequencies (one to 12 quarters), technology shocks explain between 29 and 37 percent of the forecast error variance. At the infinite horizon, technology shocks explain nearly all of the forecast error variance in real wages. Oil price shocks are also important in explaining the forecast error variance in real wages. They explain more than 15 percent of the variance in the 12-quarter ahead forecast error. Oil price shocks have permanent effects on real wages. At the infinite horizon they will explain all the variance in real wages not explained by the technology shocks. Aggregate demand shocks explain 30 percent of the one-quarter ahead forecast error variance. Their contribution to the variance in real wages falls of more rapidly than their contribution to the variance in output.

K-Quarter Ahead Forecast Errors

The k-quarter ahead forecast errors are closely related to the impulse response functions and variance decompositions described above. I analyze the forecast errors for three reasons. First, the graphs of the forecast errors in output and real wages, figure 3 panels a through d, provide information on the importance of each of the shocks in explaining the business cycles during the sample period.[14] Second, the correlation between the k-quarter ahead forecast errors of output and real wages provides a simple summary statistic for the cyclical properties of real wages in response to the different shocks. Finally, I can use the series constructed from the forecast errors to estimate the output elasticity of real wages in response to each of the shocks. I can then compare these estimated elasticities with more standard estimates available in the literature.

The decomposition of variance indicated that, on average, labor supply shocks were responsible for much of the movements in output but little of the movements in the real wage. This is clearly indicated on figure 3a. Labor supply shocks were particularly important in

[14]See Shapiro and Watson (1988) for an episodic description of the business cycles during the sample period.

causing the recession of 1975. It is difficult to infer the cyclicality of real wages in response to labor supply shocks from this figure.

Figure 3b shows that wages and output move closely together in response to oil price shocks. Figure 3b also explains why oil price shocks on average explain relatively little of the forecast error variance in output. Except in the recessions of 1974-1975 and 1980-1982, oil price shocks have relatively small effects on output or real wages. But, both of the two large oil shocks in the 1970s had substantial impacts on output. In particular, the OPEC I shock explained more than one third of the movements in output during the 1975 recession that could not be expected eight quarters earlier. The second OPEC shock also explains much of the decrease in output in the in the 1980-82 double dip recession. This figure, however, also supports Hamilton's (1983) finding that even in the pre-1973 period, increases in oil prices preceded recessions. Figure 3b suggests that oil price shocks played some role in the 1954 and 1958 recessions.

Technology shocks were important, on average, in explaining the movements in output and real wages over the 1951 - 1985 period. Figure 3c also shows the importance of technology shocks in general. Again, real wages and output move closely together. Episodic analysis of the technology shocks is less clear cut. The series of mostly negative technology shocks from 1966 to 1971 corresponds to the onset of the productivity slowdown of the late 1960s and early 1970s. It also appears that many of the movements in output from the late 1950s to the mid 1960s were due to faster than average growth in productivity. And, the 1980s seems to have been a period of frequently negative technology shocks. The series of negative productivity shocks in the 1980s may be due to the declining percentage of output from the high productivity manufacturing sector and increasing percentage from the lower productivity service sector. The importance of especially rapid technology growth in the 1960s and slower growth in the 1980s may help explain why earlier studies of real wage cyclicality found real wages to be slightly procyclical over the entire post-war period, not just over the 1970s, a period conventionally thought to have been dominated by supply shocks.

Figure 3d traces out the effects of aggregate demand shocks on output and the real wage. The recession of 1981-82 is in great part explained by a negative aggregate demand shock, validating its labeling as the "Volcker recession." There is also evidence that earlier Fed

tightenings were followed by recessions. The shifts in Fed policy identified by Romer and Romer (1989) in 1955:2, 1968:4, 1974:2, 1978:3 and 1979:4 precede decreases in output due to demand shocks by approximately 2 years. Romer and Romer (1989) find that the peak increase in unemployment in response to a Romer date is 8 to 12 quarters after the initial shift in monetary policy.

The countercyclicality of real wages in also generally apparent from this figure. Because the demand shock is really a composite of two underlying demand shocks that might have different effects on real wages, the cyclicality of real wages does vary more across time, although not in any consistent fashion, than it does in response to the other shocks. This countercyclical movement in real wages is consistent with models where nominal wages are stickier than nominal prices. One way of thinking about equation (17) is that the included right hand side variables forecast changes in the inflation rate. One part of the composite aggregate demand shock can be interpreted as an inflation surprise.[15] Then, if there is a Phillips curve, the unexpected increase in inflation should raise output. If nominal wages are stickier than prices in the short run and if there is involuntary unemployment so that workers are to the left of their short-run labor supply curves, then the inflation surprise should also reduce real wages, inducing firms to hire more labor. Similarly, if workers misperceive increases in nominal wages as increases in real wages, there is effectively a temporary shift out in the labor supply curve and output would increase and real wages would fall. The aggregate demand shocks have interesting effects on average labor productivity as well. Unlike real wages, labor productivity increases in response to aggregate demand shocks. This is consistent with firms hoarding or underutilizing their labor forces during non-expansionary periods. Hall (1988) and Rotemberg and Summers (1990) offer explanations for aggregate demand driven increases in labor productivity.

