

WILLIAM D. NORDHAUS

Yale University

Who's Afraid of a Big Bad Oil Shock?

WHEN THE U.S. INVADED Iraq in March 2003, many economists feared that the war would lead to a sharp decline in Iraqi oil production, a spike in oil prices, and a woeful U.S. economy that would follow the scripts of the oil shocks of 1973, 1978, and 1990. Real oil prices did increase, indeed more than tripled, from \$20 in 2001:Q4 to \$62 in 2006:Q3 (in 2007 dollars). But the ailments associated with earlier oil-price increases did not appear. Instead output grew rapidly, inflation was moderate, unemployment fell, and consumers remained reasonably happy.¹

Ftn. 1

Macroeconomists would be out of business if there were no surprises. The business of this paper is to inquire into the explanations for the surprising oil noncrisis of the early to mid-2000s. The robustness of the economy following the latest oil shock can perhaps be seen in the context of an important historical development in macroeconomics, namely, the Great Moderation. Over the past half-century, the economy has shown declining volatility of inflation, unemployment, and output growth.² Perhaps the moderated response of the economy to the latest oil shock should be understood as part of this overall decline in volatility.

Ftn. 2

Although much has been written about the macroeconomic impacts of price shocks, little analysis is available on the impact of the most recent oil price shock.³ Most research focuses on the role of monetary policy and

Ftn. 3

1. For data and sources as well as an elaboration of the quantitative statements in this paper, see "Accompanying Notes on Data and Methods for 'Who's Afraid of the Big Bad Oil Shock?'" September 15, 2007, available at www.econ.yale.edu/~nordhaus/homepage/Data_BPEA_20072_v2.pdf.

2. See, in particular, Stock and Watson (2003).

3. The most comprehensive study is Blanchard and Galí (2007).

inflation dynamics in the wake of a shock. There seems to be no consensus as to whether the macroeconomic fear of oil-price shocks remains warranted.

Comparing Shocks Past and Present

To begin, I define an oil-price shock as an inward shift in the supply curve for crude oil triggered by political events exogenous to the oil market and the macroeconomy. The four oil-price shocks considered here, summary statistics for which are shown in table 1, are associated with the 1973 Arab-Israeli war, the 1978 Iranian revolution, the 1990 Iraqi invasion of Kuwait, and the 2002 run-up to the American invasion of Iraq. Other periods might be considered, such as the price blip from 1998 to 2000, or the recovery from the OPEC price-cutting war of 1986, but these probably reflected cartel dynamics and were not clearly triggered by exogenous political events. Tab. 1

The first date shown for each event in the table is that when the oil shock began, and the second is that when oil prices peaked. The first numerical column shows the corresponding percentage increase in the real price of oil from start to peak. (The oil price measure used here is the average refiner acquisition cost of crude oil, except during the period of price controls, 1973:Q3–1982:Q4, when it is the import cost. The oil price is deflated by the price index of personal consumption expenditures, or PCE.) The most recent event involved a real oil-price increase of 125 percent, close to the size of the 1978–81 increase but smaller than that in 1973–75.

The next two columns show the macroeconomic response in terms of the unemployment rate and output (the latter expressed as the ratio of actual to potential GDP.) The latest shock did not appear to have the recessionary impact on the economy seen in the three earlier shocks. The unemployment rate actually fell by slightly more than 1 percentage point, and output grew slightly relative to potential. This contrasts sharply with the shocks of the 1970s, when the unemployment rate rose by between 1½ and 3½ percentage points and output declined sharply relative to potential.

Although the latest oil shock did not lead to a recession, the impacts on inflation and productivity were in the same direction as those from the shocks of the 1970s. But the impacts were different in magnitude: inflation

Table 1. Changes in Macroeconomic Indicators during Oil-Price Shocks
 Percentage points except where stated otherwise

<i>Quarter of initial oil-price shock</i>	<i>Quarter of oil-price peak</i>	<i>Real oil price (percent)</i>	<i>Change from initial shock to peak in</i>				<i>Labor productivity growth^b</i>
			<i>Unemployment rate</i>	<i>Actual output as percent of potential</i>	<i>PCE inflation^a</i>	<i>CPI inflation^a</i>	
1973:3	1975:4	191	3.5	-6.4	4.3	4.4	-1.3
1978:4	1981:1	111	1.5	-2.7	3.2	4.5	-0.6
1990:3	1990:4	27	0.4	-1.4	0.2	0.3	-0.6
2002:4	2006:2	125	-1.2	1.9	1.2	1.4	-0.6

Source: Author's calculations using data from DRI-WEFA Basic Economics Database and the Energy Information Administration.

a. Annual logarithmic change in the relevant price index over two years.

b. Annual logarithmic change in productivity averaged over three years.

during the 2002–06 shock rose about $1\frac{1}{4}$ percentage points—much less than the sharp rise in the 1970s, and labor productivity growth declined by about $\frac{1}{2}$ percentage point, less than the average in the 1970s as well.

In short, table 1 reveals that the most recent oil shock is qualitatively more similar to earlier shocks than might at first appear. Inflation rose and productivity fell, as they did in response to earlier shocks. The major difference is that the latest oil shock was followed by a continued economic expansion rather than a recession.

