

NBER WORKING PAPER SERIES

IMPERFECT COMPETITION AND THE
EFFECTS OF ENERGY PRICE
INCREASES ON ECONOMIC ACTIVITY

Julio J. Rotemberg
Michael Woodford

Working Paper 5634

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1996

We wish to thank Stephanie Schmitt-Grohé for research assistance, IGIER for its hospitality while we carried out some of this work and the NSF, the Sloan Foundation and the MIT Center for Energy and Environmental Policy Research for research support. This paper is part of NBER's research program in Economic Fluctuations and Growth. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1996 by Julio J. Rotemberg and Michael Woodford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

IMPERFECT COMPETITION AND THE
EFFECTS OF ENERGY PRICE
INCREASES ON ECONOMIC ACTIVITY

ABSTRACT

We show that modifying the standard neoclassical growth model by assuming that competition is imperfect makes it easier to explain the size of the declines in output and real wages that follow increases in the price of oil. Plausibly parameterized models of this type are able to mimic the response of output and real wages in the United States. The responses are particularly consistent with a model of implicit collusion where markups depend positively on the ratio of the expected present value of future profits to the current level of output.

Julio J. Rotemberg
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139
and NBER

Michael Woodford
Department of Economics
Princeton University
Princeton, NJ 08544
and NBER

Many authors have observed that changes in the price of oil on world markets appear to have a significant effect on economic activity. Rasche and Tatom (1981), Darby (1982), Hamilton (1983, 1995), Burbidge and Harrison (1984), Gisser and Goodwin (1986), Mork (1989), Carruth, Hooker and Oswald (1995), and others argue that oil price shocks were responsible for substantial aggregate fluctuations in recent decades. In spite of this voluminous empirical literature suggesting that oil price shocks have an important effect on economic activity, there is little consensus on the reason why this is so.

It is easily argued that an exogenous increase in the cost of a factor of production should reduce the quantity of final output that firms will choose to supply. What is less obvious is that the effect should be significant, if the factor of production in question accounts for only a small part of the total marginal cost of production, as is true of energy costs. Indeed, we present below a numerical estimate of the predicted effect of an increase in energy prices in a "calibrated" one-sector stochastic growth model, and show that while the oil price increase is predicted to contract output, the effect is only about a fifth the size of the response that we estimate using U.S. data. A ten percent innovation in the price of oil is predicted to contract private sector output by about one-half a percent; our estimates indicate instead that such an innovation has on average been associated with an output decline of 2.5 percent, five or six quarters after the innovation.

The observed effects of oil shocks are even more puzzling when the effects on real wages are considered as well. In standard growth models, the predicted contraction of the supply of output is greater the less real wages fall in response to the shock, and is greatest if real wages actually increase (perhaps because the product wage rises relative to the consumption wage). Thus high real wages play a crucial role in explanations like that of Bruno and Sachs (1985) of the effects of the oil shocks of the 1970's. Yet, like Bohi (1989) and Keane and Prasad (1991), we find that oil shocks typically reduce real wages. Our estimates suggest real wages fall by nearly one percent (again, five or six quarters after the innovation) for each ten percent innovation in oil prices. This is again nearly five times as large an effect as our calibrated growth model would predict. But more to the point, variations in the specification

of labor supply behavior (a point on which our model is obviously open to criticism) that would improve the model's ability to account for a sharp output decline (by predicting a greater degree of "real wage resistance") would result in even less ability to account for the observed decline in real wages, and vice versa. This suggests that it is the growth model's simple specification of output supply that must be rejected, rather than its model of labor supply behavior. ¹

The alternative that we explore in this paper continues to assume a simple aggregative model of output supply, but drops the assumption that firms produce for a perfectly competitive product market. Instead, we consider the effects of several simple models of imperfect competition in the product market, introduced in our previous papers (1991, 1992, 1995). We find that allowing for a modest degree of imperfect competition significantly increases the predicted effects of an energy price increase on both output and real wages. In particular, we show that a model involving implicit collusion between oligopolists can account for declines in both output and real wages of the magnitude that we estimate.

This study complements our previous work on the effects of innovations in military purchases on output and real wages (1992). As in that study, we are interested in the effects of oil price changes not simply because they appear to have been an important source of aggregate fluctuations in the U.S. in recent decades, but above all because variations in oil prices represent a particularly good example of an exogenous shock that can be directly identified in the data. As Hamilton (1985) has argued, there is little reason to believe that changes in the price of oil represent responses to U.S. economic conditions, and in particular little reason to believe that they should be correlated with changes in the U.S. production technology. Indeed, this has led authors such as Ramey (1991) and Hall (1988, 1990) to use oil price changes as demand-shock instruments for other purposes. We follow them in this identifying assumption.

We proceed as follows. In section 1, we estimate the responses of private sector output

¹In particular, we do not believe that simply replacing the neoclassical labor supply curve by an efficiency wage schedule, as proposed by Carruth, Hooker and Oswald (1995), would significantly improve upon the predictions of the neoclassical growth model.

and real wages to oil price increases. This section provides the facts that we then seek to explain. Section 2 gives an intuitive discussion of why the existence of imperfect competition accentuates the reductions in output and real wages. Section 3 presents the class of aggregative intertemporal general equilibrium models that we analyze numerically. We show that a single specification allows us to nest as special cases a model with perfectly competitive product markets (closely related to the model of Kim and Loungani (1992)), a model with monopolistically competitive product markets, a model with “customer markets” in the style of Phelps and Winter (1970), and a model with implicit collusion in the style of our (1992) paper. Section 4 discusses the calibration of the models. Section 5 compares the numerical responses of output and real wages implied by our various models with the estimated responses from section 1. Section 6 concludes.

1 The Observed Effects of Energy Price Shocks

In this section we discuss our estimates of the effects on the U.S. economy of a shock to world oil prices. In the models we discuss below, the variable that matters for the determination of output and real wages is the *real* price of energy, rather than the level of nominal energy prices. Thus it might seem that we should simply seek to identify the effects of innovations in the real price of energy. But this method would not identify a shock that we can plausibly treat as exogenous with respect to other shocks to the U.S. economy.

What is more plausibly exogenous in the period we study is the nominal price of energy. The reasons for this, as explained in Hamilton (1985), have to do with the institutions that set oil prices in this period. As he documents, the nominal U.S. price of oil in the pre-OPEC period was set to a large degree by the Texas Railroad Commission (TRC). The TRC tended to keep the nominal price constant (and allowed the quantity produced to fluctuate so demand would be met) unless a large exogenous disturbance occurred. Thus, the nominal price was changed in 1952 as a result of the Iranian nationalization of oil assets, in 1956 as a result of the Suez crisis and so on. The policy of keeping the dollar price of oil fixed between major realignments (that coincided with exogenous disturbances) was maintained in

the OPEC era. Indeed, Hamilton (1985) quotes Kuwait's oil minister as saying the nominal oil price "should be frozen so that the real price (adjusted for inflation) ... would fall for two or three years". As a result of this policy the two major changes in nominal oil prices in this era were the 1973 oil embargo in response to the Arab-Israeli war and the 1979 increase in response to the Iranian revolution.

The policy of keeping the nominal oil price nearly constant between major realignments caused by exogenous events means that innovations in the *real* price of oil can also be due to unforecastable changes in U.S. inflation. These innovations in U.S. inflation need not be exogenous with respect to U.S. technology shocks, taste shocks, and the like. We therefore consider the bivariate stochastic process for nominal oil prices and a nominal price index for the U.S., and orthogonalize the two innovations by assuming that the shock of interest to us may affect both nominal oil prices and U.S. inflation, but that *the orthogonal shock has no effect on nominal oil prices within the quarter*. Thus, it is the innovation in the *nominal* oil price that actually identifies the exogenous shock that we are interested in. But it is only the effect of this shock on the forecasted path of the *real* oil price that matters for the predictions of our theoretical models about the effects of the shock.

This identifying assumption is not equally defensible over the entire period for which we have data. We believe that it makes sense for both the pre-OPEC and the OPEC periods. But sometime in the early 1980's, OPEC lost its ability to keep the nominal price of oil relatively stable. It is reasonable to assume that after this point variations in the demand for oil (and even news about its future demand, as it is a storable commodity) began to be reflected in nominal oil prices immediately. Indeed, simple examination of the time series for nominal oil prices suggests that these prices are no longer formed in the same way in the 1980's. For instance, quarters with nominal oil prices the same as in the previous quarter no longer occur. Furthermore, the growth rate of nominal oil prices is much more rapidly mean-reverting in this period than it had been previously.