The series of k-quarter ahead forecast errors provide a computationally simple way of quantifying the correlations of output and real wages in response to the aggregate supply and aggregate demand shocks. These correlations indicate the cyclicality of real wages in response to the structural shocks.

[15]The impulse response function for the shock from the inflation equation shows that output increases and real wages decrease following the shock.

Table 3 shows the correlations of the 4-, 8, and 16-quarter ahead forecast errors of output and real wages. Tables 4a and 4b present the estimated output elasticities of real wages in response to the structural shocks. The difference between Tables 4a and 4b is in the computation of standard errors. The standard errors in Table 4a are generated by bootstrapping 500 times from estimated structural shocks and are taken from the generated distribution of the elasticities. The standard errors in Table 4b are estimated directly from the regressions of real wages on output in response to each of the shocks. Table 4b reflects the intrasample variation in the response of real wages to output, while Table 4a reflects the variation in average response across the different simulated samples.

Labor supply shocks generate no consistent cyclical pattern in the real wage. The estimated correlations are very imprecisely estimated. Similarly, the elasticities in Table 4a are very imprecisely estimated, with standard errors up to ten times the size of the estimated coefficients.

Oil price shocks generate movements in output and real wages that are nearly perfectly correlated. The correlation in the four quarter ahead forecast errors is 0.88. At the 8 and 16 quarter horizons these correlations are 0.95. At the infinite horizon, this correlation is constrained to be one because all movements in output in the long-run are due to reductions in the use of capital, so that a one percent change in output is accompanied by a one percent change in the capital-to-labor ratio and the marginal product of labor. The estimated elasticities of real wages in response to output generated by the oil price shocks are very large. They are 1.38, 0.63, and 0.76 at the three horizons studied. Each is insignificantly different from one.

Similarly, the long-run responses of output and real wages to technology shocks are identical. These responses are not constrained at the shorter horizons. Yet, they are very highly correlated. The estimated correlations are 0.65, 0.80, and 0.93, with reasonably tight standard errors at the two longer horizons. The estimated elasticities are 2.07, 0.81, and 0.87. Again, each of these is insignificantly different from one. At the infinite horizon, this elasticity must be equal to one.

Aggregate demand shocks cause countercyclical movements in real wages. The estimated correlations between real wages and output in response to demand shocks are -0.32, -0.50, and -0.53. These are less precisely estimated. However, the probabilities that these correlations are

greater than one are only 9.3 percent and 9.0 percent at the 8 and 16 quarter horizons. These correlations are unconstrained by the model at all horizons. At the longer horizons they would not be very meaningful, as little of the long-run finite movements in real wages and output will be explained by the demand shocks. The hypothetical regressions of real wages in response to aggregate demand shocks on output in response to these same shocks show that at the four, eight, and 16 quarter horizons, a one percent increase in output decreases real wages by 0.12 percent, 0.19 percent, and 0.16 percent, respectively. Again, using the standard errors from Table 4a, these are imprecisely estimated, but the probabilities tht the elasticities are greater than zero are less than 50 percent.

The structural shocks clearly generate different cyclical covariances of real wages and output. This result casts doubt on the usefulness of studies of that constrain real wage cyclicality to be the same across all business cycles. In addition, each business cycle generally has more than one cause, and all of the shocks except the oil price shocks were important over the entire sample period. Therefore, examining the covariance of real wages and output over different sample periods is not very instructive. SBP use evidence of similar cyclicalities of real wages over their sample period (1968 - 1987) and the post-war period through 1987 as evidence that their results are generalizable to the earlier period. The results in this paper cast doubt on the generalizability of the SBP findings. Similarly, macroeconomists have used the panel data results as the new stylized facts about the cyclicality of real wages. This conclusion might have been premature, as the structural shocks affect real wages differently.

Table 4 also presents evidence on the cyclicality of the forecast errors due to the sum of all of the structural shocks. The output elasticities of real wages are -0.005, 0.183, and 0.357 at the four, eight, and 16 quarter horizons. At the two longer horizons, the elasticities are larger than the 0.153 estimated by SBP for the 1948 - 1987 period, and comparable to the 0.293 estimated for the 1968 - 1987 period. Both estimates use aggregate wage data and are thus susceptible to composition bias.