Why Are Oil-Price Shocks Contractionary?

The first challenge is to understand why oil-price shocks might reduce real output in the first place. There are basically two sets of reasons: effects on productivity and effects on aggregate demand.

The effect of any energy-price shock on productivity involves substitution of other factors for energy in response to the higher energy price. In a world in which technological change is exogenous, there will actually be *no* response of multifactor productivity to price changes; the only response is substitution along a given technological frontier. However, if the focus is on partial productivity measures, such as labor productivity, then there might be a response, such as when labor is substituted for energy.

It is easy to forget how small the substitution response would be. To begin with, in a full-employment world, there would be no productivity response in the short run until substitution occurs. If people continued driving, flying, and using existing capital equipment at the same rate as before the price shock, all inputs and outputs would be identical, and there would consequently be no productivity impact.

Even after substitution occurs, the effect would be quite limited, as can be seen in an illustrative calculation. In a simplified, one-product neoclassical world, the labor productivity response to an oil-price change is approximately proportional to the elasticity of output with respect to the oil price (taken to be 3 percent here) times the price elasticity of demand for oil times the logarithmic change in the oil price. To calculate the demand elasticity, I take 1973:Q1 to 1992:Q4 as a sample period in which short-run movements in oil prices were dominated by supply changes. For this period the elasticity is estimated to be -0.015 (± 0.0052) for the current quarter, -0.047 (± 0.033) for one year, and -0.098 (± 0.054) for ten

years. (These elasticities, along with inelastic supply, suggest why crude oil prices are so volatile in the short run.) Using these numbers, I estimate that the productivity effect of a doubling of real oil prices would be a decrease of 0.11 percent for one year and 0.04 percent a year for a ten-year horizon. This is about one-tenth the productivity impacts shown in table 1. These calculations suggest the need to look elsewhere to understand the association between oil-price changes and output reactions.

A more likely mechanism is that price shocks work through aggregate demand. They can do this in two important ways. The first is through monetary policy: rising energy prices produce inflation, which may lead the Federal Reserve to tighten monetary policy and slow the economy.⁴ To explain the different behavior of the economy following the recent oil price shock, one can hypothesize that the Federal Reserve reacted less sharply in the recent episode than to earlier energy-price-driven inflations. In particular, the central bank appears to be focusing recently not on total inflation but on core inflation, which excludes the direct effect of energy-price increases.

Ftn. 4

As a technical aside, core inflation does not remove indirect energy-price increases, that is, increases that are passed through to consumer goods and services indirectly through business purchases of energy such as in airfares and apartment rents. Thus the total impact of energy prices on consumer prices with dollar-for-dollar pass-through will exceed the direct impacts that are removed in calculating core inflation rates. A rough estimate is that the total impact on consumer prices today is about 1.8 times the direct effect. (This is calculated by estimating the indirect energy included in nonenergy goods and services using input-output coefficients.) In taking only direct energy-price effects out of the price index, therefore, the Federal Reserve is implicitly squeezing profit margins or wage increases sufficiently to remove the indirect impact of energy prices.

One can think of this process in terms of the Taylor rule, which holds that the federal funds rate is a function of the deviation of inflation and unemployment from their targets. If the current Taylor rule targets core inflation, whereas earlier Federal Reserve behavior targeted overall inflation, then monetary policy will not produce contractionary effects today until and unless higher energy prices get transmitted into a further round of price or wage increases.

4. This is the mechanism described in Bernanke, Gertler, and Watson (1997).

Table 2. Estimated Regression Coefficients from Taylor Rule and Aggregate Demand Equations^a

Variable	Period covered by estimate				
	1950–70	1960–80	1970–90	1980–2000	1987–2007
<i>Taylor rule equation^b</i>					
Energy price ^c	— ^d	0.97 (0.41)	0.31 (0.44)	0.12 (0.34)	0.11 (0.13)
<i>Aggregate demand equation</i>					
Exogenous spending	0.22 (0.19)	0.66 (0.29)	1.13 (0.34)	0.77 (0.29)	0.97 (0.21)
Oil-price shock	— ^d	–0.50 (0.30)	–0.19 (0.25)	–0.06 (0.22)	–0.11 (0.18)

Source: Author's regressions.

a. Rolling regressions; see the appendix for details of the specification. Numbers in parentheses are standard errors.

b. The dependent variable is the interest rate on federal funds.

c. Contribution of energy prices to core PCE inflation.

d. Equation not estimated for this period.

e. Variable not included in estimate for this period.