The question is then when the period of exogenous nominal oil price changes ends and that of endogenous nominal oil price changes begins. Many observers agree that OPEC lost

much of its power to raise prices in the 1980's but the exact date of the break in OPEC's hold over the oil market is much more controversial. We approach this question by observing that the stochastic process for nominal oil prices is quite different in the 80's than in the earlier period, and supposing that the proper date at which to truncate our sample is the date at which this univariate process changes. As is standard in the literature (see Andrews 1993) we suppose that the most likely date of such a regime change is the point which maximizes the F-statistic for a break in regime. We thus consider a regression that explains the current quarterly percentage change in the nominal price of oil with a constant and two lags of the dependent variable. We use the producer price index for crude petroleum products as our nominal price of oil and our dependent variable runs from the fourth quarter of 1947 until the second quarter of 1989. The maximal F-statistic for a break equals 5.92 and arises when the first part of the sample includes only data until the third quarter of 1980.² The likelihood ratio for a sample break at the beginning of the OPEC regime is much smaller. The likelihood that the break occurred in 1986 when the Saudis retaliated against price chiseling by severely lowering their price is larger but still lower than the likelihood that it occurred in 1980:3. We thus use this as our break point and only consider the effect of nominal oil price changes before this date.

Analysis of the time series for nominal oil prices and the U.S. price level indicates that both series are stationary only in first differences. However their ratio, the real price of oil, appears to be stationary so that the series are cointegrated. We thus estimate a bivariate vector autoregression for the two stationary series, the growth rate of nominal oil prices and the logarithm of the real price of oil. The first of our two equations makes the current change in the logarithm of the dollar price of oil (π_{E_t}) a function of a constant, a time trend, and lags of this change as well as lags of the logarithm of the real price of oil ($p_{E_t}^r$), defined as the price of oil deflated by the U.S. private value added deflator. The second makes the

²Under the null hypothesis of no break, this is distributed according to the F-distribution with 161 and 3 degrees of freedom and is thus significant at a critical level of less than 1 per thousand (the critical value at the 1% level is 3.9). This critical level understates the size of the test because we have chosen the point where the F-statistic is maximal. Indeed, the value of this statistic is below the 10% critical value tabulated in Andrews (1993).

logarithm of our real price of oil (p_{Et}^r) a function of a constant, a time trend, and lags of this variable as well as current and lagged values of π_{Et} . The current value of π_{Et} is included so that the two innovations are orthogonal by construction. We truncated the lags in both equations when the next lag had a t -statistic below one. The estimated coefficients for these two equations are given in the first two columns of Table 1.

Treating the innovations in the first of these variables as our exogenous shock, we can combine these equations to obtain the impulse response of the real price of oil. This impulse response is plotted in Figure 1a together with confidence intervals of plus and minus two times the standard error. ³ The U.S. general price level responds very little to contemporaneous increases in the nominal price of oil so that the increase in the real price of oil is almost as large as the innovation in the nominal oil price (.98 percent for each one percent innovation in the growth rate of nominal oil prices). The real price of oil continues to rise after this point because the nominal price of oil tends to rise further. The peak real oil price occurs after 5-7 quarters (about 1.6 percent higher for each one percent initial innovation in the nominal oil price). Then, as the higher real price of oil leads the general price level to increase by more than average and leads nominal oil prices to grow by less than average, the real price of oil returns gradually to its unconditional expected value.

We also analyze the effects of this type of shock on output and the real wage. We measure the first by the real value added produced by the private sector which we compute by subtracting government value added from total GNP. Our private value added deflator is then the ratio of nominal to real private value added. Our real wage is computed by dividing hourly earnings in manufacturing by the private value added deflator. We focus on private value added rather than total GNP because our theories of pricing and production decisions (whether competitive or imperfectly competitive) do not apply to the government. ⁴ We run separate regressions explaining the logarithm of each of these two variables with two lags of the dependent variable, a time trend, the current value as well as lags of π_{Et} , and

³The standard errors are calculated using the procedure of Poterba, Rotemberg and Summers (1986).

⁴We would prefer to eliminate the U.S. oil industry as well, but we lack these data.

lags of the logarithm of p_E .⁵ The results using data from 1948:2 to 1980:3 are given in the second two columns of Table 1. Once again, we truncated the lags so that the final lag has a t -statistic greater than one. Combining the regressions in the first two columns with these latter regressions we obtain the impulse responses for the output and the real wage. These are displayed in Figures 1b and 1c, again with confidence bands of plus or minus two standard errors.

Figure 1b displays the response of real private value added. One observes that private output does indeed decline following a positive innovation in oil prices. A one percent increase in oil prices results in a reduction in output of about -.25 percent after 5-7 quarters.⁶ One interesting feature of this decline is that output is lower in the second year following the innovation than in the first (which is also when real oil prices reach their peak). Indeed, the decline is statistically significant only from quarter 3 onward.⁷

Figure 1c shows the effect on the real wage. This too declines following an increase in oil prices. Once again, the maximum decline occurs only in the second year (when it is nearly -.10 percent for each one percent increase in oil prices), although in this case the decline is statistically significant even during the first year.⁸

As noted in the introduction, the simultaneous observation of sharp declines in both output and real wages is hard to explain within the context of an aggregative competitive model. To clarify this, we first discuss a stripped-down model based on Gordon (1984). Then we turn to a more elaborate set of models and compare their quantitative predictions

⁵An alternative is to analyze the effect of changes in the nominal price of oil in a VAR consisting of changes in the nominal oil price, the real price of oil, output and the real wage. We considered such a VAR as well and obtained results that are essentially identical to those in the text. We report results based on the regressions because the fact that they contain fewer nuisance parameters makes the estimates and their standard errors easier to interpret and, perhaps, more reliable.

⁶In regressions that are not reported we also analyzed the response of hours worked in the private sector and of the unemployment rate. Oil price increases lead hours to fall and unemployment to rise as can be expected from the fall in output. The increase in unemployment, as usual, lags behind the falls in output.

⁷Carruth, Hooker, and Oswald (1995) similarly find that the effects of an oil price shock on unemployment are greatest after 7-8 quarters.

⁸Because this is the empirical finding of greatest significance for our analysis, we checked its robustness in several ways. We reran the regressions dropping the 1974 observations and we also considered separate regression for the pre- and post-1974 subsamples. All of these regressions reproduced the negative effect of nominal oil prices on the real wage though the standard errors were larger because less data were included.

to those of Figure 1.

2 Oil Price Shocks and Labor Demand: The Role of Imperfect Competition

For purposes of this illustration, we first consider the extremely simplified production structure of Gordon (1984) which abstracts from capital and materials inputs. We consider an economy with many symmetric firms and a fixed supply of capital. Each of these firms combines labor and energy to produce output using the following production function

$$Y_t = Q(V(H_t), E_t) \quad (1)$$

where H_t , E_t and Y_t represent each firm's labor input, energy input and output respectively. We assume that both the V and the Q functions are increasing in their arguments. We introduce the V function because we wish to view V as value added, which is produced with labor and capital. The introduction of the V function also allows us to assume that Q is homogeneous of degree one in its two arguments, while diminishing returns (due to the fixity of capital) are represented solely by the strict concavity of V .

Choice of the inputs E and H so as to minimize costs of production in each period implies that, at each time t , there exists a value μ_t such that

$$Q_E(V(H_t), E_t) = \mu_t p_{E_t} \quad (2)$$

$$Q_V(V(H_t), E_t)V_H(H_t) = \mu_t w_t \quad (3)$$

where p_{E_t} and w_t denote the prices of energy and labor inputs respectively, each deflated by the price of the output good. The quantity μ_t represents the inverse of the Lagrange multiplier associated with the requirement that the firm produce a given level of output. It also denotes the ratio of the price of the output good to its marginal cost of production. Thus, in the case of perfect competition, these conditions must hold in equilibrium with $\mu_t = 1$ at all times. Following Gordon (1984), we hold hours worked constant and we study the extent to which output and real wages decline when p_{E_t} rises.

