JAMES D. HAMILTON
University of California, San Diego

Causes and Consequences of the Oil Shock of 2007–08

ABSTRACT This paper explores similarities and differences between the run-up of oil prices in 2007–08 and earlier oil price shocks, looking at what caused these price increases and what effects they had on the economy. Whereas previous oil price shocks were primarily caused by physical disruptions of supply, the price run-up of 2007–08 was caused by strong demand confronting stagnating world production. Although the causes were different, the consequences for the economy appear to have been similar to those observed in earlier episodes, with significant effects on consumption spending and purchases of domestic automobiles in particular. Absent those declines, it is unlikely that the period 2007Q4–2008Q3 would have been characterized as one of recession for the United States. This episode should thus be added to the list of U.S. recessions to which oil prices appear to have made a material contribution.

The price of oil has been anything but stable over the last four decades (figure 1). A series of dramatic events in the 1970s sent the price of crude oil over $40 a barrel by the end of that decade, which would be over $100 a barrel at current prices. The price remained very volatile after the collapse in the 1980s but was still as low as $20 a barrel at the end of 2001. The next six years saw a steady increase that tripled the real price by the middle of 2007. Later that year the path of oil prices steepened sharply, sending the nominal price to an all-time high of $145 a barrel on July 3, 2008, only to be followed by an even more spectacular price collapse.[1] What caused this remarkable behavior of oil prices, and what were the effects on the economy?

To answer these questions, I begin by reviewing, in section I, some important characteristics of the demand for petroleum. Section II then

1. Crude oil prices quoted in this paper refer to the spot price of West Texas Intermediate (WTI) except where stated otherwise.

Figure 1. Real Crude Oil Prices, January 1947 to December 2008[a]

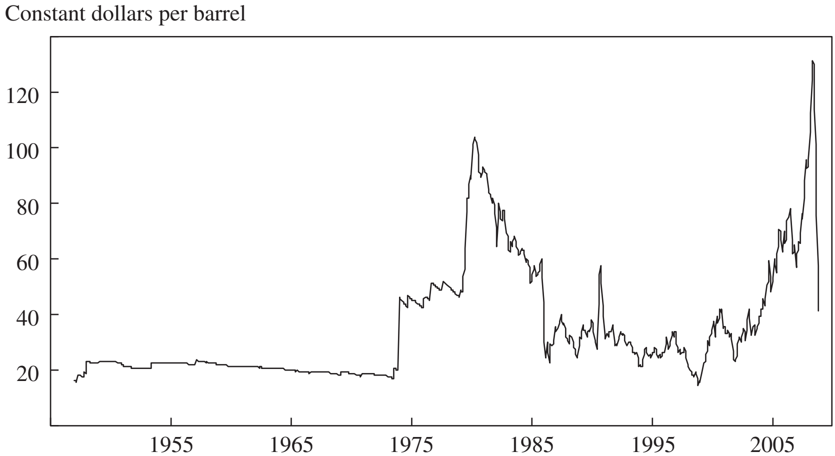

Sources: Energy Information Administration (EIA) data; Bureau of Labor Statistics data.
a. Monthly average price of West Texas Intermediate crude in dollars of November 2008.

explores the causes of several of the principal oil shocks of the late 20th century. I then turn, in section III, to an analysis of what happened to produce the dramatic price moves in 2007 and 2008. Next, section IV reviews some of the evidence on how the economy seemed to respond to earlier oil price shocks, and section V investigates the effects on the economy of the oil shock of 2007–08. Section VI briefly notes some implications for policy.

I. Some Observations on Petroleum Demand

The single most important fact for understanding short-run changes in the price of oil is that income rather than price is the key determinant of the quantity demanded. One quick way to become convinced of this fact is to examine figure 2, which is essentially a scatterplot of U.S. petroleum consumption against GDP over the last 60 years, with successive years connected by a smoothed curve. Tracing this curve from lower left to upper right thus tracks the realized combinations of real GDP and petroleum consumption as they changed over time. As the figure shows, despite the huge fluctuations in the relative price of oil over this period, petroleum consumption followed income growth remarkably steadily, with an elasticity of about 1 through 1973. There was some downward adjustment in oil use

Figure 2. Real GDP and Oil Consumption, United States, 1949–2007[a]

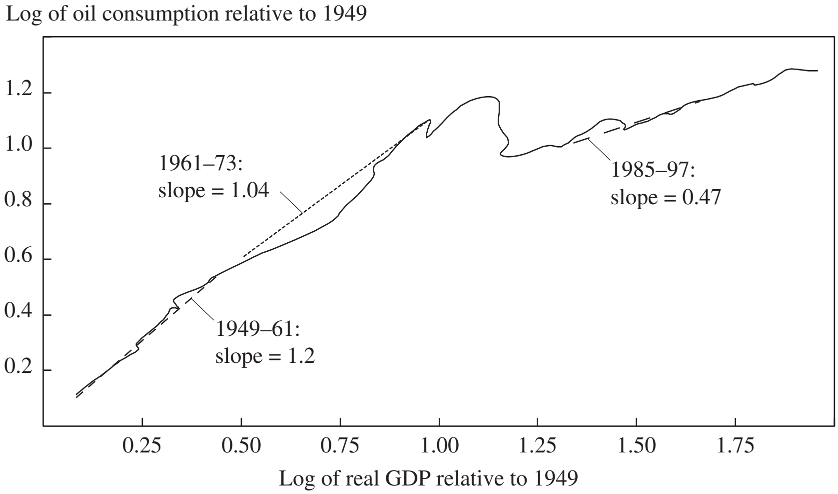

Source: Hamilton (2009), using data from Bureau of Economic Analysis, National Income and Product Accounts, table 1.1.6, “Real Gross Domestic Product, Chained Dollars,” and EIA, “Petroleum Overview, 1949–2007,” table 5.1 (www.eia.doe.gov/emeu/aer/txt/ptb0501.html).

a. Each observation plots for a single year the cumulative change in the logarithm of total petroleum products supplied to the U.S. market since 1949 against the cumulative change in the logarithm of real GDP since 1949.

between 1978 and 1982, although achieving that 20 percent (logarithmic) drop required an 80 percent (logarithmic) increase in the relative price of oil and two recessions in a three-year period (1980–82).

The slope of this path flattens after the early 1970s to reflect an income elasticity nearer 0.5, a phenomenon that some might attribute to the delayed consequences of increased energy conservation following the 1970s oil shocks. However, this flatter slope persists long after the price has fallen quite dramatically, and it seems more likely to be due to a tendency for the income elasticity of oil consumption to decline as a country becomes more developed. One sees a similar pattern of slowing growth of petroleum use in other developed countries as they become richer, and post-1990 data for several newly industrialized countries are still quite supportive of an income elasticity near unity (Hamilton 2009; Gately and Huntington 2002).

Table 1 summarizes the estimated price elasticities for gasoline and crude oil demand from a half-dozen meta-analyses or literature reviews. Since crude oil represents about half the retail cost of gasoline (Energy Information Administration 2006), one would expect that a 10 percent

Table 1. Literature Estimates of the Short-Run Price Elasticity of Demand for Gasoline and Crude Oil

Study	*Product*	*Method*	*Elasticity*
Dahl and Sterner (1991)	Gasoline	Literature survey	0.26
Espey (1998)	Gasoline	Literature survey	0.26
Graham and Glaister (2004)	Gasoline	Literature survey	0.25
Brons and others (2008)	Gasoline	Literature survey	0.34
Dahl (1993)	Crude oil	Literature survey[a]	0.07
Cooper (2003)	Crude oil	Annual time-series regression	0.05[b]

Source: Hamilton (2009).
a. Survey covers developing countries only.
b. Average of 23 countries.

increase in the price of crude would be associated with a 5 percent increase in the price of gasoline,[2] in which case the price elasticity of demand for crude oil should be about half that for retail gasoline. Most of the studies summarized in these reviews reported low estimates of the price elasticity of gasoline demand and significantly smaller elasticities for crude.