It is difficult to assess the implications of composition bias for examining the relative cyclicality of real wages in response to the shocks identified by this paper without further information on how each of the shocks affects the composition of the labor force. Although it is implausible that all business cycles are alike, it is unlikely that the differences in the cyclicality of real wages across different business cycles described by Blanchard and Watson (1986) are due to variation in composition bias. Similarly, it is impossible to conclude that, for example, composition bias is more important in response to labor supply shocks, which generate no consistent pattern of real wage cyclicality, than in response to technology shocks, which result in procyclical real wages. For this reason, I will assume that composition bias is equally important in the responses to each of the shocks.

One final issue must be addressed. It is possible that the inclusion of the measurement error shock biases the results by imposing spurious procyclicality or countercyclicality on the portion of real wage movements explained by the five structural shocks. I find little evidence for this hypothesis.

Figure 4a shows the measurement error shocks generated by this model (variance normalized to one). There is no apparent cyclical pattern. Figure 4b shows the 8-quarter ahead forecast errors in the real wage due to the measurement error shock. Again, there is no apparent cyclical pattern in these responses.

To search for any cyclical pattern in this series, I have estimated the correlations between the forecast errors in real wages due to the measurement error shock and the forecast errors in output explained by all of the shocks together. These correlations are 0.14, 0.10, and 0.16 at the 4-, 8-, and sixteen-quarter horizons, respectively. The standard errors are 0.09, 0.13, and 0.18. There is little evidence that there is a cyclical pattern in the response of real wages to the measurement error shock.

VI. CONCLUSIONS

This paper uses a variant of the dynamic model of the U.S. economy in Shapiro and Watson (1988) to study the cyclicality of real wages. The model identifies five structural shocks that drive the macroeconomy: Labor supply shocks, technology shocks, oil price shocks, and two aggregate demand shocks. I study the cyclicality of real wages in response to each of these shocks. The five shocks cannot explain all of the variance in real wages. I add a real wage equation to the model but cannot add a sixth structural shock. Instead, noting that there are strictly non-allocative movements in the real wage, I add a measurement error to the real wage

equation. The measurement error can have only transitory effects on the real wage, and no effects on any of the other variables.

There are important differences in the cyclicality of real wages in response to the structural shocks. I find that labor supply shocks cause no consistent cyclical pattern in real wages. Technology shocks and oil price shocks, which can both be interpreted as labor demand shocks, cause real wages to move procyclically. There is less agreement on how real wages should behave in response to aggregate demand shocks. Sticky nominal wage models fell out of favor with macroeconomists when recent evidence showed that real wages are procyclical. Models in which most movements in output are due to movements in aggregate demand needed to be consistent with this evidence. Sticky price models can explain procyclical movements in the real wage in response to aggregate demand shocks while sticky wage models cannot. Here, I attempt to resolve this ambiguity. Within the context of the model estimated here, aggregate demand shocks can cause real wages to move countercyclically even if the overall real wage moves procyclically, because technology shocks and oil price shocks move real wages procyclically. I find some support for the sticky wage models. Aggregate demand shocks cause real wages to move countercyclically.

These findings demonstrate that simple correlations between real wages and output, hours, or the unemployment rate do not adequately describe the facts about the cyclicality of real wages. Models where aggregate demand shocks are dominant should be less concerned about matching the stylized fact of an overall procyclical real wage. Models that predict countercyclical movements in real wages in response to aggregate demand shocks are consistent with the evidence presented here on real wage cyclicality. I find that aggregate demand shocks explain in excess of 40 percent of the forecast error variance in real output at business cycle frequencies and cause countercyclical movements in real wages.

References

Altug, Sumru. 1989. "Time-to-Build and Aggregate Fluctuations:Some New Evidence." International Economic Review. 30:4. p. 889-920.

Barsky, Robert B. 1987. "The Fisher Hypothesis and the Forecastability and Persistence of Inflation." Journal of Monetary Economics. 19:1. p. 3-24.

Bils, Mark J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." Journal of Political Economy. 99:4. p. 666-689.

Blanchard, Olivier J. 1989. "A Traditional Interpretation of Macroeconomic Fluctuations." American Economic Review. 79:5. p. 1146-1164.

Blanchard, Olivier J. and Danny Quah. 1989. "The Dynamic Effects of Aggregate Demand and Aggregate Supply Disturbances." The American Economic Review. 79:4. p. 655-673.

Blanchard, Olivier J. and Mark W. Watson. 1986. "Are All Business Cycles Alike?" in Robert J. Gordon, ed. The American Business Cycle: Continuity and Change. Chicago: University of Chicago Press.

Blough, Stephen R. 1991. "The Relationship Between Power and Level for Generic Unit Root Tests in Finite Samples." Unpublished manuscript, The Johns Hopkins University.

Cogley, Timothy. 1993. "Empirical Evidence on Nominal Wage and Price Flexibility." Quarterly Journal of Economics. 108:2. p. 475-491.