Differentiation of (1), (2) and (3), keeping μ_t equal to one, yields

$$\begin{aligned}\Delta Y_t &= s_E \Delta E_t, \\ \frac{1 - s_E}{\epsilon_{EV}} \Delta E_t &= -\Delta P_{Et}, \\ \frac{s_E}{\epsilon_{EV}} \Delta E_t &= -\Delta w_t.\end{aligned}$$

where s_E denotes the share of energy costs in the value of total output, $p_E E/Q$, ϵ_{EV} denotes the elasticity of substitution between energy inputs and value added V , and Δ denotes the logarithmic derivative (i.e., ΔX is the derivative of $\log X$). These equations imply that

$$\begin{aligned}\Delta w_t &= -\frac{s_E}{1 - s_E} \Delta P_{Et}, \\ \Delta Y_t &= -\frac{s_E}{1 - s_E} \epsilon_{EV} \Delta P_{Et}.\end{aligned}$$

As Gordon (1984) argues, the elasticity of substitution of energy for value added, ϵ_{EV} , must be less than one. Otherwise, the model would be inconsistent with the rise in the share of energy as a fraction of total costs that follows increases in energy prices. The percentage declines in both output and the real wage deflated by the price of output must thus be smaller than the ratio of energy costs to value added ($s_E/1 - s_E$) times the percentage increase in energy prices. In the U.S., the ratio of energy costs to value added is about 0.04, so that the decline in both quantities must be quite small.⁹ While it is possible to obtain larger declines in output if employment falls, such reductions in employment would require that the real wage fall by even less than is indicated by the above calculation. Thus one cannot obtain substantial declines in *both* output and real wages.

In fact, such real wage declines as it is possible to obtain in this model occur only when one deflates by the price of gross output rather than a value-added deflator. It is useful to define the value-added price deflator

$$p_{Vt} \equiv \frac{Y_t - p_{Et} E_t}{V_t(H_t)}. \quad (4)$$

⁹See section 4 for further discussion of the size of this parameter.

In the case where the Q function takes the Leontieff form, this corresponds to the standard GDP deflator. For other production functions the two do not coincide, but p_{Vt} is an “ideal” (Divisia) value-added deflator. In the case of perfect competition, (2)-(4) can be combined to yield

$$V_H(H_t) = \frac{w_t}{p_{Vt}} \quad (5)$$

Thus, in terms of the value-added-deflated real wage, we obtain a labor demand curve that is invariant with respect to changes in the price of energy. This means that the competitive model can account for a fall in output and employment *only* if the real wage in terms of value added rises.¹⁰

This labor demand curve provides a useful point of view from which to see why allowance for imperfect competition matters. When μ_t differs from 1, equations (2)-(4) yield instead

$$V_H(H_t) = \mu_t \frac{w_t}{p_{Vt}} \left[\frac{1 - s_{Et}}{1 - \mu_t s_{Et}} \right] \quad (6)$$

where s_{Et} is the time t energy share, $p_{Et}E_t/Q_t$. Equation (6) gives two reasons for the value added-deflated real wage associated with a given level of employment to decline when energy prices rise. The first is that this real wage would fall if the increase in energy prices led to an increase in the markup μ_t . The second is that, even with a fixed markup, the term in square brackets will rise as long as the energy share s_{Et} rises and the markup exceeds one. In particular, holding the markup and employment constant, we obtain

$$\Delta(w_t/p_{Vt}) = \frac{-(\mu - 1)}{(1 - s_E)(1 - \mu s_E)} \Delta s_{Et} \quad (7)$$

As long as the elasticity of substitution $\epsilon_{EV} < 1$, the share of energy s_{Et} rises with an increase in the price of energy, so that the real wage declines. Moreover, the required percentage decline is bigger, the larger is the markup.

The intuition for this result is clearest in the case where there is a fixed amount of energy needed to produce each unit of final good. Suppose that, initially, the production of

¹⁰Thus an energy price increase is not equivalent to an adverse technology shock, which would shift this labor demand curve. Energy price increases are treated as equivalent in the standard textbook view.

a product requires one dollars' worth of energy and that the price of energy inputs rises by 20 percent. Perfectly competitive firms would be willing to keep their employment constant with a constant nominal wage as long as their output prices rise by 20 cents. Such an increase in prices would keep the value-added deflator constant. Imperfectly competitive firms whose employment stayed constant would raise their price by more because they mark up their entire marginal cost. With a markup of 1.5, the 20 cent increase in their unit costs leads them to raise their prices by 30 cents. Thus, the value-added-deflated wage of these firms must fall if they are to keep their employment constant.

3 A Dynamic General Equilibrium Simulation Model

We now consider a more general production function and construct a general equilibrium model that, except for considering imperfect competition, is similar to those analyzed by Kim and Loungani (1992) and Finn (1991), especially the former. In particular, we follow the real business cycle literature in assuming that consumption and labor supply decisions are made by a representative household.

We consider an economy with many symmetric firms. Each firm has a production function of the general form

$$Y_t = Q(V_t, G(E_t, M_t)) \quad (8)$$

where Y_t , E_t and M_t represent each firm's gross output, energy input and materials input respectively, and V_t is an index of primary inputs (capital and labor) that represents an ideal index of value added. Both aggregator functions Q and G are assumed to be homogeneous degree one, increasing in their arguments, and concave. This specification generalizes that of Bruno and Sachs (1985) to allow for materials costs. The separate inclusion of materials has an important quantitative effect on our results when markets are imperfectly competitive because of the "double marginalization" distortion that arises when intermediate inputs are not priced at marginal cost. ¹¹

¹¹See Rotemberg and Woodford (1995) and Basu (1995) for further discussion.

The value-added index is assumed to be given by

$$V_t = F(K_t, z_t H_t) - \Phi_t \quad (9)$$

Here K_t represents capital inputs, H_t represents hours, z_t is an index of labor augmenting technical progress, and Φ_t represents fixed costs of production. Both z_t and Φ_t are exogenous parameters from the point of view of the firm. We assume a deterministic time trend in z_t in order to account for the observed trend growth in per capita U.S. output. In each of our imperfectly competitive models, we assume a positive value for Φ_t , so that the model reproduces the apparent absence of significant pure profits in U.S. industry despite the presence of market power. A time trend is allowed for the fixed costs as well, so that we can have a steady state equilibrium growth path in which the share of fixed costs in total costs is constant over time.

Choice of the inputs E, M, and H so as to minimize costs of production in each period (given the capital stock and the quantity produced) then implies that, at each time t , there exists a markup μ_t such that

$$Q_G(V_t, G(E_t, M_t))G_E(E_t, M_t) = \mu_t p_{Et} \quad (10)$$

$$Q_G(V_t, G(E_t, M_t))G_M(E_t, M_t) = \mu_t p_{Mt} \quad (11)$$

$$z_t Q_V(V_t, G(E_t, M_t))F_H(K_t, z_t H_t) = \mu_t w_t \quad (12)$$

where p_{Et} , p_{Mt} , and w_t denote the prices of energy, materials and labor inputs respectively, each deflated by the price of the output good. In our symmetric equilibrium, we set the price p_{Mt} equal to one at all times because each firm's materials are the output goods of other firms.

The economy also contains a large number of identical infinite-lived households. The representative household seeks to maximize

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t N_t U \left(\frac{C_t}{N_t}, \frac{H_t}{N_t} \right) \right\} \quad (13)$$

where β denotes a constant positive discount factor, N_t denotes the number of members per household in period t , C_t denotes total consumption by the household in period t , and H_t denotes total hours worked by members of the household in period t . By normalizing the number of households at one, we can use N_t also to represent the total population, C_t to denote aggregate consumption, and so on. The trend growth of N_t is another source of long-run growth in our model. We assume, as usual, that U is a concave function, increasing in its first argument, and decreasing in its second argument.

The additively separable preference specification (13) implies that consumption demand and labor supply by the household are given by time-invariant Frisch demand and supply curves of the form

$$\frac{C_t}{N_t} = C(w_t, \lambda_t)$$

$$\frac{H_t}{N_t} = H(w_t, \lambda_t)$$

where λ_t denotes the marginal utility of wealth in period t .¹² In terms of the Frisch demand functions the conditions for market clearing in the labor market and the product market respectively can be written as

$$H_t = N_t H(w_t, \lambda_t) \tag{14}$$

$$N_t C(w_t, \lambda_t) + [K_{t+1} - (1 - \delta)K_t] + G_t = Y_t - M_t \tag{15}$$

where δ is the constant rate of depreciation of the capital stock, satisfying $0 < \delta \leq 1$, and G_t denotes government purchases of produced non-energy goods. Equation (15) is the standard GNP accounting identity, except that we do not count value added by the government sector or by the domestic oil industry as part of either G_t or Y_t . Note that we assume that the materials used in each firm's production come out of other firms' production: a single produced good is both an intermediate good (materials) and a final good (consumption, investment, and government purchases).¹³

¹²See Rotemberg and Woodford (1992) for further discussion.