It will be helpful to see if the Federal Reserve behaves as it speaks. For this purpose I estimate a rolling-regression Taylor rule with separate terms for PCE inflation without energy prices and the contribution of energy prices to PCE inflation (see the appendix for the exact equation). The top panel of table 2 shows the estimated coefficients on the energy term for a succession of twenty-year periods. The energy-price contribution has a large coefficient during the first two oil shocks, but the coefficient is essentially zero in the 1980–2000 and 1987–2007 periods. This result shows that the Federal Reserve has indeed moved toward focusing on core inflation in its policy decisions. Therefore, oil prices are less contractionary today in part because monetary policy accommodates one-time direct shocks to energy prices. Tab. 2

The second mechanism, and one that was emphasized in the 1970s by members of the Brookings Panel, can be called the “oil price as tax increase” mechanism. The logic is that the marginal propensity to consume (MPC) of energy buyers is much higher than that of energy sellers. Buyers tend to be consumers, and disproportionately low-income consumers, because of the regressive nature of energy spending, and so have high MPCs. Sellers, by contrast, are corporations, high-income individuals, and OPEC countries, who in total have low marginal propensities to spend on U.S. production. The net effect of a permanent oil-price increase might therefore be almost as powerful as a permanent tax increase.

Measuring Oil Shocks

The second challenge in studying oil-price shocks is to determine an appropriate shock variable. The standard approach in studies of oil prices and the macroeconomy has been to use either nominal or real oil prices in dollars.⁵ This approach is clearly defective, as it does not scale oil prices for their economic importance. Ftn. 5

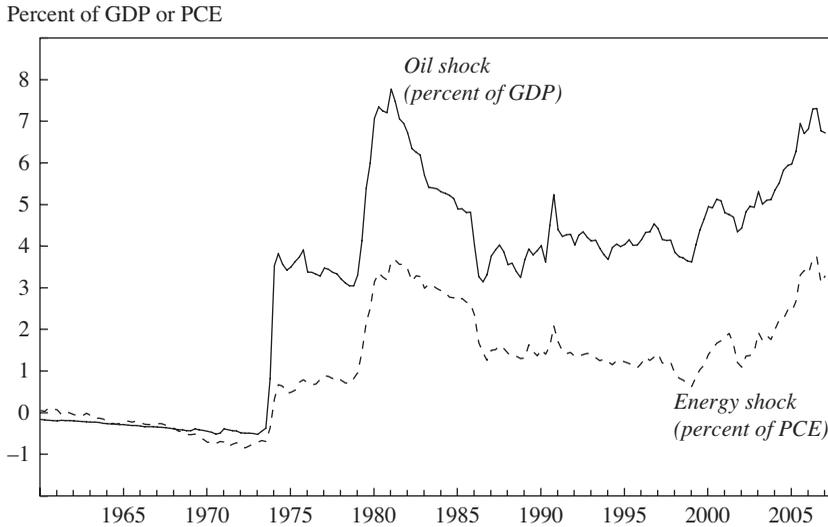
It is useful to go back to first principles in defining a shock variable. All three of the mechanisms discussed in the last section would suggest scaling price shocks by the overall share of oil or energy in the economy. In the production function approach, the elasticity of output with respect to oil or energy (usually measured by the income share) will be the key variable. Similarly, in the oil-price-as tax-increase theory, the impact of changing prices will come through the effect on real income, that is, through the standard income effect.

I therefore construct a shock variable as the value of oil purchases relative to total income times the change in the price. More precisely, I define the oil shock variable as the logarithmic change in price times the lagged share of oil purchases in GDP: $\text{OilShock}_t = [\ln(P_t/P_{t-1})](P_{t-1}X_{t-1}/\text{GDP}_{t-1})$, where P_t is the nominal price of crude oil and X_t is total domestic consumption of petroleum products. I then chain the series of quarterly shock variables to obtain a cumulative shock variable. This cumulative shock variable has the same dimension as a chained high-employment surplus as a percent of GDP.

This definition is incomplete because it considers only oil. I therefore define a second, broader energy-shock variable. This is conceptually the same as the oil-shock variable but uses the price and quantity of personal consumption expenditures on energy products (primarily gasoline, electricity, and heating fuels), with PCE itself rather than GDP as the denominator. The advantage of the energy-shock variable is that it includes all energy sources and focuses on the impact on consumers. The disadvantage is that it omits second-round effects, such as the effect of oil prices on jet fuel and thence on airfares.

Figure 1 shows the time paths of the cumulative energy- and oil-shock variables since 1960. Table 3 reports values of the shock variables both as cumulative totals and annualized over the shock periods. The table shows Fig. 1
Tab. 3

5. See Hamilton (1983) and Blanchard and Galí (2007) as examples.

Figure 1. Cumulative Oil- and Energy-Price Shocks, 1960–2007

Source: Author's calculations using data from the DRI-WEFA Basic Economics Database and the Energy Information Administration.

that the 2002–06 shock was about three-quarters as large as the 1970s shocks when measured by the cumulative “tax increase” but was only half as large when measured as the “tax increase per year.”

The historical context suggests that the first price shock had a far more powerful economic impact than later shocks. The 1970s shocks were dislocations that fundamentally changed Americans’ view of energy and resources. They gave rise to an “age of scarcity,” which replaced perceptions of seemingly bottomless oil wells and a promise of energy “too cheap to meter.” Given the abundance of oil and the historical stability of oil prices before 1973, the changes that ensued were almost unimaginable outcomes that stunned businesses, consumers, and governments alike. They were similar in psychological effect to the collapse of the World Trade Center towers or the inundation of New Orleans by Hurricane Katrina.