The price elasticity of petroleum demand has always been small, and it is hard to avoid any conclusion other than that it became even smaller in the United States in the 2000s. One can barely detect any downward deviation from the trend in petroleum consumption in figure 2 for that period despite the enormous price increase through 2007. Jonathan Hughes, Christopher Knittel, and Daniel Sperling (2008) estimate that the short-run price elasticity of gasoline demand was (in absolute value) in the range of 0.21 to 0.34 over 1975–80, but between only 0.034 and 0.077 for 2001–06.

Another key parameter for determining the consequences of an energy price increase is the value share of energy purchases in total expenditure. The fact that the U.S. income elasticity of demand has been substantially below unity over the last quarter century induces a downward trend in that share: for a given relative price, if the percentage growth in energy use is less than the percentage growth in income, total dollar expenditure on energy will decline as a percentage of income. On the other hand, the very low short-run price elasticity of demand causes the value share to move in

2. The regression coefficient relating the logarithm of the nominal U.S. gasoline retail price to the log of the nominal WTI price in a monthly cointegrating regression estimated over April 1993–August 2008 is 0.62. Data from Energy Information Administration (EIA), "Spot Prices for Crude Oil and Petroleum Products" (tonto.eia.doe.gov/dnav/pet/pet_pri_spt_s1_m.htm).

Figure 3. Consumer Energy Expenditure as a Share of Total Personal Consumption Expenditure, United States, January 1959 to September 2008

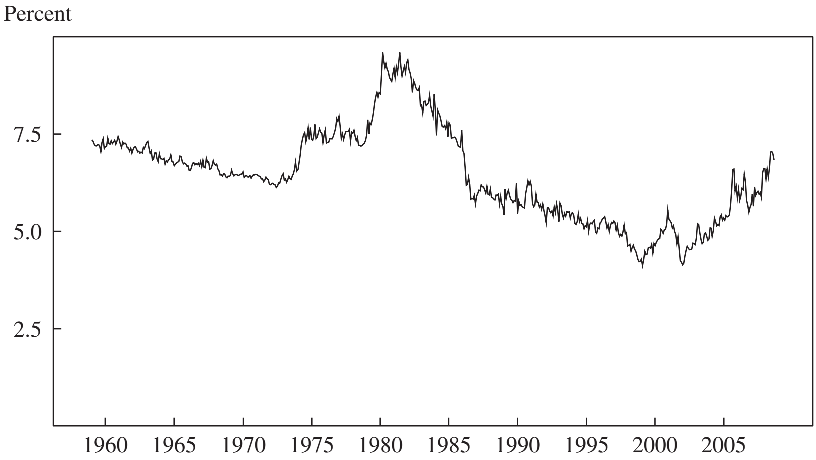

Source: Bureau of Economic Analysis, National Economic Accounts, table 2.3.5U, "Personal Consumption Expenditures by Major Type of Product."

the same direction as the relative price: if the percentage increase in price is greater than the percentage decrease in quantity demanded, dollar spending as a share of income will rise when the price of energy goes up.

Figure 3 displays the net effect of these two factors on spending by consumers on energy goods and services, measured here as a percentage of total consumption spending. The income elasticity effect imparts a chronic downward trend, so that by 2002 energy purchases had fallen to a little over 4 percent of a typical consumer's total budget. However, subsequent energy price increases produced a dramatic reversal of this trend, with the share in 2008 rising to almost twice the 2002 value.

Figure 3 also serves as a reminder that a price elasticity cannot remain below unity in all circumstances. A consumer who fails to reduce the quantity purchased of an item by as much in percentage terms as its price goes up will find that the item comes to consume an ever-larger fraction of her budget. If her price elasticity were constant at less than unity, an arbitrarily large price increase would ultimately bring her to a point where 100 percent of her budget was going to energy, in which case ignoring the price would no longer be possible. The low energy expenditure share in the early part of this decade may be part of the explanation for why Americans largely ignored the early price increases: we didn't change our behavior much because most of us could afford not to. By 2007–08, however, the situation had changed, as energy once again returned to an importance in the typical budget that had not been seen since the 1970s.

II. Causes of Past Oil Supply Disruptions

Figure 4 plots monthly oil production for three Middle Eastern countries that have recurrently appeared in the news over the last 35 years. Three events over this period—the Iranian revolution in the fall of 1978, Iraq's invasion of Iran in September 1980, and Iraq's invasion of Kuwait in August 1990—resulted in dramatic and immediate disruption of the flow of oil from key global producers. Another episode, not evident in figure 4 but that I will nevertheless include in the set of historical oil shocks discussed, began with the Yom Kippur War on October 6, 1973. Although that conflict did not directly prevent any significant shipments of oil, the Organization of Arab Petroleum Exporting Countries (OAPEC) announced on October 17 that it would cut production by 5 percent until Israeli forces "are completely evacuated from all the Arab territories occupied in the June 1967 war, and the legitimate rights of the Palestinian people are restored."[3] In a previous paper (Hamilton 2003), I included the Suez crisis of 1956 as a fifth significant oil shock. However, the price increase in that episode was much more modest, and data for the kinds of comparisons performed below are not readily available for that period, so this paper will use the above four episodes only.

Each row of graphs in figure 5 focuses on one of these four episodes. The four left-hand panels record the drop in oil production during each episode, for OAPEC as a group in the 1973–74 episode (top left panel) and for the producing country or countries most affected in the others. In each case the production shortfall is expressed as a percentage of total global production just before the shock.[4] As the figure shows, each of these events removed between 6 and 9 percent of world supply.

In each episode, increased production in other countries partially mitigated the consequences. To indicate the net consequences of the disruptions for global production, the left-hand panels of figure 5 also show the percentage decline in world oil production following each event. Production increases from other countries were rather minor in 1973–74 but quite substantial in 1990–91.

3. OAPEC ministers' press release, as quoted in Al-Sowayegh (1984, p. 129).

4. These numbers differ slightly from those reported in table 4 of Hamilton (2003) because of small differences in the estimates of total global oil production used, and because here the Iranian shortfall is dated as beginning in October rather than September 1978.

Figure 4. Oil Production of Selected Middle Eastern Countries, January 1973 to June 2007

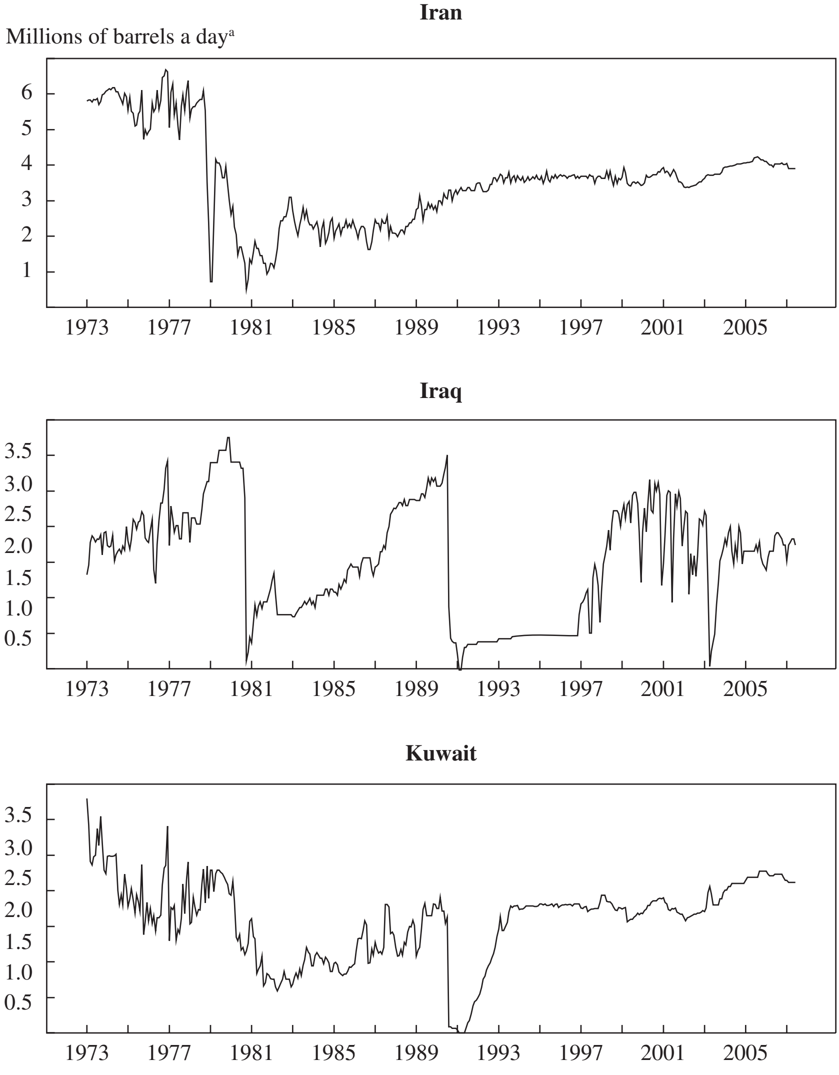

Source: EIA, *Monthly Energy Review,* table 11.1a, "World Crude Oil Production: OPEC Members" (tonto.eia. doe.gov/merquery/mer_data.asp?table=T11.01a).

a. Monthly averages, including lease condensate.