¹³To be more precise, we assume that there are many differentiated goods (and services), but consider only a symmetric equilibrium in which each is produced in the same quantity and sold for the same price. In such a setup there is no problem in assuming that firms must purchase other firms' products to use in

Equations (14) and (15) assume that there are no resource costs associated with energy production. No hours or final output must be devoted to the energy sector. Thus, equation (15) assumes implicitly that energy is freely available at no cost to the oligopolistic firms that sell it; the exogenous variations in p_{E_t} represent variations in the degree to which they succeed in colluding to keep the price of energy high (here taken as given rather than modeled). The rents earned by the producers of energy are distributed to the shareholders of their firms, which is to say to the representative household.¹⁴

Finally, voluntary accumulation of physical capital (or claims to it) by households requires that

$$1 = \beta E_t \left\{ \left(\frac{\lambda_{t+1}}{\lambda_t} \right) \left[\frac{F_K(K_{t+1}, z_{t+1} H_{t+1}) Q_V(V_{t+1}, E_{t+1})}{\mu_{t+1}} + (1 - \delta) \right] \right\} \quad (16)$$

at all dates.

To complete the model, we need only discuss the determination of the markup μ_t that appears in equations (10)-(12). In the case of perfect competition, $\mu_t = 1$ at all times. If firms have market power in their product markets, μ_t can exceed one. The three imperfectly competitive theories of the markup that we consider are presented at more length in Rotemberg and Woodford (1991, 1995). The first, which is based on monopolistic competition and on the assumption of homothetic tastes over bundles of differentiated goods, results in a constant markup μ greater than one. This markup depends upon the elasticity of substitution among the differentiated goods and the homotheticity of preferences implies that this elasticity is always the same in a symmetric equilibrium.

We also show in Rotemberg and Woodford (1991) that two quite different types of models with varying markups imply that desired markups depend only on the ratio of X_t to Y_t where X_t is the present discounted value of profits gross of fixed costs. Thus both models imply their own production, even though these intermediate goods are sold for a price higher than their marginal production cost.

¹⁴This specification has the benefit of great simplicity even if it is not as realistic as one would wish. We obtained essentially identical results when we followed Kim and Loungani (1992) and replaced (15) by an equation that made $Y_t - M_t - p_{E_t} E_t$ equal to $N_t C(w_t, \lambda_t) + [K_{t+1} - (1 - \delta)K_t] + G_t$. In this alternative specification, one deducts the cost of energy from $Y_t - M_t$ before one obtains the output available for consumption, investment or government purchases.

that

$$\mu_t = \mu(X_t/Y_t) \quad (17)$$

where

$$X_t = E_t \sum_{j=1}^{\infty} \alpha^j \beta^j \frac{\lambda_{t+j}}{\lambda_t} \left(\frac{\mu_{t+j} - 1}{\mu_{t+j}} \right) Y_{t+j} \quad (18)$$

The parameter α in (18) measures the rate at which new products are created as well as the probability that any collusive agreement or stock of loyal customers will survive until the next period. Its role is discussed at length in Rotemberg and Woodford (1992).

According to the customer market model of Phelps and Winter (1970), the function μ is decreasing in its argument. The reason is that firms in the customer market model set prices by trading off the benefit from exploiting existing customers (whose elasticity of demand is very low) with the benefit from expanding their customer base by attracting new customers (whose elasticity of demand is higher). Expanding the customer base is attractive because these customers will, at a later date, have low elasticities of demand. Thought of in this way, it is apparent that such firms will set high prices when demand by current customers is high relative to the demand that can be expected by future customers. Also, prices will be low if the profits from future sales are more valuable because interest rates are low. Thus, high values of X/Y which represent either high sales in the future, low interest rates or low sales today, lead to low markups.

By contrast, the function μ is increasing in its argument in the implicit collusion model of Rotemberg and Saloner (1986) and Rotemberg and Woodford (1992). In this model, the markup is set at the highest level consistent with having no firm deviate from the collusive understanding. The deviations are prevented by the threat that they will be followed by periods of very low profitability (price wars). The most effective of these punishments (and also the simplest to analyze) is such that, starting the period after the deviation, the present discounted value of profits is zero. This means that a deviating firm gives up X_t . On the other hand, deviations are more attractive in the present period when sales are higher (because a deviating firm captures more sales from its competitors) and when the markup is higher (because this means that there are more profits to obtain by undercutting the going

price slightly). Thus, high values of X_t/Y_t imply that the firms can afford to have higher values of μ_t and still avoid deviations.

To summarize, four theories of markup determination can be subsumed by (17). The first two theories simply assume that $\mu(X/Y)$ is a constant function. A rational expectations equilibrium is then a set of stochastic processes for the endogenous variables $\{Y_t, V_t, K_t, H_t, E_t, M_t, w_t, \mu_t, X_t, \lambda_t\}$ that satisfy (8)-(12) and (14)-(18), given the exogenous processes $\{p_{Et}, G_t, z_t, N_t, \Phi_t\}$.

4 Calibration of Model Parameter Values

In the next section we report the predicted responses of output and real wages to small changes in p_{Et} , in both the competitive and imperfectly competitive versions of the model just described. We analyze the response to shocks by log-linearizing the equilibrium conditions derived in the previous section, as in Rotemberg and Woodford (1992, 1995). The coefficients of the log-linearized equilibrium conditions involve various parameters, many of which are standard from real business cycle models, but others of which arise only because of our explicit treatment of energy and materials, or because of our allowance for imperfect competition and increasing returns. We “calibrate” these parameter values to be consistent with various measured features of the U.S. economy. Finally, the simulations depend on specifying theoretical constructs that correspond to our empirical measures of output and the real wage. We deal with these issues in turn.

We ensure that the equilibrium involves stationary fluctuations in suitably rescaled state variables, despite trend growth of population and productivity, by making certain additional homogeneity assumptions. First, we assume that the representative household’s preferences imply that there exists a $\sigma > 0$ such that $H(w, \lambda)$ is homogeneous of degree zero in $(w, \lambda^{\frac{1}{\sigma}})$, and that $C(w, \lambda)$ is homogeneous of degree one in $(w, \lambda^{\frac{1}{\sigma}})$.¹⁵ Second, we assume that the exogenous forcing variables $\{G_t/z_t N_t\}$, $\{\Phi_t/z_t N_t\}$, $\{N_{t+1}/N_t\}$, $\{z_{t+1}/z_t\}$ and $\{p_{Et}\}$ are each stationary, even though $\{z_t\}$ and $\{N_t\}$ are only difference-stationary. Given these

¹⁵The family of utility functions U with this property is discussed in King, Plosser and Rebelo (1988).

assumptions, our equilibrium conditions can all be written, in a time-invariant form, in terms of suitably detrended endogenous variables, such as $\tilde{Y}_t \equiv \frac{Y_t}{z_t N_t}$, and thus admit a stationary solution in terms of the detrended variables.

The economy's steady state growth path is a set of constant values for the detrended variables (\tilde{Y}_t , etc.) that satisfy the transformed equilibrium conditions in all periods. We do not need to solve the equations for these steady state values, as we are not interested in explaining them in terms of more fundamental determinants, and indeed we do not bother to specify the global properties of the utility and production functions, that determine these steady state values. Our concern is rather with the model's predictions regarding the co-movements of the percentage deviations of the detrended variables from their steady state growth path; we intend to compare these predictions with the observed percentage deviations of the corresponding variables from their trend growth paths, as reported in section 1. The numerical values assumed for the coefficients of the linearized equilibrium conditions (written in terms of percentage deviations of the various detrended state variables from their steady-state values) are crucial for this. These coefficients are all functions of various shares and elasticities, evaluated at the steady-state values of the detrended state variables.

The coefficients of the log-linearized conditions are all functions of the model parameters listed in Table 2. We first discuss the parameters relating to the production function. The only properties of the function G that matter for the log-linearized equilibrium conditions are the steady-state value of s_E/s_M , where s_E and s_M denote the respective shares of energy and materials in the value of gross output, and ϵ_{EM} , the elasticity of substitution between energy and materials inputs, also evaluated at the steady-state factor mix. Similarly, the only relevant properties of the functions F and Q are summarized by the steady-state values of four more parameters: the ratio of labor costs to capital costs, the ratio of intermediate input costs to the cost of value added, and the two elasticities of substitution ϵ_{KH} and ϵ_{VG} .