Figure 2 illustrates this graphically using a “surprise index” for real oil prices and (for comparison) stock prices. The surprise index is defined as the absolute value of the quarterly price change relative to the twenty-year moving average volatility of price changes (volatility is measured as the standard deviation of price changes, with all changes expressed as absolute

Fig. 2

Table 3. Impact of Oil Price Shocks on Output and Personal Consumption Expenditure
Percent

<i>Period</i>	<i>Cumulative (percent)</i>		<i>Annualized (percent a year)</i>	
	<i>Oil^a</i>	<i>Energy^b</i>	<i>Oil</i>	<i>Energy</i>
1973:3–1975:4	4.3	1.5	1.9	0.7
1978:4–1981:1	4.7	2.8	2.1	1.3
1990:3–1990:4	0.7	0.5	3.0	1.9
2002:4–2006:2	2.4	2.2	0.7	0.6

Source: Author's calculations using data from DRI-WEFA Basic Economics Database and the Energy Information Administration.

a. Percent of GDP.

b. Percent of PCE.

values of logarithmic differences). Circles identify the periods when the surprise was more than five moving standard deviations (“five sigmas”). Intuitively, a high level of surprise indicates that the price change was large relative to the historical norm of price changes.

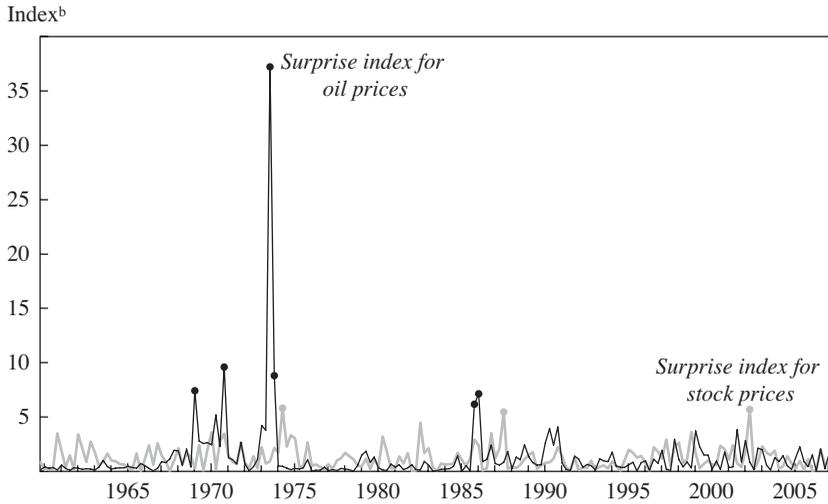
The surprise for oil prices peaked near forty sigmas in the 1973–74 shock and has never approached this level again since then. In contrast, during the 2002–06 shock the surprise index for oil barely reached three sigmas. Levels of surprise around five sigmas are the maximum observed for stock prices throughout the fifty-year period. In other words, the surprise index for oil prices was stratospherically high in the 1970s, but by the 2000s it was about the same as for a run-of-the-mill asset.

The mid-1970s spike in the surprise index reflects the major forecasting errors of that time. My recollection is that no serious forecaster in 1972 foresaw anything approaching the oil-price changes that occurred over the next decade. By contrast, the oil-price increase of the early to mid-2000s was large, but not outside the range of reasonable pessimism. For example, George Perry's study of the economic effects of terrorist events indicated that oil at \$75 a barrel would represent a “worse case” and not even the “worst case.”⁶ My own study of the economic effects of the war in Iraq also suggested an oil-price spike to \$75 a barrel for my “high” case, which further assumed that the price would decline as the war wound down.⁷

Ftn. 6,
7

6. Perry (2001).

7. Nordhaus (2002).

Figure 2. Surprise Indexes for Oil and Stock Prices, 1960–2007^a

Source: Author's calculations using data from the DRI-WEFA Basic Economics Database and the Energy Information Administration.

- a. Circles (gray for stock prices, black for oil prices) indicate index values exceeding 5 moving standard deviations.
 b. One-quarter absolute change in the logarithm of the price, divided by the twenty-year average volatility (standard deviation of the absolute change in the logarithm).

My subjective synthesis of the various components suggests that the economic impact of the 2002–06 shock was somewhere between one-third and one-half as large as those of the 1970s. However, there is some ambiguity about the relative magnitude, depending on whether the cumulative shock or the rate of shock is thought to be the appropriate measure, and whether the surprise index is also weighed in the balance.

Mechanisms for a Declining Macroeconomic Impact of Oil Shocks

It is worth pausing to consider why oil shocks might have a smaller impact on the economy today than in earlier periods. One can preliminarily consider three possible reasons: smaller shocks, smaller multipliers, and offsetting forces.

A first reason why oil-price shocks might have a smaller impact today than yesterday is that the shocks themselves are smaller. As already indicated, however, the decline in the magnitude of shocks can account for a

smaller macroeconomic response, but it cannot account for the continued economic *expansion* observed during 2002–06.