Fischer, Stanley. 1977. "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule." Journal of Political Economy. 85:4. p. 191-205.

Geary, Patrick T. and John Kennan. 1982. "The Employment-RealWage Relationship: An International Study." Journal of Political Economy. 90:4. p. 854-871.

Hall, Robert E. 1988. "The Relationship Between Price and Marginal Cost in US Industry." Journal of Political Economy. 96:5. p. 921-947.

Hamilton, James D. 1983. "Oil and the Macroeconomy since WorldWar II." Journal of Political Economy. 91:2. p. 228-248.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson. 1991. "Stochastic Trends and Economic Fluctuations." American Economic Review. 81:4. p. 819-840.

Kydland, Finn and Edward C. Prescott. 1990. "Business Cycles: Real Facts and a Monetary Myth." Federal Reserve Bank of Minneapolis Quarterly Review. 14:2. p. 4-18.

Liederman, Leonardo. 1983. "The Response of Real Wages to Unanticipated Money Growth." Journal of Monetary Economics. 11:1. p. 73-88.

McCallum, Bennett T. 1986. "On 'Real' and 'Sticky-Price' Theories of the Business Cycle." Journal of Money, Credit, and Banking. 18:4. p. 397-414.

Mankiw, Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." Quarterly Journal of Economics. 100:2. p. 529-539.

Mork, Knut Anton. 1989. "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results." Journal of Political Economy. 97:3. p. 740-744.

Neftci, Salih N. 1978. "A Time-Series Analysis of the Real Wages-Employment Relationship." Journal of Political Economy. 86:2, pt. 1. p. 281-291.

Romer, Christina D. and David H. Romer. 1989. "Does MonetaryPolicy Matter? A New Test in the Spirit of Friedman and Schwartz." NBER Macroeconomics Annual 1989. Cambridge: MIT Press. p. 121-170.

Rotemberg, Julio J. and Lawrence H. and Summers. 1990. "Inflexible Prices and Procyclical Productivity." Quarterly Journal of Economics. 105:3. p. 851-874.

Rotemberg, Julio J. and Michael Woodford. 1991. "Markups and the Business Cycle." NBER Macroeconomics Annual 1991. Cambridge: MIT Press. p. 63-129.

Rotemberg, Julio J. and Michael Woodford. 1992. "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity." Unpublished manuscript.

Sargent, Thomas J. 1978. "Estimation of Dynamic Labor Demand Schedules under Rational Expectations." Journal of Political Economy. 86:6. p. 1009-1044.

Shapiro, Matthew D. 1993. "Federal Reserve Policy: Cause and Effect." Forthcoming in Monetary Policy. N. Gregrory Mankiw, Ed. Chicago: University of Chicago Press.

Shapiro, Matthew D. and Mark W. Watson. 1988. "Sources of Business Cycle Fluctuations" NBER Macroeconomics Annual 1988. Cambridge: MIT Press. p. 111-156.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1992. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" National Bureau of Economic Research Working Paper No. 4202.

Sumner, Scott and Stephen Silver. 1989. "Real Wages, Employment, and the Phillips Curve." Journal of Political Economy. 97:3. p. 707-720.

Taylor, John. 1979. "Staggered Wage Setting in a Macro Model." American Economic Review. 69:2. p. 108-113.

Table 1
Unit Root Descriptive Statistics

Series	Dickey-Fuller t-statistic[1]	Implied Largest Autoregressive Root
h_t	-3.21	0.94
y_t	-2.16	0.95
y_t-h_t	-1.04	0.99
w_t	-1.41	0.99
w_t-(y_t-h_t)	-3.20	0.88
π_t	-2.64	0.80
i_t	-2.05	0.96
r_t	-3.08	0.77

[1]Augmented Dickey-Fuller tests performed over the period from 1948:1 to 1992:2 using six lags. Tests for a unit root in h_t, y_t, (y_t-h_t), w_t, and (w_t-(y_t-h_t)) include a constant and a time trend. The five and ten percent critical values, taken from Harvey (1990) are -3.45 and -3.15, respectively. Tests for a unit root in π_t, i_t, and r_t include a constant. The five and ten percent critical values are -2.89 and -2.57, respectively.

Table 2
Decomposition of Forecast Error Variance[1]

Fraction of Output Explained by Shock to

Quarter	Labor Supply	Oil	Technology	Aggregate Demand
1	40.2	0.2	2.1	57.5
	(18.0)	(2.0)	(9.1)	(28.4)
4	45.4	1.3	1.2	54.2
	(18.3)	(3.0)	(6.5)	(23.4)
8	38.4	9.6	9.2	42.8
	(16.8)	(6.9)	(9.0)	(15.7)
12	38.8	9.3	13.0	38.9
	(16.6)	(7.1)	(9.9)	(13.5)
20	44.9	8.0	17.3	29.8
	(15.5)	(6.9)	(10.3)	(10.9)
32	55.7	6.1	20.2	20.0
	(13.9)	(6.0)	(10.5)	(8.1)

[1]Standard errors, in parentheses, are calculated by first bootstrapping 500 draws from the series of structural shocks, next generating simulated series, and then reestimating the model. Actual oil prices are used in all simulations.