The log-linearization of the transformed (9) involves coefficients that depend also upon the steady state value of Φ/V , the ratio of fixed costs to value added. Hence this ratio is another parameter that must be given a numerical value; it indicates the degree to which

there are increasing returns in the production of value added. As in Rotemberg and Woodford (1992), which we follow in the calibration of several parameters, we assume that this ratio takes the value required for zero pure profits in the steady state.¹⁶ We thus assume that

$$\frac{\Phi}{V} = \frac{\mu - 1}{1 - \mu(s_E + s_M)} \quad (19)$$

where μ denotes the ratio of price to marginal cost in steady state. As a result, the ratio Φ/V is not listed among the parameters that must be fixed independently in Table 2. Equation (19) also implies that the steady-state shares of the various factor costs in the value of gross product are equal to their shares in total costs. Thus the share ratios just referred to are all derivable from the shares s_H , s_K , s_E , and s_M listed in Table 2, and the latter quantities represent only three independent parameters, as they must sum to 1.

We assign numerical values to the six independent production function parameters as follows. In the U.S., the value of oil inputs is at most 4 percent of total value added. Value added in the mining of oil amounts to 1.8% of GDP on average. Imports of crude petroleum, mineral fuels, and lubricants are another 1.6% on average. Thus the value of oil inputs is about 3.4% of total value added. Even if one counts other energy inputs that might be thought to be close substitutes for oil (so that their prices increase to a similar extent, relative to other goods, when oil prices increase), the figure does not become much larger; for example, mining of coal amounts to only 0.4% of GDP. Hence $s_E/(1 - s_E - s_M)$ should equal approximately .04.

Next, we assume that materials constitute 50% of costs. This is less than the 60% share indicated by the Berndt and Wood (1979) data for the U.S. manufacturing sector. We have used a slightly lower number on the grounds that many service sector industries appear to have lower materials requirements. Thus we set s_M equal to 0.5 and s_E equal to 0.02. It follows that materials and energy, together, account for 52% of costs. We then obtain the labor and capital shares by assuming, as in Rotemberg and Woodford (1992), that labor

¹⁶This is presumably ensured by entry decisions over the long run, not explicitly modeled on the assumption that they are of little importance for short-run dynamics. Note that it requires that fixed costs grow at the same trend rate as output, presumably through an increase in the number of firms at a constant scale of operation. See Rotemberg and Woodford (1995) for further discussion.

accounts for 75% of value added. Thus s_H and s_K are set equal to 0.36 and 0.12 respectively

We follow the real business cycle literature in assuming that the elasticity of substitution ϵ_{KH} between capital and labor in the production of value added equals 1. The other two elasticities of substitution are less often considered. We consider two different values for each. One possibility is to assume very little opportunity for substitution away from either materials or energy inputs (which we represent by making both elasticities equal .0001). This is suggested by the estimates of Berndt and Wood (1979). On the other hand, other studies, such as Pindyck and Rotemberg (1983), suggest some degree of substitutability. We do not attempt to use their parameter values, as their production function specification is not consistent with the one we assume here. Instead, we have directly estimated the elasticities of substitution ϵ_{VG} and ϵ_{EM} under assumptions consistent with our model specification, using data for 20 two-digit U.S. manufacturing sectors. This estimation is described in the Appendix. Based on these estimates, our values for the two elasticities are $\epsilon_{EM} = 0.18$ and $\epsilon_{VG} = .69$.

As shown in Rotemberg and Woodford (1992), the homogeneity assumptions described above imply that all the aspects of preferences that matter for our analysis can be described by the two parameters ϵ_{Hw} and the σ . The former is a measure of the response of labor supply to a temporary real wage change, that is sometimes called the “intertemporal elasticity of labor supply” in the labor literature (e.g., Card, 1994). The latter parameter (introduced above in our statement of our homogeneity assumptions) corresponds to the elasticity of consumption growth with respect to changes in the real rate of return, holding constant hours worked in both periods. We assume an intertemporal elasticity of labor supply of 1.3, and a σ of 2. These values follow Rotemberg and Woodford (1992), except that σ has been reduced from 3 to 2, for closer conformity with the type of preferences assumed in the real business cycle literature.¹⁷

¹⁷Our finding that the competitive version of the model does not predict output and real wage declines as large as those we measure is robust to variation in these values.

Other parameters that enter the linearized equilibrium conditions include the steady-state shares of consumption, investment, and government purchases in private value added, the steady-state growth rate of real output, the steady-state real rate of return, and the rate of depreciation of the capital stock. All of these quantities have direct correlates in the national income accounts, and so we calibrate them by assuming that the steady-state values coincide with average values for the U.S. over the postwar period. Similarly, we use employment data to calibrate the steady-state share of hours hired (or conscripted by) the government. The values used here again follow Rotemberg and Woodford (1992), where they are discussed in more detail.

Finally, we must specify three parameters relating to the equilibrium behavior of markups. The first is the steady-state value of the markup itself, μ . The value that we use for all of our imperfectly competitive models is $\mu = 1.2$. This implies that for the typical firm, price is 20% higher than marginal cost, while fixed costs account for one-sixth of its total costs. As discussed in Rotemberg and Woodford (1995), this is well within the range of values for both market power and increasing returns indicated by a number of studies of U.S. industry. The second is the elasticity of the function $\mu(X/Y)$, evaluated at the steady-state value of X/Y . This parameter is the one that distinguishes our several imperfectly competitive models. In the static monopolistic competition model, as in the competitive model, ϵ_μ is zero, as the markup is constant. For the customer market model we let ϵ_μ equal -1 while we assume that it equals .15 for the implicit collusion model. (We must assume a positive value less than .2 in the latter case, for theoretical consistency, as explained in Rotemberg and Woodford, 1992.) The third parameter that must be calibrated is the α appearing in equation (18). (Note that this parameter matters only in the case of the two variable-markup models.) Here we follow Rotemberg and Woodford (1992) in setting $\alpha = .9$.

When we present our simulations, we report the response of real private value added and of a real wage to energy price shocks. These simulated response are intended to correspond to the responses of the time series that are measured by the U.S. Department of Commerce. Thus, we do not report the simulated responses of the ideal value added index V_t , or of the real

wage w_t (deflated by the price of output), nor of the wage deflated by the ideal value-added deflator. The available Commerce Department series differ from these constructs because there exists a domestic energy sector. Lack of disaggregated data at quarterly intervals prevents us from removing this sector from our measures of private value added and of the value-added deflator. We thus report the theoretical responses of variables which, like those measured by the Commerce Department, include a domestic energy sector.

We suppose that nominal value added is

$$Y_t - p_{E_t}E_t - M_t + p_{E_t}E_t^d$$

where E_t^d represents domestic energy production. Even though we assume that all income is received by a single representative household, we are free to assume that for accounting purposes, not all of the energy inputs used are treated as part of domestic output. Real value added is instead

$$Y_t - p_{E_0}E_t - M_t + p_{E_0}E_t^d$$

where time 0 is the base period for the GDP accounts. We furthermore assume that the fraction of energy inputs that are counted as domestic production is a constant, *i.e.*, that

$$E_t^d = s_D E_t \tag{20}$$

for some fraction $0 \leq s_D \leq 1$. This equation is somewhat arbitrary (since we do not here model the production or pricing decisions of energy producers explicitly). In the simulations reported, we set $s_D = .5$, as this represents the approximate share of U.S. oil usage that is domestically produced. Note that the specification (20) probably overstates the negative effects on U.S. energy production of a reduction in U.S. energy demand due to a price increase. Thus our results probably exaggerate the extent to which the models (both competitive and imperfectly competitive) predict a reduction of U.S. private value added following such a shock.

The corresponding private value added deflator is defined by dividing the above nominal value added measure by the real measure. Hence the real wage plotted in the figures is not

w_t , but rather

$$\frac{w_t(Y_t - (1 - s_D)p_{E0}E_t - M_t)}{Y_t - (1 - s_D)p_{Et}E_t - M_t}$$

5 Results of Simulations

Figures 2a and 2b display the theoretical responses of output and the real wage respectively, under the parameter values just discussed.¹⁸ As with all the results we will present, the response is calculated for ten quarters following a unit innovation in $\hat{\pi}_{Et}$. The predictions of four theoretical models are compared: the competitive model, the static monopolistic competition model, the customer market model, and the implicit collusion model. In these figures we also reproduce the estimated impulse responses (with confidence bands) from Figures 1b and 1c for purposes of comparison.