Figure 5. Cumulative Changes in World Oil Production, Prices, and Inventories after Four Major Disruptions

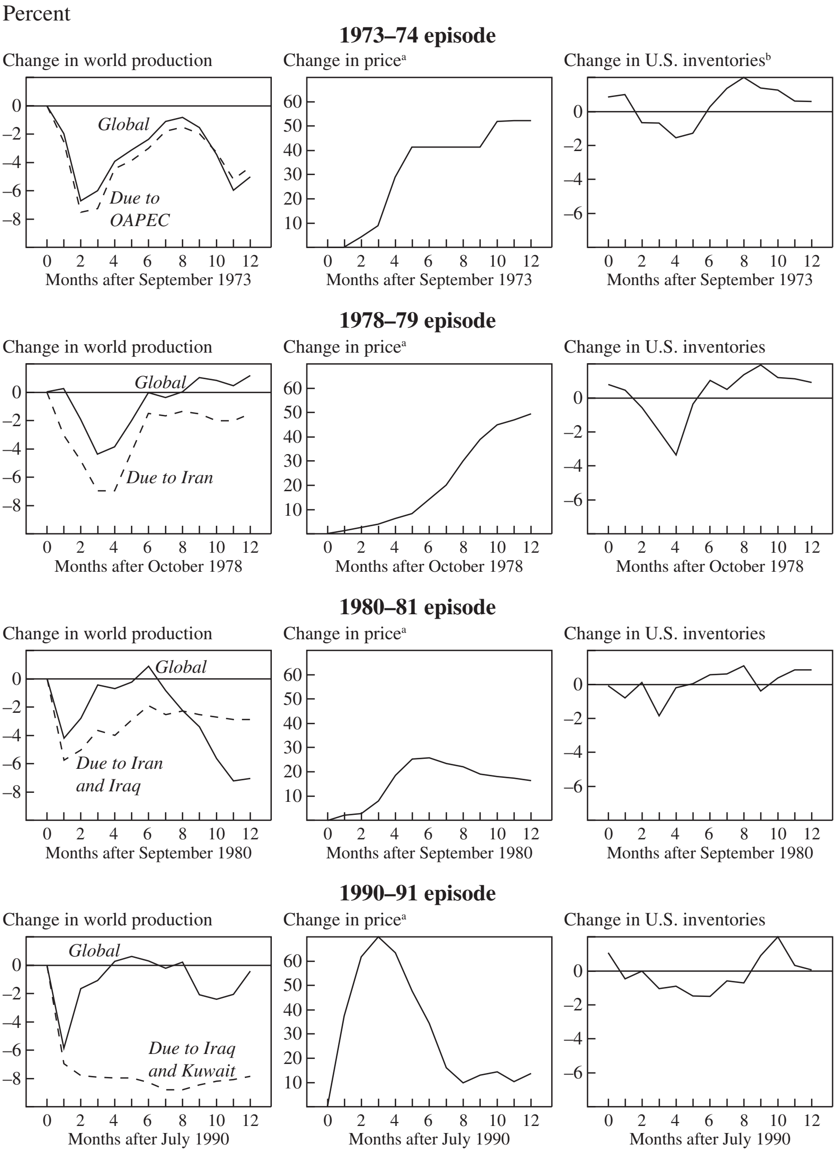

Sources: EIA, *Monthly Energy Review,* tables 11.1a (tonto.eia.doe.gov/merquery/mer_data.asp?table=T11.01a) and 11.1b (tonto.eia.doe.gov/merquery/mer_data.asp?table=T11.01b); EIA, "Refiner Acquisition Cost of Crude Oil" (tonto.eia.doe.gov/dnav/pet/pet_pri_rac2_dcu_nus_m.htm); EIA, "Total Stocks" (tonto.eia.doe.gov/dnav/pet/pet_stoc_wstk_dcu_nus_m.htm); Bureau of Labor Statistics data; author's calculations.

a. For the 1973–74 episode, change relative to the indicated starting month in 100 × the log of the seasonally unadjusted producer price index for crude petroleum; for other episodes, cumulative change in 100 × the log of the monthly refiner acquisition cost for crude petroleum.

b. In all episodes, monthly change in end-of-month U.S. stocks of crude oil and petroleum products, as a share of world production.

Table 2. Quantity and Price Changes in Past Oil Shocks

Episode	*Supply reduction (percent)*[a]	*Increase in price (percent)*[b]	*Implied price elasticity of demand*[c]
October 1973–March 1974	4.0	41.3	0.10
November 1978–July 1979	1.3	38.7	0.03
October 1980–March 1981	1.2	25.8	0.05
August 1990–October 1990	2.9	71.6	0.04

Sources: EIA, *Monthly Energy Review,* table 11.1a (tonto.eia.doe.gov/FTPROOT/multifuel/mer/00350807.pdf); EIA, "Refiner Acquisition Cost of Crude Oil" (tonto.eia.doe.gov/dnav/pet/pet_pri_rac2_dcu_nus_m.htm); Bureau of Labor Statistics data; author's calculations.

a. Average monthly shortfall of global production of crude petroleum over the episode as a share of global production in the month before the episode.

b. Peak value during the episode of the cumulative change in price, calculated as 100 times the logarithm of the ratio of the current monthly refiner acquisition cost to the cost at the beginning of the episode.

c. Ratio of the supply reduction to the change in price.

The contemporaneous path of oil prices is depicted in the middle column of panels in figure 5. Each episode was associated with a significant increase in the price of oil, for example by 25 percent (logarithmically) in 1980 and 70 percent in 1990. However, price controls in effect during the first three episodes spread the consequences over time.

Lutz Kilian downplays the contribution of these supply disruptions to the price movements portrayed in figure 5, instead attributing much of the historical fluctuation in the price of oil to what he describes as "precautionary demand associated with market concerns about the availability of future oil supplies" (Kilian 2009, p. 1053). He identifies precautionary demand as any movement in the real price of oil that cannot be explained statistically by his measures of shocks to supply and aggregate demand. One might also try to measure the contribution of precautionary demand by looking at changes in inventories. The four right-hand panels of figure 5 record the monthly change in U.S. inventories of crude oil and petroleum products in each episode, again measured as a percentage of total global production. In each of these episodes, inventories were going down, not up, at the time of the sharpest price movements, suggesting that inventory changes were serving to mitigate rather than aggravate the price shocks. Positive inventory investment typically came much later, as firms sought to restock inventories that had been earlier drawn down.

One can also explore whether the supply disruptions alone suffice to explain the observed price movements on the basis of plausible elasticities. Table 2 compares the average decline in global oil production during each of the four episodes with the observed price change and calculates the implied price elasticity of demand assuming zero shift in demand from growing income and no other influences on prices during the episode. These elasticities are a bit smaller than might have been expected from the consensus estimates in table 1, but none of them render it implausible to attribute most of the price change to the supply shortfall itself.