A second possible reason for a smaller impact is that the transmission mechanism has changed, tending progressively to moderate the macroeconomic impact of a given shock. The classical Frischian view of the economy distinguishes exogenous and policy shocks (in which category the first reason lies) from the complicated, nonlinear, stochastic dynamic system that takes these shocks and generates from them endogenous variables like unemployment, prices, and output. In this second hypothesis, a given variability of an exogenous variable such as the oil price leads to lower variability of output than it did in earlier periods. In the simplest macroeconomic framework, the oil-price multipliers on output and inflation are smaller.

A third possible reason for a smaller impact, discussed in detail in a later section, is that other variables might have reinforced the contractionary and inflationary impacts of higher oil prices in the 1970s but instead tended to offset those impacts in the 2000s. Economic histories of the 1970s point to commodity shortages, poor grain harvests, political business cycles, the collapse of the Bretton Woods system, disappearing anchovies, and an army of plagues, all of which reinforced the inflationary impact of the 1973 shock. In addition, the anti-inflation wars of the 1979–82 period, in part motivated by the shock that began in 1978, surely reinforced the contractionary impact of that shock.

The Trajectory of Shocks and Responses

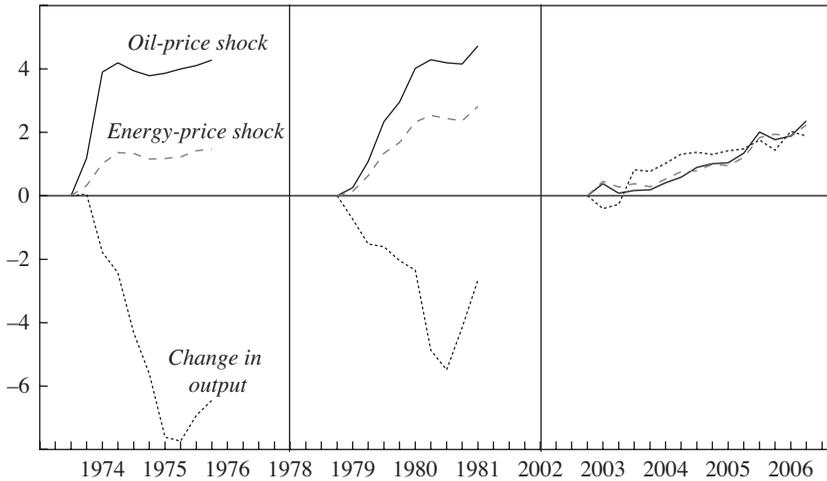
In this section I take a closer look at the pattern of shocks and responses in the three major episodes of oil-price shocks (omitting the much smaller 1990 shock). For this purpose I examine the time series of inflation and output along with the cumulative measures of energy and oil shocks defined above. To do this, I plot the cumulative oil- and energy-price shock variables just described along with the trajectories of output or inflation.

Figure 3 shows the time series for the price shock variables along with that of output. For purposes of comparing episodes I use as the output variable the ratio of real output to potential real output as defined by the Congressional Budget Office. I have normalized the variables to be zero on the shock date, so that the plot shows the cumulative change since the

Fig. 3

Figure 3. Price Shocks and Output^a

Percentage points



Source: Author's calculations using data from the DRI-WEFA Basic Economics Database and the Energy Information Administration.

a. Normalized to equal zero in the first quarter of the shock period.

shock date. The output anomaly for the 2002–06 shock is clearly visible. Whereas output fell in opposition to the positive price shocks of the 1970s, the relationship reversed in the 2000s: output continued to rise.

Figure 4 similarly plots the energy shocks alongside inflation, the latter defined both in terms of the consumer price index, or CPI (which policy-makers tended to follow in the 1970s) and in terms of the PCE price index. The relationship between inflation and the price shocks is clearly positive for each of the three periods.

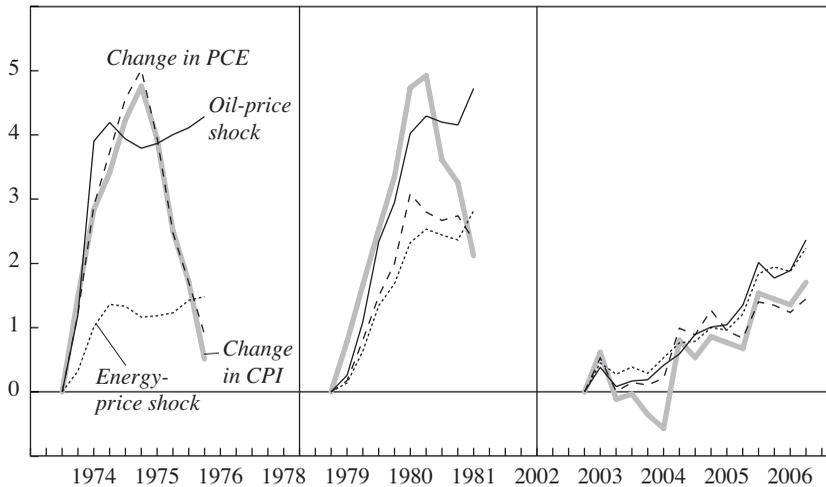
Looking closer, however, one sees some interesting patterns. Perhaps the most interesting relationship is between the energy-price shock and PCE inflation. By construction, if the only factor changing during the period was direct energy-price inflation, and there were no indirect effects, the two lines should coincide (almost) exactly, because the PCE price index is a superlative index that uses virtually the same construction as my energy-price shock variable.