Table 2, Continued

Fraction of Real Wage Explained by Shock to

Quarter	Labor Supply	Oil	Technology	Aggregate Demand	Measurement Error
1	0.0	0.1	29.2	30.0	41.7
	(7.5)	(2.1)	(14.1)	(15.4)	(6.3)
4	0.1	7.9	31.6	24.4	36.9
	(7.3)	(6.6)	(14.8)	(14.2)	(7.7)
8	0.5	12.6	32.3	19.0	35.6
	(6.9)	(9.5)	(14.9)	(11.9)	(8.3)
12	0.5	15.3	37.2	13.5	33.6
	(6.2)	(10.7)	(14.5)	(9.1)	(8.0)
20	1.6	17.0	42.8	8.4	30.2
	(5.7)	(11.7)	(15.0)	(6.4)	(8.0)
32	3.7	18.4	45.9	6.4	25.5
	(6.6)	(12.3)	(16.4)	(5.1)	(7.9)

Table 3
Correlations of Real Wages and Output in Response to Shocks[1]

In Response to shock in:	4-Quarter Ahead Forecast Error	8-Quarter Ahead Forecast Error	16-Quarter Ahead Forecast Error
All Shocks	-0.01 (0.18)	0.35 (0.21)	0.58 (0.23)
Labor Supply	0.73 (0.86)	0.76 (0.80)	-0.20 (0.68)
Oil Price	0.88 (0.54)	0.95 (0.13)	0.95 (0.16)
Technology	0.65 (0.64)	0.80 (0.35)	0.93 (0.19)
Aggregate Demand	-0.32 (0.40)	-0.50 (0.38)	-0.53 (0.39)

[1]The k-quarter ahead forecast errors are constructed as the moving average coefficients on real wages and output times the realizations of the structural shocks. The forecast errors due to aggregate demand shocks reported are the sum of the two forecast error series due to the individual aggregate demand shocks. Standard errors, in parentheses, are calculated by first bootstrapping 500 draws from the series of structural shocks, next generating simulated series, and then reestimating the model. Actual oil prices are used in all simulations.

Table 4a
Elasticities of Real Wages with Respect to Output in Response to Shocks[1]

In Response to shock in:	4-Quarter Ahead Forecast Error	8-Quarter Ahead Forecast Error	16-Quarter Ahead Forecast Error
All Shocks	-0.005 (0.090)	0.183 (0.109)	0.357 (0.152)
Labor Supply	0.026 (0.270)	-0.053 (0.235)	-0.021 (0.187)
Oil Price	1.380 (1.157)	0.629 (0.313)	0.762 (0.330)
Technology	2.068 (1.195)	0.814 (0.483)	0.865 (0.373)
Aggregate Demand	-0.116 (0.187)	-0.190 (0.190)	-0.160 (0.195)

[1]The reported values are the coefficients on output from regressions of the k-quarter ahead forecast error in the real wage on a constant and the k-quarter ahead forecast error in output. Standard errors, in parentheses, are calculated by first bootstrapping 500 draws from the series of structural shocks, next generating simulated series, and then reestimating the model. Actual oil prices are used in all simulations.

Table 4b
Elasticities of Real Wages with Respect to Output
in Response to Shocks[1]

In Response to shock in:	4-Quarter Ahead Forecast Error	8-Quarter Ahead Forecast Error	16-Quarter Ahead Forecast Error
All Shocks	-0.005	0.183	0.357
	(0.044)	(0.045)	(0.045)
Labor Supply	0.026	-0.053	-0.021
	(0.002)	(0.004)	(0.010)
Oil Price	1.380	0.629	0.762
	(0.068)	(0.019)	(0.022)
Technology	2.068	0.814	0.865
	(0.216)	(0.055)	(0.031)
Aggregate Demand	-0.116	-0.190	-0.160
	(0.031)	(0.030)	(0.023)

[1]See Table 4a for a description of the reported values. Standard errors of the coefficients from the regressions using the actual data are in parentheses.