Kilian (2009) also argues that the declines in individual-country (or OAPEC) production in the left-hand panels of figure 5 overstate the magnitude of the supply disruptions in these four episodes. He observes, for example, that Iraq increased production significantly in anticipation of both its war with Iran in 1980 and its invasion of Kuwait in 1990, so that measuring the decline in Iraqi production relative to levels just before the conflict overstates the shock (see the middle panel of figure 4). Note, however, that this is not a factor in the global production figures used in figure 5 or the calculations in table 2, which are based on the observed global decline after the indicated date. Moreover, despite high Iraqi production in the months before the 1980 war, global production in September 1980 was 2.9 percent below its level three months earlier and 5.4 percent below that of six months earlier. Likewise, global production in July 1990 was down 2.1 percent and 0.7 percent from its values three and six months earlier, respectively. Hence, a comparison of global production in these episodes with values earlier than the September 1980 and July 1990 reference dates used would have made the imputed quantity reductions in table 2 even more significant.

Kilian (2009) and Robert Barsky and Kilian (2002) argue, quite correctly in my view, that demand pressures also contributed to higher oil prices in several of these episodes. In particular, it would be irresponsible to claim that the nominal oil price increase in 1973–74 had nothing to do with the general inflation and the boom in the prices of other commodities also observed at that time. Nevertheless, I share Alan Blinder and Jeremy Rudd's (2008) doubts about whether inflationary pressures are the primary explanation for why OAPEC chose to reduce its oil output by 5 percent within weeks of the onset of the Yom Kippur War. My overall conclusion thus supports the conventional interpretation: the oil price shocks of past decades were primarily caused by significant disruptions in crude oil production brought about by largely exogenous geopolitical events.

III. Causes of the Oil Shock of 2007–08

Figure 6 plots five different measures of energy prices during the last quarter of 2007 and the first half of 2008. By any measure this episode qualifies as one of the biggest shocks to oil prices on record. However, the

Figure 6. Five Measures of the Change in Energy Prices, August 2007 to June 2008

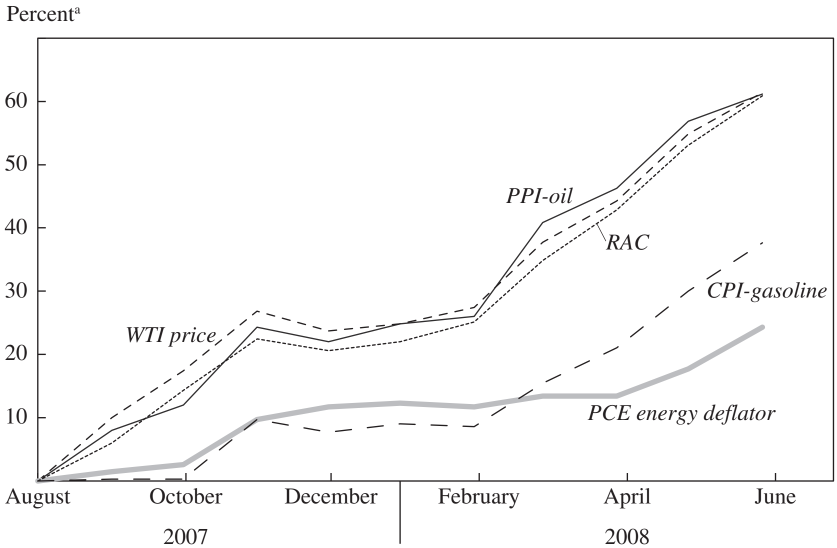

Sources: Bureau of Labor Statistics, series PW561 and PZU471, via Webstract; FRED, Federal Reserve Bank of St. Louis, series OILPRICE; EIA, "Refiner Acquisition Cost of Crude Oil" (tonto.eia.doe.gov/dnav/pet/pet_pri_rac2_dcu_nus_m.htm); Bureau of Economic Analysis, series PINRGG3, via EconStats.

a. Cumulative change in price, calculated as 100 × the logarithm of the ratio of the current value to the value in August 2007. PPI, producer price index for crude petroleum; WTI price, monthly average price of West Texas Intermediate crude oil; RAC, refiner's acquisition cost for crude petroleum; CPI-gasoline, consumer price index for gasoline; PCE energy deflator, implicit price deflator for personal consumption expenditure on energy goods and services.

causes were quite different from those associated with the four episodes examined above.

III.A. Supply

Despite occasional dramatic news, such as hurricanes in the Gulf of Mexico in September 2005, turmoil in Nigeria in 2006–08, and ongoing strife in Iraq, global oil production in the 2000s has been remarkably stable (figure 7). The big story has been not a dramatic reduction in supply of the kind summarized in figure 5, but rather a failure of production to increase between 2005 and 2007.

Why did global production stagnate? In any given producing field, pressure eventually falls and daily output begins to decline. Increasing global production requires moving on to new producing areas. The United States' endowment of crude oil has been extensively explored and developed, and total U.S. production is now about half what it was in 1971 (top panel of

Figure 7. World Crude Oil Production, January 2003 to October 2008

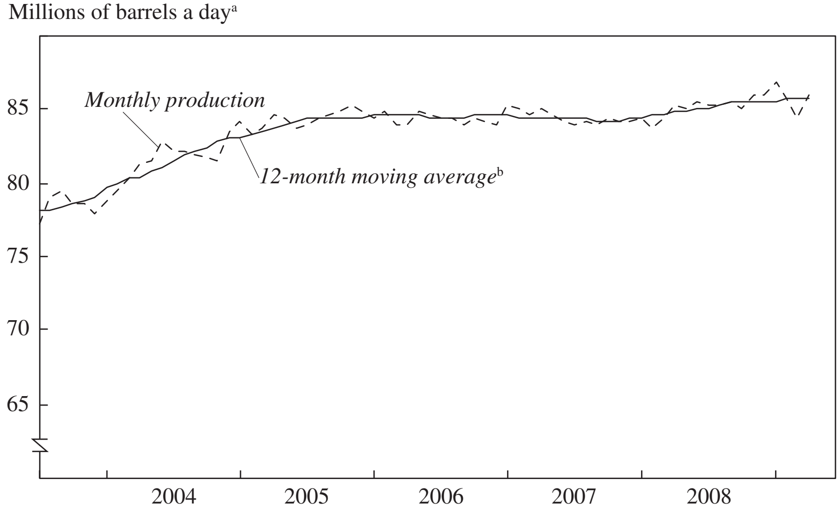

Source: EIA, "Total Oil Supply," table 1.4 (www.eia.doe.gov/emeu/ipsr/t14.xls).

a. Includes lease condensate, natural gas plant liquids, other liquids, and refinery processing gain.

b. Average is centered at the indicated date, with end-of-sample values representing the average of $\{x_{t-6}, \ldots, x_{t+s}\}$ for feasible s.

figure 8). World production has nevertheless increased substantially since then as new fields have been developed. On the other hand, figure 8 also shows that several of the new fields are now in significant decline, including the North Sea (which accounted for 8 percent of world production in 2001) and Mexico's Cantarell field (formerly the world's second-largest producing field). Production declines caused former OPEC member Indonesia to become an oil importer; it dropped out of OPEC in 2008.

The world's most important oil exporter has for many years been Saudi Arabia, whose monthly production is plotted in the bottom panel of figure 8. Saudi oil output has historically been quite volatile, not because of depletion but because the Saudis followed a deliberate strategy of adjusting production in an effort to stabilize prices. For example, the kingdom's decision to increase production sharply in late 1990 was a reason why the oil price shock of 1990 was so short-lived: increased Saudi output accounts for much of the early rebound in the bottom-left panel of figure 5.