In 1973–74, however, the change in PCE inflation was about three times the energy-price shock. This indicates that other factors (imports,

Fig. 4

Figure 4. Price Shocks and Inflation^a

Percentage points



Source: Author's calculations using data from the DRI-WEFA Basic Economics Database and the Energy Information Administration.

a. Normalized to equal zero in the first quarter of the shock period. Inflation is defined as the simple year-over-year inflation rate.

indirect effects, other commodities, wages) were operating to reinforce the energy-price shock in this period. Recall as well that the CPI bias was small in this period. In the second period, virtually all of the rise in PCE inflation can be accounted for by direct energy-price inflation. Note, however, that the CPI in the 1978–81 period rose much more rapidly than the PCE price index, partly because of the CPI's flawed construction due to inappropriate treatment of housing and the index's Laspeyres structure. In the third period the contribution of the energy-price shock was about $1\frac{1}{2}$ times PCE inflation. During this period, therefore, other factors were retarding inflation slightly. Of particular note, but not shown in the figure, is that energy prices were not passed through into higher wage inflation after 2002.

To sum up the evidence on the trajectories of inflation and output: the 2002–06 episode was completely anomalous for output, and the major news on recent inflation is that there appears to have been no pass-through of higher energy prices into other prices and wages.

A Declining Impact of Shocks?

It seems clear that the disappearance of energy crises is not completely due to a lack of energy shocks. This leaves the second and third potential causes identified above as possible explanations: The lack of a response might be due to smaller amplifiers in the macroeconomic transmission mechanism, or to offsetting forces stemming from macroeconomic policy or exogenous events. Sorting out these two forces is clearly a major task, but a few suggestions can be offered.

Perhaps the most important question is whether the macroeconomy has become less sensitive to exogenous or policy shocks (including oil-price shocks) in recent years. Investigating this issue requires taking a stance on the structure of the macroeconomy. In this section I take a modest first step by examining a simple, one-equation model for aggregate demand. This equation, presented in the appendix, uses as forcing variables the cumulative oil shock variable, interest rates, and exogenous spending. As with the Taylor rule equation, I estimate this equation for rolling twenty-year periods.

The bottom panel of table 2 reports the estimated coefficients on the oil shock and exogenous spending variables. The impact of exogenous spending continues to be relatively strong, perhaps even increasing slightly over the period. On the other hand, the impact of the oil shock variable declines sharply over time and is close to zero after 1980—a period that excludes the price shocks of the 1970s.

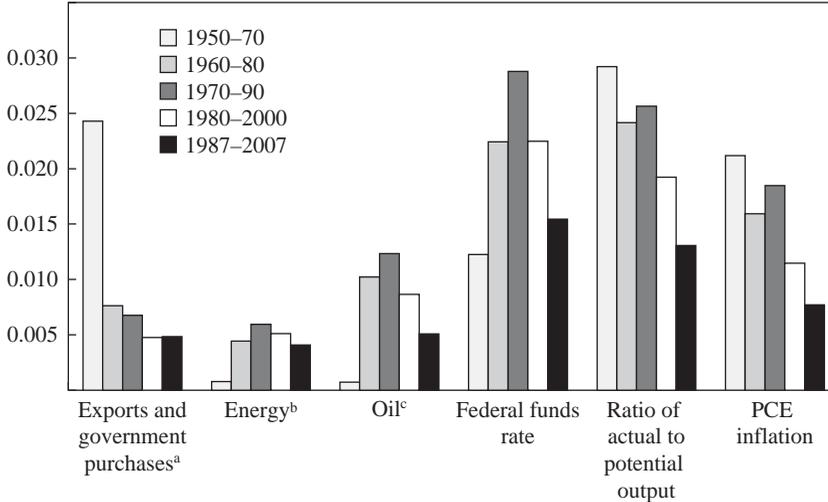
Rolling estimates of volatility over time provide another view of the contributions of different factors. Figure 5 shows, for each of five sub-periods, the standard deviations of four-quarter changes in four policy or shock factors and in inflation and output. As suggested by the theory of the Great Moderation, the variability of output over this range of observations declined by almost two-thirds while that of inflation declined by more than half.

Figure 5 also shows the increase and then decrease in the volatility of oil- and energy-price shocks over the postwar period. The increase in volatility from the 1950–70 period to later periods is extremely dramatic, and there has been only a modest decline in energy-price volatility, and a slightly larger decline in oil-price volatility, since the 1970–90 peak. In terms of macroeconomic impact, the decreasing macroeconomic sensitivity of output to oil-price shocks has accompanied declining shocks over

Fig. 5

Figure 5. Volatility of Two Shock Variables, Two Policy Variables, Output, and Inflation, by Subperiod

Four-quarter standard deviation



Source: Author's calculations using data from the DRI-WEFA Basic Economics Database and the Energy Information Administration.

a. As a ratio to potential output.

b. Scaled energy-price shock variable (from figure 1).

c. Scaled oil-price shock variable (from figure 1).

the period since 1970. This combination of smaller shocks and reduced sensitivity to a given shock has unambiguously led to a smaller impact of oil- and energy-price shocks on the economy since the 1970s.