Figure 1. Log of Real Oil Prices
1951:1 = 100

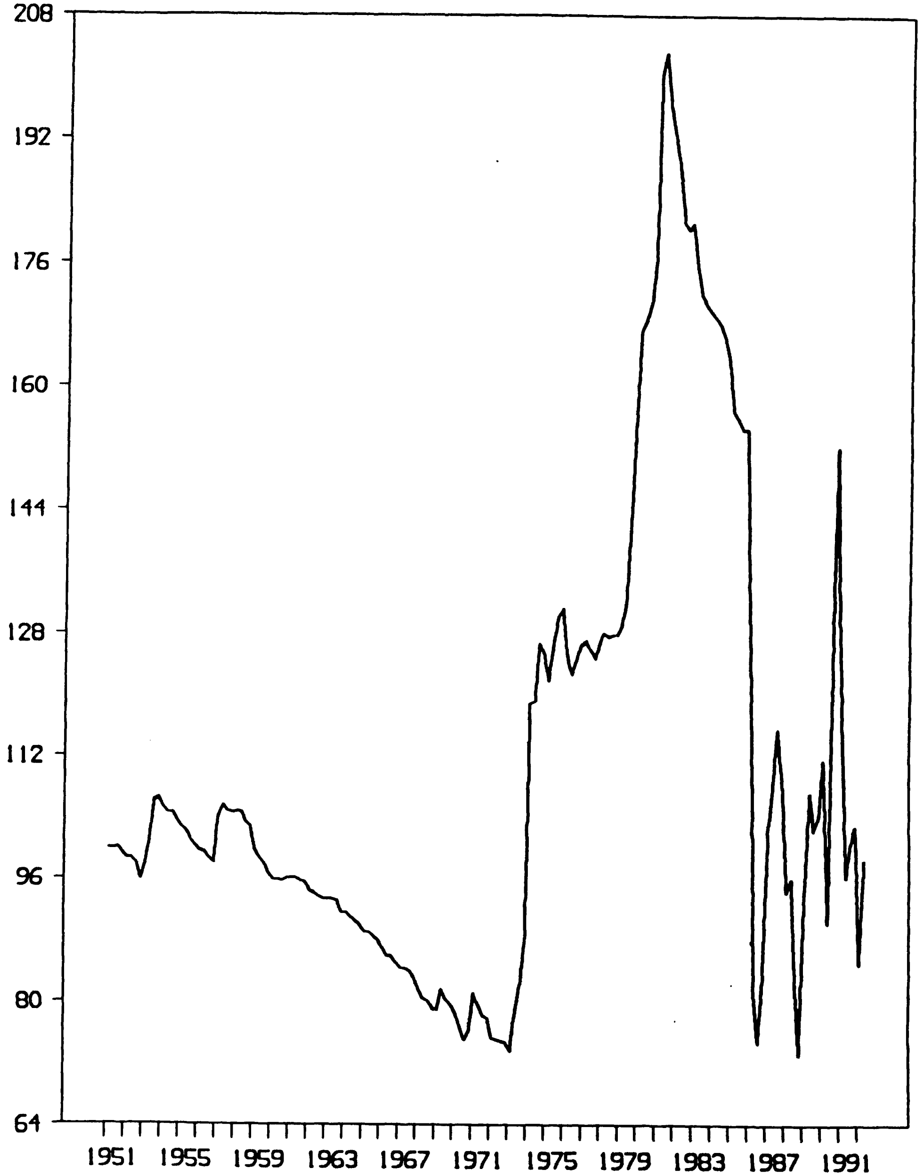

Figure 2.a. Response of Output to Labor Supply Shocks

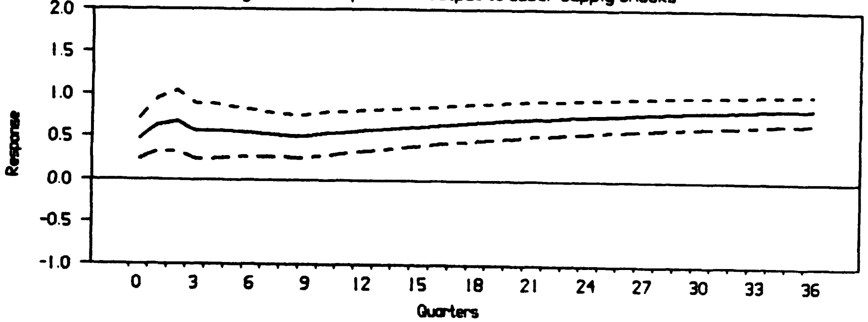

Figure 2.b. Response of Wages to Labor Supply Shocks

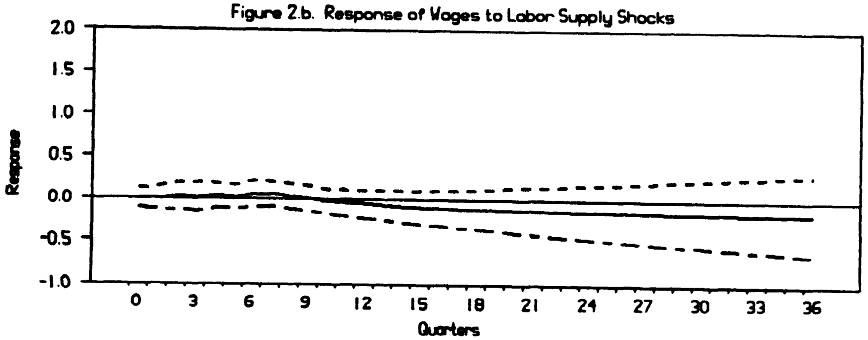

Figure 2.c. Response of Average Labor Productivity to Labor Supply Shocks

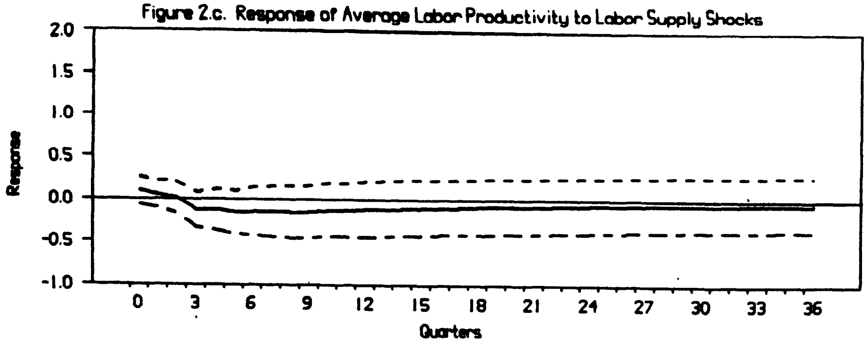