Because the Saudis had historically used their excess capacity to offset short-run supply shortfalls elsewhere, many analysts assumed that they would continue to do so in response to the longer-run pressure of growing world demand, and so most forecasts called for continuing increases

Figure 8. Crude Oil Production from Selected Countries and Fields

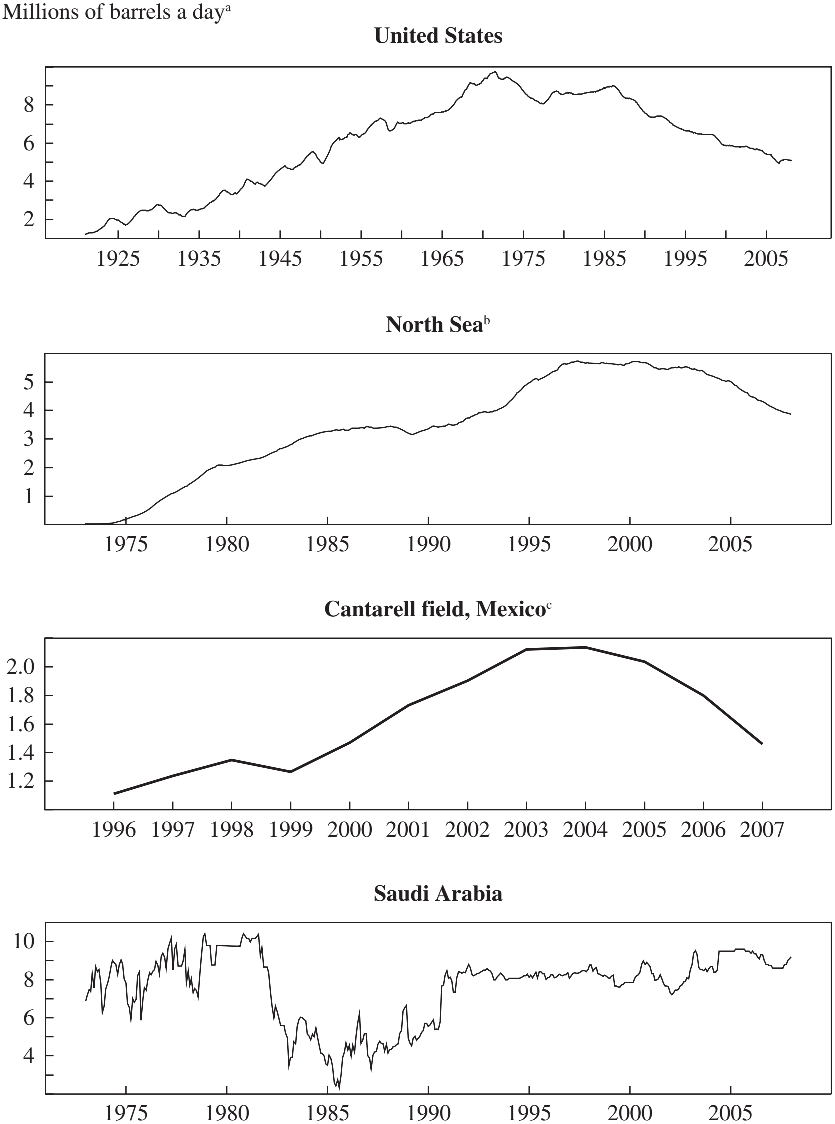

Sources: Adapted from figures 11, 13, and 14 in Hamilton (2009), using data from EIA, *Monthly Energy Review*, tables 11.1a and 11.1b; Pemex, *Statistical Yearbook 2007*, "Oil Production by Region"; Green Car Congress, "Mexico's Cantarell Continues Steep Decline in Oil Production in 2007" (www.greencarcongress.com/2008/01/mexicos-cantare.html).

a. Monthly data except where noted otherwise.

b. Moving average of the preceding 12 months of production by the United Kingdom and Norway.

c. Annual data.

in Saudi production over time. For example, the International Energy Agency's *World Energy Outlook 2007* was projecting that the Saudis would be pumping 12 million barrels a day (mbd) by 2010. In the event, Saudi production for 2007 was about 850,000 barrels a day lower than it had been for 2005.

It is a matter of conjecture whether the decline in Saudi production should be attributed to depletion of the country's Ghawar oilfield (the world's largest), to a deliberate policy decision in response to a perceived decline in the price elasticity of demand, or to the long-run considerations discussed below. Whatever its cause, the decline was certainly one important factor contributing to the stagnation in world oil production over 2005–07. It also unambiguously marks the beginning of a new era for oil pricing dynamics: without the Saudis' willingness or ability to adjust production to smooth out price changes, any disturbance to supply or demand will have a much larger effect on prices than in earlier periods.

III.B. Demand

Even as global supply stagnated, global demand was growing strongly. Particularly noteworthy is oil consumption in China, which has been growing at a 7 percent compound annual rate (calculated logarithmically) since 1990 (figure 9). Chinese consumption in 2007 was 870,000 barrels a day higher than just two years earlier.

Figure 9. Oil Consumption, China, 1980–2007

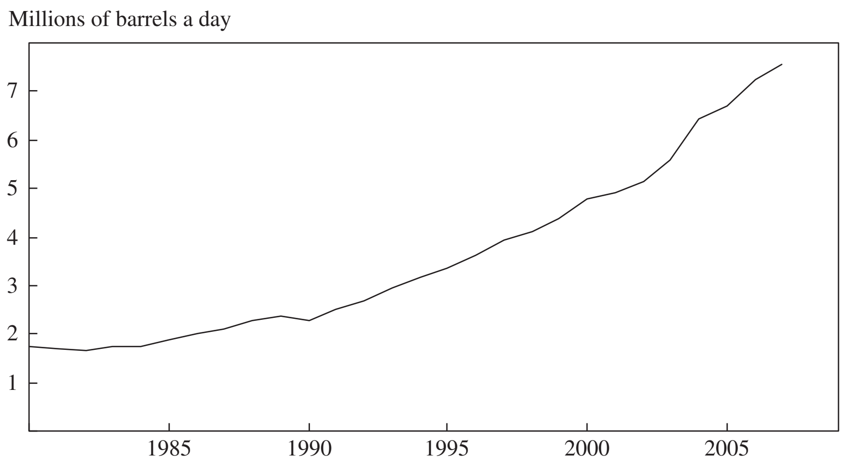

Source: EIA, "World Petroleum Consumption, Most Recent Annual Estimates, 1980–2007" (www.eia.doe.gov/ emeu/international/RecentPetroleumConsumptionBarrelsperDay.xls).

With global oil production flat, China's increased consumption meant that consumption in other regions had to decline. U.S. consumption in 2007 was 122,000 barrels a day below its level in 2005; meanwhile Europe's daily consumption dropped by 346,000 barrels and Japan's by 318,000 barrels. What persuaded residents of these countries to reduce oil consumption in the face of rising incomes? The answer is a rising price of oil.

Consider some ballpark estimates of the size of the price increase required to reduce consumption in the rest of the world by enough to accommodate China's increased consumption. According to the International Monetary Fund (*World Economic Outlook: October 2008,* table A1), real gross world product grew by a total of 9.4 percent in 2004 and 2005. As noted above, the income elasticity of petroleum demand in countries like the United States is currently about 0.5, whereas in the newly industrialized countries it may be above unity (Hamilton 2009; Gately and Huntington 2002). World petroleum production, at 85 mbd, was 5 mbd higher in 2005 than in 2003, a 6 percent increase. Thus, it is entirely plausible to attribute the 6 percent increase in world oil consumption between 2003 and 2005 to a shift in the demand curve caused by growth in world income.

Real gross world product grew an additional 10.1 percent in 2006 and 2007. Hence, it seems reasonable to suppose that if oil had remained at the 2005 price of $55 a barrel, the quantity demanded would have increased by at least another 5 mbd by the end of 2007. Economic growth slowed significantly in the first half of 2008 but remained positive, and one can conservatively assume that economic growth would have added at least another 0.5 mbd to the quantity demanded in that period, more than enough to absorb the slight increase in global oil production that finally appeared in the first half of 2008. Under these assumptions, the price had to rise between 2005 and the first half of 2008 by an amount sufficient to reduce the quantity that otherwise would have been demanded by 5 mbd, that is, to 85.5 mbd (top panel of figure 10).