What of exogenous spending shocks? The stunning point is that the variability of the exogenous component has declined by a factor of five since the first period. This sharp decline is fighting against a roughly constant multiplier, with the likely outcome being a lower volatility of output.

An additional question is how oil-price shocks get transmitted into inflation. One possible reason for the observed pattern of effects is that wage earners are less inclined today to try to maintain real wages, and an energy shock therefore generates less of a wage response. A quick look at the dynamics of compensation using a standard accelerationist Phillips curve finds little evidence that wages have responded less to cyclical conditions in recent years than in earlier periods. However, there is clear evidence that the wage reaction is moving from a price-wage dynamic to

a wage-wage dynamic. In other words, wage increases appear to be highly inertial today, whereas they were more responsive to price inflation in earlier decades. This shift does confirm the impression that the oil-price shocks of the 1970s got passed through to wages much more than has been the case in the 2000s.

The net assessment is that the declining impact of oil- and energy-price shocks on the economy results from a combination of both declining sensitivity of the economy to energy-price shocks and declining variability of energy and oil prices since the 1970s. The first factor appears quantitatively more important, but estimates of the sensitivity parameter are hardly well determined.

Headwinds and Tailwinds

The final question is whether the conventional association of oil shocks with recession is just the post hoc ergo propter hoc fallacy writ large. Is it not just as likely that the oil shocks of the 1970s came at a point when the economy was already battling gale-force headwinds of tightening money, declining exports, shrinking government spending, and natural resource and commodity scarcities? By contrast, might not the oil shock of the 2000s have benefited from the gentle tailwinds of recovery from recession, fiscal expansion from tax cuts and war spending, and a flood of liquidity from a glut of global savings?

The evidence on this issue is mixed. James Hamilton argued strenuously at the time that the early oil shocks were statistically independent of other macroeconomic events.⁸ Moreover, there is little reason to believe that the four events examined here were actually triggered by cyclical economic forces. For example, I know of no evidence that the timing of the invasion of Iraq in 2003 was motivated by a desire to provide an economic stimulus just before the 2004 election. So the argument is really that the sample of observations is much too small (just two for the 1970s and only one for the past fifteen years) to allow sweeping generalizations about the strong contractionary impact of oil-price increases.

The most recent evidence suggests that the strength and direction of other forces is a major part of the difference between the experience of the

Ftn. 8

8. Hamilton (1983).

1970s and that of the 2000s. Blanchard and Galí conclude that, in addition to other reasons, the evidence is “consistent with the hypothesis that other (non-oil) shocks have coincided in time with the major oil shocks, either reinforcing the adverse effects of the latter in the 1970s, or dampening them during the more recent episodes.”⁹

Ftn. 9

One can address this issue using the simple aggregate demand equation discussed above. When the equation is fit to the period 1970:Q1 to 2007:Q2, the results show that (assuming constant coefficients) exogenous and policy forces (exports, government purchases, and interest rates) were highly contractionary in both the 1970s episodes, but neutral in the 2002–06 period. Much further study is needed to determine the strength and direction of the prevailing economic winds in the different energy-price storms. But this surely looks like an important part of the story.

Conclusion

It is clear that the economy behaved differently after the most recent oil-price shock than after the earlier ones. The behavior of output after the latest shock was completely different from earlier episodes. Indeed, the sign was opposite: output continued to grow relative to potential output after the shock, and unemployment continued to fall. The impact of the recent shock on inflation was qualitatively similar to that in past episodes although quantitatively different. Unlike with the shocks of the 1970s, there appears to have been no substantial pass-through of the energy-price increases into wages or other nonenergy prices.

Why did the economy perform so much better after the oil-price shock of the 2000s than after the earlier shocks? I have identified three reasons: the shock itself was smaller, the transmission mechanism between shocks and the rest of the economy was weaker, and other forces in the macroeconomic environment were working against rather than with the shock.

To begin with the first reason: the oil shock of 2002–06 was different from the earlier ones: it was less of a surprise and occurred more gradually. I have suggested that a reasonable reading of the various measures was that the economic impact of the 2002–06 shock was between one-third and one-half as large as those of the 1970s. The number is clouded by interpretive

9. Blanchard and Galí (2007).

uncertainty, however. The 2002–06 shock appears larger if one ignores the surprise element and considers the cumulative impact to be the relevant measure, and smaller if only the annual average impact is thought important and the analysis incorporates the huge surprise in 1973–74.

Turning to the second reason, there is evidence that the transmission mechanism from energy prices to output has changed from negative to neutral over the last three decades. The reasons for this declining macroeconomic sensitivity to energy prices are not completely understood, but two underlying causes seem plausible.