In Figure 2a, we show the predicted response of private value added. For our parameter values, the competitive model does predict a contraction of output following a positive innovation in oil prices. However, this contraction is much smaller than is indicated in Figure 1b. Output never falls by more than .06 percent in response to a one percent innovation in oil prices, which is only one-fourth of the effect that we estimate for quarters 5-7 following the innovation. Consistent with our heuristic discussion of section 2, the predicted response for the competitive model lies above the +2s.e. boundary of the confidence band in quarters 4 through 9.

The competitive version of our model also fails to predict that the decline in output should be significantly greater in the second year than in the quarters immediately following the impact. This means that the erosion of the capital stock following an energy price increase does not substantially increase the predicted output decline. Hence, to a useful approximation, the predicted effect of an oil shock in the competitive model can be determined in a framework where the capital stock is treated as given (as in our informal discussion in section 2). It also suggests that our oversimplified treatment of investment demand, ab-

¹⁸We also considered simulations where we varied the assumed values of the elasticities of substitution ϵ_{EM} and ϵ_{VG} . These variations had only a small effect on our results.

stracting from adjustment costs of any kind, is probably innocuous, at least for our analysis of the competitive model.

Imperfectly competitive models are able to account for a more severe contraction. Simply assuming a constant markup of 1.2 results in a predicted output decline of -.13 by quarters 5-8, which is twice as large as the one we obtain when $\mu = 1$. The implicit collusion model with $\mu = 1.2$ and $\epsilon_\mu = .15$ predicts an even larger decline, more than -.20 from quarter 5 onward. The allowance for endogenous markup variation thus makes the maximum output contraction 50% larger without any change in the assumed steady-state markup, and makes it comparable to the estimated decline.

The model with implicit collusion implies larger output declines because it predicts an increase in markups. Markups rise for two reasons. First, the increased price of energy inputs lowers the return to capital. In the event of a permanent increase in energy prices, the equilibrium capital stock would eventually fall as a result, but in the transition period, real interest rates would be lower than normal. As a result, the present value of future profits increases. This raises X_t/Y_t and, as a result, markups are higher until the capital stock adjusts and the real rate returns to its steady-state value.

Second, as is clear from the estimates in Table 1 and as is shown in Figure 1a, a shock to energy prices is generally followed by further increases in nominal and real energy prices. Starting around six quarters after the shock, real energy prices are expected to decline back to their usual value. These expected declines further increase X_t/Y_t at that time. The reason is that they imply that sales at that point are low and production costs are high relative to the values these variables are expected to have in the future. This means that the temptation to undercut the implicitly collusive agreement at the risk of a future breakdown in collusion is unusually low, and the degree of collusion that can be sustained is accordingly unusually high. Thus this model correctly predicts that the main contraction of output should occur only in the second year following the innovation, since it is at this time, when real energy prices are not only high but are also expected to decline, that X_t/Y_t is significantly above

its steady-state value.¹⁹

The customer market model, by contrast, predicts a larger immediate contraction than does the constant-markup model, but less of a contraction in the second year. This is again because X_t/Y_t rises in the second year, which in this model implies markup reductions, as firms sacrifice current profits to compete more vigorously for their future customer base. Thus the assumption of customer markets results in a less successful prediction, even compared to the static model of monopolistic competition. Given our parameter values, the implicit collusion model is the only one whose predicted path for output is always within the confidence band.

The competitive version of our model has particular difficulty in explaining the observed decline in real wages following an oil price increase. In the case of the wage deflated by the Commerce Department's value-added deflator (Figure 2b), the competitive model predicts a very small real wage decline, only a fourth of the estimated decline in the second year. This decline is entirely due to the domestic production energy, which raises the value-added deflator. As in the previous figure, this model predicts little additional decline in the second year (the decline by the middle of the second year is only about a third larger than the decline that has occurred by the second quarter following the innovation). The predicted path of wages is above the +2s.e. boundary of the confidence band in each of the quarters 4-7.

We find that a higher μ alone, regardless of any markup variation, helps to explain the real wage decline. The static model with a constant markup $\mu = 1.2$ implies that the value added deflated wage eventually falls by -.06 percent for each percent increase in the price of oil. This response is inside, but near the edge of, the two standard error confidence band from the estimated response. The implicit collusion model predicts an even greater decline. Indeed, in the case of $\mu = 1.2$, $\epsilon_\mu = .15$, the predicted decline is even slightly greater than the estimated response in the second year. Furthermore, this model again predicts a much

¹⁹In fact, the model predicts that markups actually fall in the quarter of the innovation, preventing any output decline at all. This is because X_t/Y_t falls, due to the expectation of even higher energy prices in the next several quarters, despite the first effect mentioned, which raises X_t/Y_t even in the first quarter.

sharper decline in the second year than in the first, so that the predicted path of wages tracks the estimated path reasonably well. The customer market model, by contrast, predicts the lowest real wage in the quarter of the innovation, with wages gradually returning to normal thereafter. Thus the implicit collusion model again best matches the estimated response, although the predictions of all three imperfectly competitive models are within the confidence band in this case.

6 Conclusion

We have shown that imperfectly competitive models, and in particular a model involving implicit collusion in the product market, can explain the estimated effect of oil price increases on output and real wages to a much greater extent than can a stochastic growth model that assumes a perfectly competitive product market. In this conclusion, we briefly discuss how our theory relates to other simple aggregative models which seek to explain for the output reductions that followed oil price increases.

It is sometimes argued that the recessions following the oil shocks of the 1970's were actually due to the tightening of monetary policy on these occasions, rather than an effect of the higher oil prices themselves. (See, e.g., Darby (1982), Bohi (1989).) From this point of view, our development of a non-monetary model with imperfect competition might seem to be unnecessary, and our analysis of the competitive case with no allowance either for a feedback rule for monetary policy or for nominal rigidities misleading. We cannot engage at this point in a complete discussion of models where money has important effects. But it does seem to us that models where monetary policy matters cannot avoid our conclusions, at least without adding considerable complications.

Suppose that over our sample period, oil price increases did lead systematically to reduced growth of the money supply over subsequent quarters. Suppose furthermore that one were to model the real effects of changes in monetary policy by postulating imperfectly indexed nominal wage contracts. In this case the neoclassical labor supply curve would be replaced by a perfectly elastic labor supply at a real wage that depends upon the nominal price level

for non-energy output. In this case, an unexpectedly low money supply, and consequently an unexpectedly low nominal price level, would result in a contraction of employment and output. But a condition like (5) would still apply (in the case that firms are perfect competitors in their product markets), and this contraction would occur only insofar as the real wage divided by the ideal deflator for value added *rose*. Thus it is hard to see how the hypothesis of a coincident monetary tightening could explain the sharp *decline* in real wages that accompanies the observed contraction of output.

If one supposes that the real effects of monetary policy are instead due to nominal price rigidity, one probably has to consider models in which product markets are imperfectly competitive in the first place (as in Rotemberg, 1982, and Blanchard and Kiyotaki, 1987). This still might seem an alternative to the particular kinds of imperfectly competitive models developed in this paper. However, as Blinder (1981) and Rotemberg (1983) have observed, models with sticky prices generally do not imply that oil price increases should have a large effect on output. Instead, the price stickiness should buffer the economy from such shocks, by comparison with what would happen in the case of flexible prices, since markups should be squeezed at a time of sharply rising nominal marginal costs. Thus such an explanation would require one to argue not simply that the monetary contraction adds to the contractionary impact of the oil shock, but that the monetary contraction is really the whole story, since the oil price increase alone would have had little effect on output at all. A large enough monetary contraction could certainly produce effects upon output and real wages as large as those we estimate; but it remains unclear why (given the small effect of the oil price shock upon costs and hence upon inflationary pressures) such a large monetary contraction should follow oil price increases.

A leading alternative hypothesis, of course, is that the aggregate effects of energy price increases depend crucially upon the fact that such shocks affect different sectors differently. Among this class of explanations, one must mention the sectoral reallocation model of Hamilton (1988), as well as the sticky-price model of Ball and Mankiw (1992). A quantitative comparison of explanations of this kind with the one offered here must be left for further

research.