It is worth commenting on what was new about the contribution of Chinese and world economic growth over this period. Although China's economy has been growing at a remarkable rate for a quarter century, it has only recently become big enough relative to the global economy to make a material difference. Compare, for example, the 4.9 percent average annual growth in world product achieved during 2004–07 with the 2.9 percent average over the 1990s, a period of robust growth but before China emerged as an important contributor. And judging from the reported gap

Figure 10. Short-Run Price Determination in the World Oil Market, 2005–08[a]

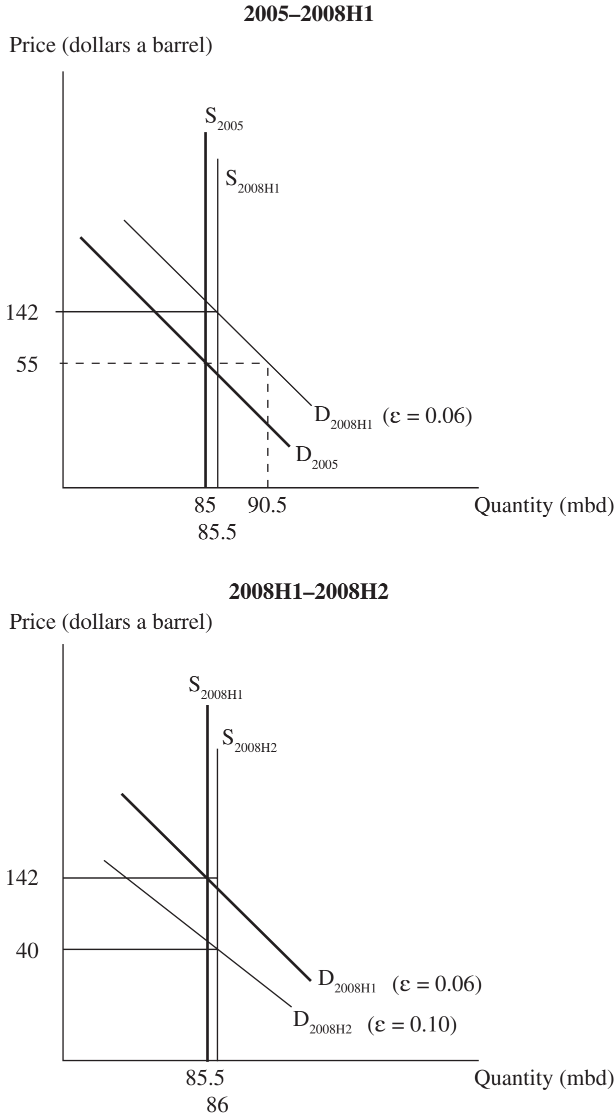

Source: Author's model described in the text.

a. Depictions of slopes are exaggerated for clarity. H1 and H2, first and second half of year.

between China's total petroleum production and its consumption,[5] China was a net exporter of petroleum up through 1992, and its imports were still only 800,000 barrels a day in 1998. By 2007, however, China's net petroleum imports were estimated to be 3.7 mbd, making it the world's third-largest importer and a dominant factor in world markets.

Global growth in petroleum demand in recent years is thus quite exceptional. And although there have been other episodes when global oil production stagnated over a two-year period, these were inevitably either responses to falling demand during recessions or the physical supply disruptions detailed above.

Although figure 10 is drawn with vertical short-run supply curves, the analysis here does not require any particular assumptions about the short-run supply elasticity. I simply take it as an observed fact that as a result of some combination of shifts of or movements along the short-run supply curve, the quantity supplied in the first half of 2008 was essentially the same as in 2005, and that the price and output for each date represent an intersection of supply and demand. The above exercise simply explores the necessary adjustments if the strong growth of gross world product between the two periods is presumed to have shifted the demand curve to the right by 5.5 mbd. The question is then, What price increase would have been necessary to move the quantity demanded along that second demand curve to a point where it would have been as low as 85.5 mbd?

The answer depends, of course, on the slope of that second demand curve. If, for illustration, the short-run price elasticity of demand were $\varepsilon = 0.06$, then the price would have been predicted to rise to \$142 a barrel under the above scenario:

$$\varepsilon = \frac{|\Delta \ln Q|}{\Delta \ln P} = \frac{\ln 90.5 - \ln 85.5}{\ln 142 - \ln 55} = 0.06.$$

On the other hand, such numerical calculations are extremely sensitive to the assumption one makes about the elasticity. If instead the elasticity were $\varepsilon = 0.10$, the price would only have needed to rise to \$97 a barrel to prevent global quantity demanded from increasing.

5. EIA, "World Petroleum Consumption, Most Recent Annual Estimates, 1980–2007" (www.eia.doe.gov/emeu/international/RecentPetroleumConsumptionBarrelsperDay.xls) and "World Production of Crude Oil, NGPL, and Other Liquids, and Refinery Processing Gain, Most Recent Annual Estimates, 1980–2007" (www.eia.doe.gov/emeu/international/Recent TotalOilSupplyBarrelsperDay.xls).

Figure 11. Crude Oil Stocks of U.S. Refiners, 2007, 2008, and Historical Average[a]

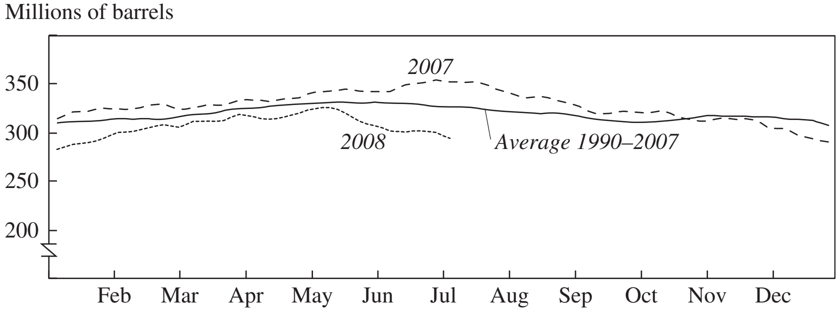

Source: EIA, "U.S. Stocks of Crude Oil and Petroleum Products" (tonto.eia.doe.gov/dnav/pet/xls/pet_stoc_wstk_dcu_nus_w.xls).
a. Excludes the Strategic Petroleum Reserve.

Which is closer to the correct short-run elasticity, 0.06 or 0.10? Recalling tables 1 and 2, one could easily defend either value, or numbers significantly smaller or bigger. Moreover, as noted by Hughes, Knittel, and Sperling (2008), the relevant elasticity for 2007–08 could have been much smaller than those that governed other episodes. One key variable to help answer this question is the value of inventories. If the price increase between 2005 and the first half of 2008 was greater than needed to equate supply with demand, inventories should have been piling up, whereas if the price increase was too small, inventories would have been drawn down.

Reliable data on worldwide stores of oil are lacking, but reasonably good measures are available on the inventories of crude oil held by U.S. refiners. Figure 11 plots the average seasonal pattern of these inventories, along with the actual values in 2007 and 2008. In the first half of 2007, inventories were a bit above trend. But in late 2007 and the first half of 2008, when the price increase was most dramatic, inventories were significantly below normal, suggesting that an assumed elasticity of 0.10 was too high and that the increase in price through the end of 2007 was not sufficient to bring quantity demanded down to equal quantity supplied.

Just as academics debate the correct value for the price elasticity of crude oil demand, so do market participants. Many observers have wondered what could have been the nature of the news that sent the price of oil from $92 a barrel in December 2007 to its all-time high of $145 a barrel in July 2008. Clearly it is impossible to attribute much of this move to a major surprise, such as that economic growth in the first half of 2008 was faster than expected or that oil production gains were more modest than anticipated. The main uncertainty, I would argue, was over the value of ε. The big news of the first half of 2008 was the surprising observation that even $100 oil was not going to be sufficient to prevent global quantity demanded from increasing above 85.5 mbd, and that no more than 85.5 mbd was going to be available.

This explanation of the price shock also requires that market participants could have had little inkling in the first half of 2008 of the massive economic deterioration just ahead. In this they certainly would have had good company. For example, here is the analysis offered publicly by European Central Bank president Jean-Claude Trichet on July 3, 2008:[6]

> On the basis of our regular economic and monetary analyses, we decided at today's meeting to increase the *key ECB interest rates* by 25 basis points. . . . [Inflation is] expected to remain well above the level consistent with price stability for a more protracted period than previously thought. . . . While the latest data confirm the expected weakening of real GDP growth in mid-2008 after exceptionally strong growth in the first quarter, the economic fundamentals of the euro area are sound.