Figure 2.d. Response of Output to Oil Price Shocks

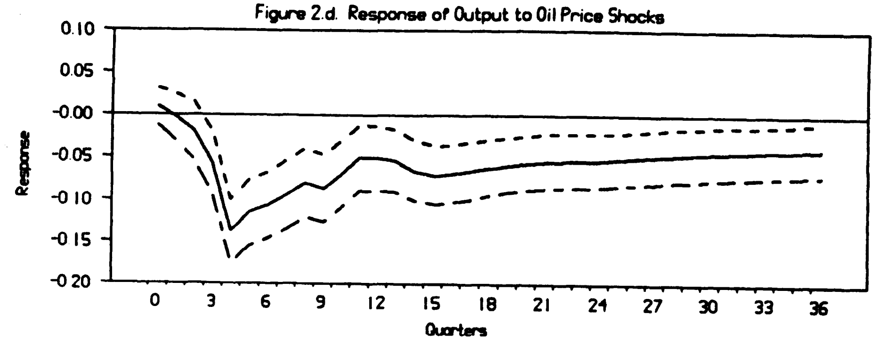

Figure 2.e Response of Wages to Oil Price Shocks

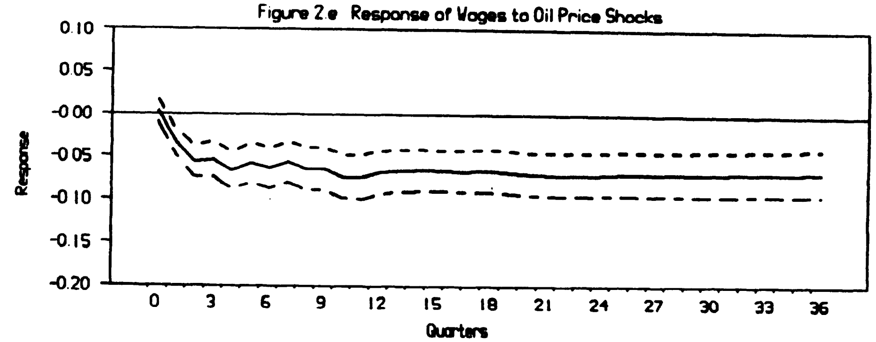

Figure 2.f Response of Average Labor Productivity to Oil Price Shocks

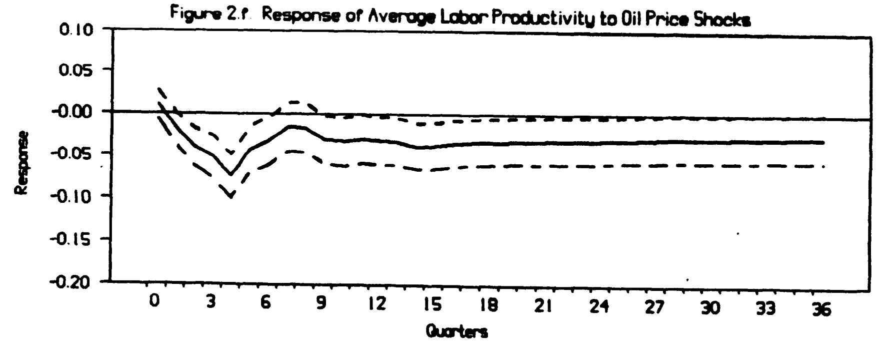

Figure 2.g. Response of Output to Technology Shocks