7 Bibliography

- Andrews, Donald W. K., "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 61, July 1993, 821-56.
- Ball, Laurence, and N. Gregory Mankiw, "Relative-Price Changes as Aggregate Supply Shocks," *NBER Working Paper n. 4168*, September 1992.
- Basu, Susanto, "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 85, 1995, 512-531.
- Berndt, Ernst, and David Wood, "Engineering and Econometric Interpretations of Energy-Capital Complementarity," *American Economic Review*, 69, 1979, 342-354.
- Blinder, Alan S., "Monetary Accommodation of Supply Shocks under Rational Expectations," *Journal of Money, Credit and Banking*, 13, November 1981, 425-38.
- Bohi, Douglas R., *Energy Price Shocks and Macroeconomic Performance*, Washington: Resources for the Future, 1989.
- Bruno, Michael, and Jeffrey Sachs, *Economics of Worldwide Stagflation*, Harvard University Press, Cambridge, 1985.
- Burbidge, John, and Alan Harrison, "Testing for the Effects of Oil-Price Rises Using Vector Autoregressions," *International Economic Review*, 25, 1984, 459-484.
- Card, David, "Intertemporal Labor Supply: An Assessment," in C.A. Sims, ed., *Advances in Econometrics: Sixth World Congress*, Cambridge: Cambridge University Press, 1994.
- Carruth, Alan A., Mark A. Hooker, and Andrew J. Oswald, "Unemployment Equilibria and Input Prices: Theory and Evidence from the United States," mimeo, University of Kent, January 1995.
- Darby, Michael, "The Price of Oil and World Inflation and Recession," *American Economic Review*, 72, 1982, 738-751.

- Finn, Mary G., "Energy Price Shocks, Capacity Utilization and Business Cycle Fluctuations," *Institute for Empirical Macroeconomics Discussion Paper no. 50*, Federal Reserve Bank of Minneapolis, September 1991.
- Gisser, Micha, and Thomas H. Goodwin, "Crude Oil and the Macroeconomy: Tests of Some Popular Notions," *Journal of Money, Credit and Banking*, 18, 1986, 95-103.
- Gordon, Robert J., "Supply Shocks and Monetary Policy Revisited," *American Economic Review*, 74, May 1984, 38-43.
- Hall, Robert E., "The Relation Between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, 96, October 1988, 921-48.
- , "Invariance Properties of Solow's Productivity Residual," in P.A. Diamond, ed., *Growth / Productivity / Unemployment*, Cambridge, Mass.: M.I.T. Press, 1990.
- Hamilton, James D., "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91, 1983, 228-248.
- , "Historical Causes of Postwar Oil Shocks and Recessions," *The Energy Journal*, 6, 1985, 97-116.
- , "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy* 96, 1988, 593-617.
- , "This is What Happened to the Oil Price - Macroeconomy Relationship," Discussion Paper no. 95-36, Department of Economics, University of California - San Diego, September 1995.
- Keane, Michael P. and Eswar S. Prasad, "The Employment and Wage Effects of Oil Price Shocks: A sectoral Analysis," *Federal Reserve Bank of Minneapolis Institute for Empirical Macroeconomics Discussion Paper 51*, September 1991.
- Kim, In-Moo, and Prakash Loungani, "The Role of Energy in Real Business Cycle Models," *Journal of Monetary Economics*, 29, 1992, 173-89.
- King, Robert G., Charles I. Plosser and Sergio Rebelo, "Production, Growth and Business Cycles, I: The Basic Neoclassical Model," *Journal of Monetary Economics*, 21, 1988,

195-232.

- Mork, Knut A., "Oil and the Macroeconomy when Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 97, 1989, 740-744.
- Phelps, Edmund S., and Sidney G. Winter, "Optimal Price Policy under Atomistic Competition," in E. S. Phelps, ed., *Microeconomic Foundations of Employment and Inflation Theory*, W. W. Norton Co., New York, 1970.
- Pindyck, Robert S., and Julio J. Rotemberg, "Dynamic Factor Demands and the Effects of Energy Price Shocks," *American Economic Review*, 73, 1983, 1066-1079.
- Poterba, James M., Julio J. Rotemberg, and Lawrence H. Summers, "A Tax-Based Test for Nominal Rigidities," *American Economic Review*, 76, 1986, 659-675.
- Ramey, Valerie, "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, 99, 1991, 306-334.
- Rasche, Robert H., and Tatom, John A., "Energy Price Shocks, Aggregate Supply and Monetary Policy: The Theory and the International Evidence," *Carnegie-Rochester Conference Series on Public Policy*, 14, 1981, 9- 94.
- Rotemberg, Julio J., "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies*, 44, 1982, 517-531.
- , "Supply Shocks, Sticky Prices and Monetary Policy," *Journal of Money Credit and Banking*, 15, November 1983, 489-98.
- Rotemberg, Julio J., and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review*, 76, June 1986, 390-407.
- Rotemberg, Julio J. and Michael Woodford, "Markups and the Business Cycle," *NBER Macroeconomics Annual*, 1991, 63-128.
- and —, "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100, December 1992, 1153-1207.
- and —, "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets," in T.F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton:

Princeton University Press, 1995.

Appendix I

Production Function Parameter Values

Here we discuss our estimates of the elasticities of substitution ϵ_{EM} and ϵ_{VG} , that are used in the simulations. We estimate these elasticities under the assumption of perfect competition, because we are especially concerned to correctly calibrate the competitive model, the empirical inadequacy of which we document in this paper. The same parameter estimates are then used in all of the simulations, as we wish to display the consequences of variations solely in our assumptions about markup determination. (We note, however, that assumption of a significant departure from perfect competition ought to change our estimates of ϵ_{EM} and ϵ_{VG} as well.)

The elasticity ϵ_{EM} is defined as the coefficient in the log-linear approximation

$$\Delta G_{Et} - \Delta G_{Mt} = -(\Delta E_t - \Delta M_t)/\epsilon_{EM}$$

Here ΔX denotes demeaned first difference of $\log X$, for each of the variables. (Because our model implies that both E/M and G_E/G_M are stationary variables, such a log-linear approximation should be valid in the case of sufficiently small equilibrium fluctuations, for any smooth aggregator. If G is a CES function, of course, the log-linear relationship is exact.) Cost minimization by firms then implies that in equilibrium

$$\Delta p_{Et} - \Delta p_{Mt} = -(\Delta E_t - \Delta M_t)/\epsilon_{EM} \quad (21)$$

This follows from equilibrium conditions (10) - (11) of the text, except that we do not assume that $p_{Mt} = 1$ because we analyze sectoral data.

Our data includes separate observations for materials and services inputs. To ensure that the estimation yields parameters for our theoretical model, we aggregate these two inputs into a single “materials” category Divisia aggregation. Thus

$$\begin{aligned} \Delta M_t &= \frac{s_{MEt}}{s_{MEt} + s_{MS_t}} \Delta M_{Et} + \frac{s_{MS_t}}{s_{MEt} + s_{MS_t}} \Delta M_{St} \\ \Delta p_{Mt} &= \frac{s_{MEt}}{s_{MEt} + s_{MS_t}} \Delta p_{MEt} + \frac{s_{MS_t}}{s_{MEt} + s_{MS_t}} \Delta p_{MS_t} \end{aligned}$$

where ΔM_{Et} and ΔM_{St} denote the growth rates for non-energy inputs that are materials and services respectively, Δp_{MEt} and Δp_{MS_t} are the growth rates for the corresponding price indices and s_{MEt} and s_{MS_t} are the corresponding cost shares. These aggregates are then used together with ΔE_t and Δp_{Et} to estimate (21).

The elasticity ϵ_{VG} is correspondingly defined as the coefficient in

$$\Delta Q_{Gt} - \Delta Q_{Vt} = -(\Delta G_t - \Delta V_t)/\epsilon_{VG}$$

If we assume that all factors are variable and that the firm is a price taker in each factor market, then cost minimization again implies

$$\Delta p_{Gt} - \Delta p_{Vt} = -(\Delta G_t - \Delta V_t)/\epsilon_{VG} \quad (22)$$

where price changes Δp_{Et} and Δp_{Vt} are again constructed as Divisia aggregates. However, we do not observe a rental price series for capital. Moreover, adjustment costs for capital (which, admittedly are neglected in our theoretical model) create a wedge between “cost of capital” series constructed along Jorgensonian lines and the current marginal product of capital.