And although a growth slowdown in the United States was certainly acknowledged at that point, many were unpersuaded that it would become serious enough to qualify as a true recession. Edward Leamer (2008, p. 29), for example, wrote in August that U.S. economic indicators would "have to get much worse to pass the recession threshold."

One may be able to rationalize the dramatic oil price spike of 2007–08 as a potentially appropriate response to fundamentals. But what about the even more dramatic subsequent collapse? Certainly Trichet, Leamer, and everyone else changed their minds about those assessments of real economic activity as the disastrous economic news of the fall of 2008 came in. But economic collapse alone is not a sufficient explanation for the magnitude of the oil price decline, if the analysis in the top panel of figure 10 is correct. Even a 10 percent drop in global economic activity would undo only the effects of the rightward shift of the demand curve since 2005. Bad as the news in the second half of 2008 was, it did not come close to that magnitude, yet the price of oil by the end of December was down to $40 a barrel, well below the 2005 price of $55. Instead, one would need to attribute a significant part of the price collapse to yet another shift in the elasticity. Whereas a short-run price elasticity of 0.06 might be needed to make sense of the developments of the first half of 2008, a

6. European Central Bank, "Introductory Statement with Q&A," (www.ecb.int/press/pressconf/2008/html/is080703.en.html, emphasis in original).

higher intermediate-run elasticity, as petroleum users made delayed adjustments to the earlier price increases, has to be postulated as another factor contributing to the price decline in the second half of the year (bottom panel of figure 10).

It is hardly controversial to suggest that the long-run demand responses to price increases are more significant than the short-run responses. The more-fuel-efficient vehicles sold in the spring and summer of 2008 are going to mean lower consumption, at least from those vehicles, for many years to come. The Energy Information Administration (EIA) reported that U.S. petroleum and petroleum products supplied in 2008Q3 were 9.6 percent lower (logarithmically) than in 2007Q3, a far larger drop in percentage terms than the presumed 6.3 percent rightward shift from the 2005 to the first-half 2008 world demand curve assumed in the top panel of figure 10, and again far in excess of anything attributable to the drop in income alone.

III.C. The Role of Speculation

One can thus tell a story of the oil price shock and subsequent collapse that is driven solely by fundamentals. But the speed and magnitude of the price collapse lead one to give serious consideration to the alternative hypothesis that this episode represents a speculative price bubble that popped. One proponent of this view has been Michael Masters, manager of a private financial fund, who has been invited a number of times to testify before the U.S. Senate. Masters (2008) blames the oil price spike of 2007–08 on the actions of investors who bought oil not as a commodity to use but as a financial asset; he claims that by March 2008, commodity index trading funds held a quarter trillion dollars worth of futures contracts. A typical strategy is to take a long position in a near-term futures contract, sell it a few weeks before expiry, and use the proceeds to take a long position in a subsequent near-term futures contract. When commodity prices are rising, the sell price should be higher than the buy, and the investor can profit without ever physically taking delivery. As more investment funds sought to take positions in commodity futures contracts for this purpose, so that the number of buys of next contracts always exceeded the number of sells of expiring ones, the effect, Masters argues, was to drive up the futures price, and with it the spot price. This “financialization” of commodities, according to Masters, introduced a speculative bubble in the price of oil.

The key intellectual challenge for such an explanation is to reconcile the proposed speculative price path with what was happening to the physical quantities of petroleum demanded and supplied. To be concrete about the nature of this challenge, consider a representative refiner who purchases a quantity Z_t of crude oil at price P_t per barrel, of which X_t is used up in current production of gasoline and the remainder goes to increase inventories I_t:

(1) $$I_{t+1} = I_t + Z_t - X_t.$$

This is simply an accounting identity: if the quantity of oil consumed by users (in this case, X_t) is smaller than the quantity physically produced (Z_t), inventories must accumulate. If one hypothesizes that, by whatever process, financial speculation results in some particular price P_t, that price necessarily has implications for those who use the product (X_t) and those who produce it (Z_t). It seems impossible to discuss a theory of price P_t that makes no reference to the physical quantities produced, consumed, or held in inventory.

To explore this issue more fully, consider the following simple model. Suppose that the refiner produces a quantity of gasoline y_t, which it sells at price G_t (where both P_t and G_t are measured in real terms), according to the production function $y_t = F(X_t, I_t)$. The second term reflects the idea that refiners cannot operate efficiently with zero inventories. A positive value for the derivative $F_I(X_t, I_t)$ introduces a "convenience yield" from inventories—in other words, a motive for the firm to hold a positive level of inventory even if it anticipates falling crude oil prices ($P_{t+1} < P_t$). The refiner faces a real interest rate of r_t and a cost of physically holding inventories $C(I_{t+1})$. The refiner's objective is thus to choose $\{Z_t, X_t, I_{t+1}\}_{t=0}^{N}$ so as to maximize

$$\sum_{t=0}^{N} \frac{1}{\prod_{\tau=1}^{t}(1 + r_\tau)} \left[G_t F(X_t, I_t) - C(I_{t+1}) - P_t Z_t\right],$$

taking I_0 and $\{P_t, G_t\}_{t=0}^{N}$ as given. Note that I pose this as a perfect-foresight problem, since the complications introduced by uncertainty are not relevant for the points I want to make here, and liquid futures markets exist for P_t and G_t.

The first-order conditions for this optimization problem are[7]

(2) $$G_t F_X(X_t, I_t) = P_t$$

(3) $$P_t + C'(I_{t+1}) = \frac{1}{(1 + r_t)}\left[G_{t+1} F_I(X_{t+1}, I_{t+1}) + P_{t+1}\right].$$

7. Specifically, the values of $\{X_t, Z_t, I_{t+1}\}_{t=0}^{N}$ are determined as functions of $\{G_t, P_t\}_{t=0}^{N}$ from equation 2 for $t = 0, \ldots, N$; equation 1 for $t = 0, \ldots, N$; equation 3 for $t = 0, \ldots, N - 1$, and the terminal condition $I_{N+1} = 0$.

Equation 2 is the optimality condition associated with the firm purchasing one more barrel of crude oil whose marginal cost is P_t and using the crude immediately to refine and sell more gasoline, whose marginal benefit to the firm is $G_t F_X(X_t, I_t)$. Equation 3 is the condition required for optimal inventory management. If the firm buys one more barrel of crude today to store as inventory, the marginal cost is $P_t + C'(I_{t+1})$. If the inventory is then used to reduce next period's crude purchases, the discounted marginal benefit is $(1 + r_t)^{-1}[G_{t+1}\, F_I(X_{t+1}, I_{t+1}) + P_{t+1}]$.

If the firm were to face an increase in price ($P_{t+1} > P_t$), with all other prices remaining constant, it would respond by increasing I_{t+1} until the condition given in equation 3 was restored. This plan would be implemented by increasing current crude purchases Z_t and decreasing Z_{t+1}. In the market equilibrium that I will finish spelling out shortly, that would put upward pressure on P_t and downward pressure on P_{t+1}. But it is interesting to comment here on the limiting case of a constant physical storage cost, $C'(I_{t+1}) = h,$ and constant convenience yield, $F_I(X_{t+1}, I_{t+1}) = c,$ the latter including as a special case zero convenience yield, $F_I(X_{t+1}, I_{t+1}) = 0$, that is, a situation where inventories are already so high that no sales gain would accrue from building them even higher. In this case equation 3 becomes

$$P_t + h = \frac{1}{(1 + r_t)}\left[G_{t+1}c + P_{t+1}\right]. \tag{4}$$

In this limiting case, equation 4 becomes an equilibrium condition that would have to characterize the relationship between P_t and P_{t+1} in any equilibrium with nonzero inventories. If, for example, the right-hand side of equation 4 exceeded the left, there would be an infinite increase in the demand for crude Z_t and an infinite decrease in Z_{t+1}, to which the equilibrium prices P_t and P_{t+1} would have to respond until equation 4 was restored.