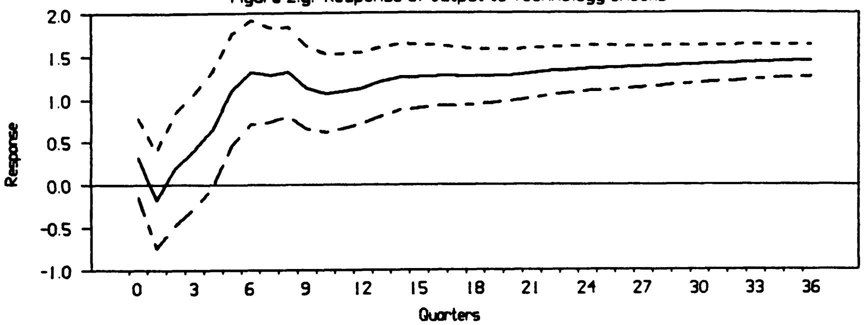

Figure 2.h. Response of Wages to Technology Shocks

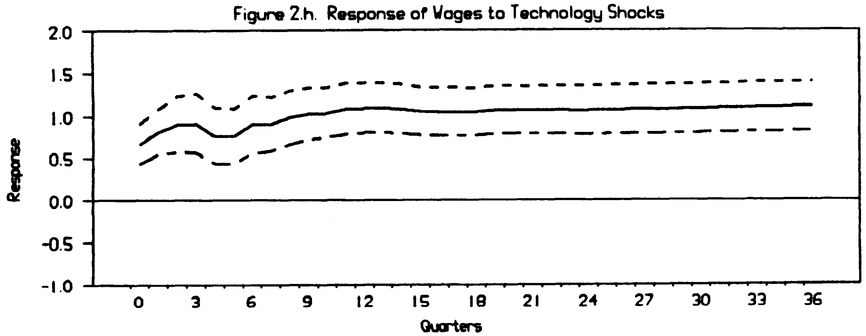

Figure 2.i. Response of Average Labor Productivity to Technology Shocks

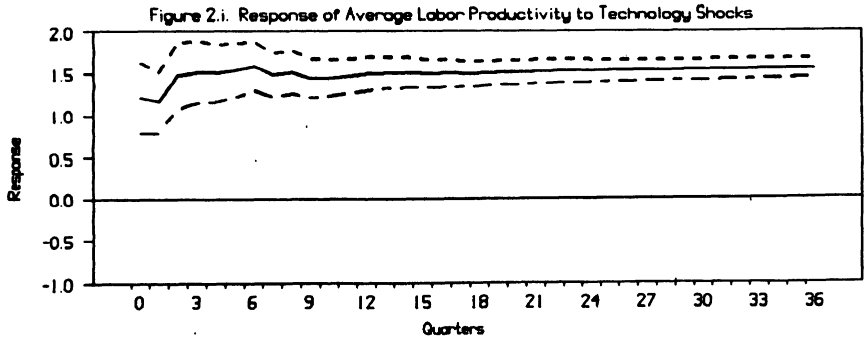

Figure 4.a. Measurement Error Shocks

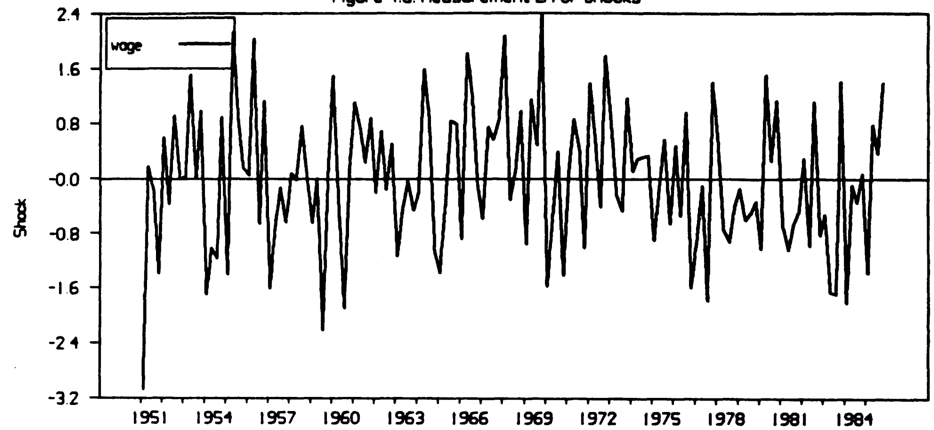

Figure 4.b 8 Quarter Ahead Forecast Errors due to Measurement Error Shocks