An alternative estimating equation is accordingly more convenient. Equation (22) implies

$$\Delta s_{Gt} - \Delta s_{Vt} = -(\Delta G_t - \Delta V_t) \left(1 - \frac{1}{\epsilon_{VG}}\right)$$

where s_{Gt} and s_{Vt} are the cost shares of intermediate inputs and value added respectively. Furthermore, under the assumption of perfect competition, we can replace the cost shares s_{Gt} and s_{Vt} by the shares of these input costs in the value of total output, and still write $s_{Vt} = 1 - s_{Gt}$, so that the above equation becomes

$$\left(1 - \frac{1}{s_{Gt}}\right) \Delta s_{Gt} = (\Delta G_t - \Delta V_t) \left(1 - \frac{1}{\epsilon_{VG}}\right) \quad (23)$$

We use this equation to estimate ϵ_{VG} using $s_{GT} = s_{Et} + s_{Mt}$.²⁰ The quantity growth rates ΔG_t and ΔV_t are constructed as Divisia indices

$$\Delta G_t = \frac{s_{Et}}{s_{Et} + s_{Mt}} \Delta E_t + \frac{s_{Mt}}{s_{Et} + s_{Mt}} \Delta M_t$$

²⁰Because of the use of cost shares in the value of total output, e.g., $s_{Et} = p_{Et}E_t/Y_t$, rather than shares in total cost, (23) is correct only under the assumption of perfect competition.

$$\begin{aligned}\Delta V_t &= (1 - s_{Et})^{-1}[\Delta Y_t - s_{Gt}\Delta G_t] \\ &= (1 - s_{Et} - s_{Mt})[\Delta Y_t - s_{Et}\Delta E_t - s_{Mt}\Delta M_t]\end{aligned}$$

while ΔM_t is constructed as indicated above.²¹

We estimate equations (21) and (23) using the KLEMS data for 20 two-digit U.S. manufacturing sectors supplied by BLS Division of Productivity Research. We impose common elasticities on these 20 sectors to obtain relatively precise estimates that we can use in the calibration of our symmetric model. We have also examined independent sectoral regressions, and found qualitatively similar results for most sectors, but with large standard errors for the coefficient estimates.

We use the cumulative changes over two years for the growth rates appearing in those equations. We construct these two-year changes by summing the annual changes for two consecutive years, where the annual changes are computed as indicated above. (We have data for 17 such periods, from 1950/51 through 1987/88.) Two-year growth rates are used because adjustment of the factor mix to relative price changes appears not to occur entirely within a single year.²² Since our simulation exercise aims to explore the effects on the economy that occur during the first two years following an innovation in energy prices, we seek a medium term elasticity of substitution rather than one that is valid only for the first four quarters. Indeed, our figures show that the largest effects of energy price increases occur in the second year after the shock.

Finally, we allow for the possibility of stochastic variation in the aggregator functions Q and G , which would add error terms to equations (21) and (23). Hence, we de-trend (as well as de-meaning) all of our growth rates, and we estimate (21) and (23) with an instrumental variable estimator. The instrument is the growth of nominal oil prices over the same two-year period. As discussed in the text, we regard this as a largely exogenous process, and so

²¹Even though it is the Divisia version of the standard deflator of “value added”, the second of these equations, is again valid only under the assumption of perfect competition. The reason is, again, that we use cost shares.

²²When we experimented with one year changes, we found the results much more sensitive to the normalization of the second stage regression because the instrument is much poorer.

expect it to be uncorrelated with stochastic shocks in the Q and G aggregators, just as we expect it to be uncorrelated with the labor-augmenting technical shock variable z .

Regression coefficients for regressions of the left and right hand side of (21) and (23) on the contemporaneous nominal oil price changes are given below.

Dependent Variable	Regression Coefficient	Standard Error
$\Delta p_{E_t} - \Delta p_{M_t}$.259	(.030)
$\Delta E_t - \Delta M_t$	-.046	(.034)
$(1 - s_{G_t})^{-1} \Delta s_{G_t}$.122	(.021)
$\Delta G_t - \Delta V_t$	-.267	(.044)

We observe that the proposed instrument is a statistically significant predictor of all of the changes that we are interested in, except $\Delta E_t - \Delta M_t$. Even if this particular low t-statistic indicates that this is a poor choice of instrument (i.e., one not really correlated with the shifts we are interested in, and correlated with $\Delta p_{E_t} - \Delta p_{M_t}$ for accidental reasons), then our use of it can bias our estimate of ϵ_{EM} towards zero.

The first-stage regressions just reported imply IV estimates of $\epsilon_{EM} = .177$ and $\epsilon_{VG} = .686$. These are the baseline values used in the paper. Using the change in the price of oil over the previous two-year period as an additional instrument has no material consequences on these results. ²³

²³In general, the use of several instruments implies that the results depend on the side of the equation that projected against the instruments in the first stage. In the case of ϵ_{EM} , the resulting estimate is 0.184 no matter which side of (21) is projected on the instruments. The estimate of ϵ_{VG} equals 0.659 if $(1 - s_{G_t})^{-1} \Delta s_{G_t}$ is projected on the instruments while it equals 0.670 if $\Delta G_t - \Delta V_t$ is projected. We prefer the estimates with only one instrument because the lagged oil price change does not result in a significant coefficient in any of the first-stage regressions, except that for $\Delta p_{E_t} - \Delta p_{M_t}$.

Table 1
The Effects of Changes in Nominal Oil Prices 47:2-80:3

Explanatory Variable	Nominal Oil Price (π_{Et})	Real Oil Price (p_{Et}^r)	Output	Real Wage
Constant	0.075 (0.07)	-0.002 (0.01)	1.03 (0.3)	0.122 (0.03)
Trend	6.2e-5 (1e-6)	-6.5e-5 (2e-5)	1.2e-3 (3e-4)	1.9e-4 (7e-5)
Own First Lag			1.08 (0.08)	0.771 (0.089)
Own Second Lag			-0.22 (0.08)	0.163 (0.086)
π_{Et}		0.98 (0.02)	-0.04 (0.03)	-0.0312 (0.016)
π_{Et-1}	1.51 (0.46)	-0.32 (0.08)	0.45 (0.16)	-0.133 (0.016)
π_{Et-2}	0.15 (0.09)	-0.06 (0.02)	-0.03 (0.03)	-0.0160 (0.016)
π_{Et-3}			-0.07 (0.03)	0.0068 (0.015)
π_{Et-4}			-0.08 (0.03)	-0.0214 (0.015)
p_{Et-1}^r	-1.22 (0.47)	1.32 (0.08)	-0.44 (0.16)	
p_{Et-2}^r	1.19 (0.47)	-0.32 (0.08)	0.43 (0.16)	

Standard errors in parenthesis.

Table 2
The Calibrated Parameters

Parameter	Defined by	Values	Description
g	$\gamma_2 \gamma_N - 1$	0.008	Steady state growth rate (per quarter)
s_C	$\frac{\bar{C}}{Y-M}$	0.697	Share of private consumption expenditure in $Y - M$
s_G	$\frac{\bar{G}}{Y-M}$	0.117	Share of government purchases of goods in $Y - M$
s_I	$(g + \delta) \frac{\bar{K}}{Y-M}$	0.186	Share of private investment expenditure in $Y - M$
s_D	E_t^d / E_t	0.5	Share of energy that is domestically produced
δ		0.013	Rate of depreciation of capital stock (per quarter)
s_E	$\frac{Q_G G_E \bar{E}}{\mu Q}$	0.02	Share of energy costs in total costs
s_M	$\frac{Q_G G_M \bar{M}}{\mu Q}$	0.5	Share of materials costs in total costs
s_H	$(1 - s_E - s_M) \frac{F_H \bar{H}^p}{F}$	0.36	Share of labor costs in total costs
s_K	$(1 - s_E - s_M) \frac{F_K \bar{K}}{F}$	0.12	Share of capital costs in total costs
θ^g		0.17	Share of hours hired by the government
r	$\frac{F_K Q_V}{\mu} - \delta$ or $\gamma_2^g \beta^{-1} - 1$	0.014	Steady state real rate of return (per quarter)
ϵ_{KH}	$\frac{F_K F_H}{F_{KH} F}$	1	Elasticity of substitution between capital and hours
ϵ_{VG}	$\frac{Q_V Q_G}{Q_V G Q}$	0.69, 0.0001	Elasticity of substitution between value added and G
ϵ_{EM}	$\frac{G_E G_M}{G_{EM} G}$	0.18, 0.0001	Elasticity of substitution between energy and materials
$1/\sigma$		0.5	Elasticity of consumption growth with respect to real return holding hours worked constant
ϵ_{Hw}		1.30	Intertemporal elasticity of labor supply
μ		1, 1.2	Steady state markup (ratio of price to marginal cost)
ϵ_μ	$\frac{\bar{X} \mu'}{Y \mu}$	-1, 0, 0.15	Elasticity of the markup with respect to X/Y
α		0.89	Expected rate of growth of market share




FIGURE 1a




FIGURE 1b




FIGURE 1c

FIGURE 2a




FIGURE 2b