First, there is evidence that the Federal Reserve has reacted more sensibly to energy prices in the 2000s than it did in the 1970s, looking at core inflation rather than total inflation and attending to the superlative PCE price index rather than the flawed CPI of the 1970s and early 1980s. Perhaps experienced hands at the Federal Reserve remembered history and decided not to repeat it. Perhaps they judged, correctly in my view, that the energy-price shock was a one-time shock to the price level that would not be passed through into wage inflation.

A second and more speculative reason for the muted macroeconomic reaction is that today's consumers, businesses, and workers may see oil-price increases as volatile and temporary movements rather than the earth-shaking changes that they were thought to be in the 1970s. There is evidence that shocks tend to be slightly more transient today than were those of the 1970s. Returning to the oil-price-as-tax-increase theory, consumers might view oil-price increases as temporary rather than permanent "tax" increases, and if so their reactions would be correspondingly smaller. Similarly, businesses today might build energy-price volatility into their plans and investments, so that large price changes do not upset operations as much as in earlier periods. A similar set of factors might lead workers and unions to absorb declines in real wages induced by energy prices rather than go on strike in an attempt to recapture losses that might turn out to be transitory. All of these factors would tend to reduce the impact of energy-price shocks on the macroeconomy.

Finally, the strength and direction of the prevailing economic winds in the different episodes—the third reason—are also an important part of the difference. External cyclical forces such as government purchases, exports, and financial conditions were clearly slowing the economy during the oil shocks of the 1970s, whereas they were roughly neutral, by the measures used here, from 2002 to 2006.

Additionally, economists were deriving their view of the contractionary influence of oil shocks from two or three observations. So although it may be plausible that the origins of the oil shocks were genuinely exogenous political events such as the Iranian revolution, the sample size was too small to make the sweeping generalizations that many did about the strong contractionary impacts of oil-price increases. Economists should perhaps have been more suspicious, given the analytical point that energy-price changes have no effect on multifactor productivity and very little effect on labor productivity. The history of the profession's interpretation of oil shocks is a reminder that we have many fewer degrees of freedom than our time-series statistics indicate.

A cautious reading today suggests that policymakers should not be afraid of a Big Bad Oil Shock. The economy weathered an increase in real oil prices of 125 percent from 2002 to 2006 without any major strain. The evidence suggests that the main thing we have to fear about oil-price shocks is the fearful overreactions of the monetary authority, consumers, businesses, and workers.

APPENDIX

Equations Used in Table 2

The equations estimated in the two panels of table 2 are an aggregate demand equation (equation A.1 below) and a modified Taylor rule (equation A.2):

$$(A.1) \quad cu_t = a_0 + a_1 exog_t + a_2 cu_{t-2} + a_3 TB3_{t-2} + a_4 CumOilshock_t$$

$$(A.2) \quad ffff_t = \beta_0 + \beta_1 u_t + \beta_2 pi_core_exenergy_t \\ + \beta_3 pi_energy_{t-2} + \beta_4 ffff_{t-2}$$

Here cu_t is the ratio of output to potential output, $exog_t$ is the ratio of exports plus government spending to potential output, $TB3_t$ is the nominal three-month Treasury bill rate, $CumOilshock_t$ is the cumulated oil-shock variable, $ffff_t$ is the federal funds interest rate, u_t is the unemployment rate, $pi_core_exenergy_t$ is the PCE inflation rate with energy prices removed, and pi_energy_t is the contribution of energy prices to PCE inflation. Equation A.1 is estimated using two-stage least squares, with lagged values of

all variables except cu_t used as instruments. Both equations were estimated with autoregressive (AR1) corrections.

References

- Bernanke, Ben S., Mark Gertler, and Mark Watson. 1997. "Systematic Monetary Policy and the Effects of Oil Price Shocks." *BPEA*, no. 1: 91–142.
- Blanchard, Olivier J., and Jordi Galí. 2007. "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?" In *International Dimensions of Monetary Policy*, edited by Mark J. Gertler and Jordi Galí. University of Chicago Press (forthcoming).
- Hamilton, James D. 1983. "Oil and the Macroeconomy Since World War II." *Journal of Political Economy* 91, no. 2: 228–48.
- _____. 1996. "This Is What Happened to the Oil Price-Macroeconomy Relation." *Journal of Monetary Economics* 38, no. 2: 215–20.
- Hamilton, James D., and Ana Maria Herrera. 2004. "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: A Comment." *Journal of Money, Credit, and Banking* 36, no. 2: 265–86.
- Nordhaus, William D. 2002. "The Economic Consequences of a War with Iraq." In *War with Iraq: Costs, Consequences, and Alternatives*, edited by Carl Kaysen and others. Cambridge, Mass.: American Academy of Arts and Sciences.
- Perry, George L. 2001. "The War on Terrorism, the World Oil Market, and the US Economy." Analysis Paper no. 7. Brookings (October 24).
- Stock, James H., and Mark W. Watson. 2003. "Has the Business Cycle Changed and Why?" In *NBER Macroeconomics Annual 2002*, edited by Mark Gertler and Kenneth Rogoff. MIT Press.