More generally, if $C'(I_t)$ and $F_I(X_t, I_t)$ are relatively flat functions of I_t, then the effect of equation 3 is to force P_t and P_{t+1} to move closely together. In crude oil markets, the futures price corresponding to P_{t+1} serves an information discovery role, any changes in which translate instantaneously into a corresponding movement in spot prices. For example, figure 12 plots f_{1d}, the price of crude oil for the nearest-term futures contract on day d, and f_{3d}, the price of oil for the futures contract expiring two months after the expiration of the contract associated with f_{1d}. The two series move very closely together. On only 7 percent of the 6,421 business days between April 5,

Figure 12. Prices of Oil Futures Contracts, 1983–2008

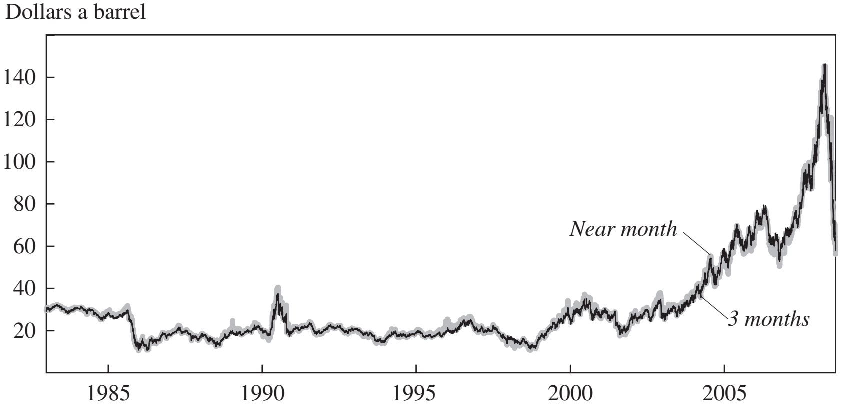

Source: EIA, "NYMEX Futures Prices" (tonto.eia.doe.gov/dnav/pet/pet_pri_fut_s1_d.htm).

1983, and November 12, 2008, were the changes Δf_{1d} and Δf_{3d} of opposite signs. A regression of $\Delta \ln f_{3d}$ on $\Delta \ln f_{1d}$ has an R^2 of 0.86. Thus, this part of Masters' claim—that if speculation affected the futures price, the spot price would be forced to move with it—is very much consistent with both theory and evidence.

The model can be closed by specifying that crude oil is exogenously supplied,

$$Z_t = \bar{Z}_t, \tag{5}$$

and that gasoline demand has a price elasticity of β:

$$\ln F(X_t, I_t) = \gamma - \beta \ln G_t. \tag{6}$$

The system of equations 1–3, 5, and 6 then determines $\{Z_t, X_t, I_{t+1}, P_t, G_t\}_{t=0}^{N}$ as functions of $\{\bar{Z}_t\}_{t=0}^{N}$.

Notice that if the marginal storage cost $C'(I_{t+1})$ is negligible, then equations 1–3 and 5 are homogeneous of degree 0 in $\{P_t, G_t\}_{t=0}^{N}$. Without equation 6—that is, if there were no response of gasoline demand to the price of gasoline—the price of crude oil would be indeterminate. Suppose the market was initially in a situation where all five equations were satisfied, and consider the limiting case where the demand for gasoline is perfectly price inelastic ($\beta = 0$). Suppose now that for some reason speculators bid up the futures price of crude (P_{t+1} increases). By inventory arbitrage (equation 3),

P_t would have to rise with it. In this sense, one might claim to have a theory of how financial speculation in P_{t+1} could be the determining factor in the price of oil.

On the other hand, when the price elasticity $\beta > 0$, the above analysis no longer applies. In response to the hypothesized increase in P_{t+1} and P_t, the price of gasoline G_t would rise (from equation 2), the quantity of gasoline demanded would decline, and the amount of crude X_t that must be used to meet this demand would be lower. An increase in P_{t+1} and P_t induced by speculation would thus cause crude inventories I_{t+1} to accumulate relative to the firm's desired path.

If the price elasticity is small but not zero, this feedback will be subtle, and it might conceivably take some time before mispricing arising from the futures markets would be recognized and corrected. It is interesting to note, however, that the same condition needed to rationalize a speculation-based interpretation of the oil shock of 2007–08—a very low price elasticity of oil demand—is exactly the same condition that allows the event to be attributed to the fundamentals alone.

The other possible way in which advocates of the price bubble interpretation might attempt to reconcile their story with the physical side of the petroleum market is by hypothesizing a mechanism whereby the quantity of oil supplied $\bar{Z}_t$ is itself influenced by the futures price. Given the pressures for growth in petroleum demand from countries like China to continue, if it remains difficult to increase global production, the price pressures of 2008 are only the beginning of the story. Recalling the Hotelling (1931) principle, it would in this situation pay the owners of the resource to forgo current production, in order to be able to sell the oil at the higher future price. One might then hypothesize that oil-producing countries were misled by the speculative purchases of oil futures contracts into reducing current production $\bar{Z}_t$ in response, by this mechanism reconciling the postulated speculation with the physical dynamics of oil supply and demand (equation 1).[8]

If so, such miscalculation by oil producers could not have been based on comparing the longer-term futures price with the spot price in 2008. Figure 13 plots the term structure of prices implied by futures contracts on the New York Mercantile Exchange (NYMEX) at the price peak in July 2008. Prices of the very near term contracts sloped modestly upward (for example, the December 2008 contract sold for a higher price than

8. For further discussion see Jovanovic (2007).

Figure 13. Term Structure of Oil Futures Contracts, July 2008[a]

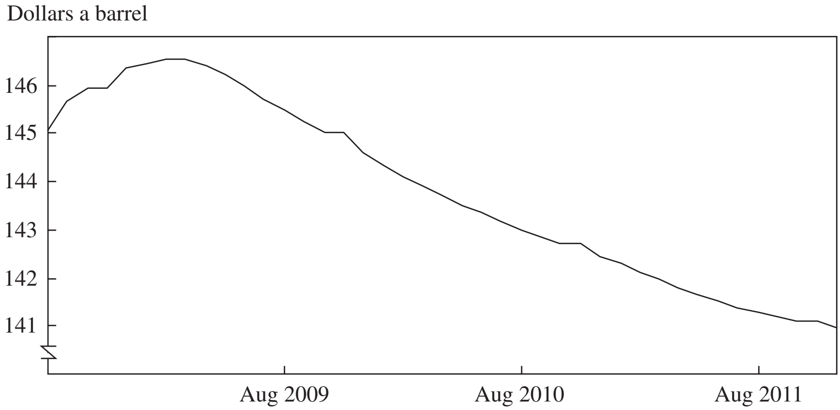

Source: Norma's Historical Data.

a. Closing prices on July 11, 2008, of NYMEX light sweet crude contracts for settlement in the indicated month.

that for August 2008), but that slope turned distinctly downward after the February 2009 contract, meaning that any producers who used the futures markets to sell their oil forward could expect to sell future production at a lower price than current production. This downward slope from early 2009 onward is inconsistent with a natural Hotelling interpretation of why producers might keep oil in the ground. One might argue instead that producers distrusted the futures markets or felt they could not use them as a significant hedge given the volumes involved. In any case, the high spot price in 2008 meant that a country that had held off production from 2001 to 2008 would have been richly rewarded, and that experience might have persuaded some producers of the benefits of not producing at capacity in 2008 either. Of interest is this Reuters report from April 2008:

> Saudi Arabia's King Abdullah said he had ordered some new oil discoveries left untapped to preserve oil wealth in the world's top exporter for future generations, the official Saudi Press Agency (SPA) reported. "I keep no secret from you that when there were some new finds, I told them, 'no, leave it in the ground, with grace from god, our children need it'," King Abdullah said in remarks made late on Saturday, SPA said.[9]

9. Reuters News Service. "Saudi King Says Keeping Some Oil Finds for Future," April 13, 2008 (uk.reuters.com/article/oilRpt/idUKL13968772008041З).

With hindsight, it is hard to deny that the price of oil rose too high in July 2008 and that this miscalculation was influenced in part by the flow of investment dollars into commodity futures contracts. It is worth emphasizing, however, that the two key ingredients needed to make such a story coherent—a low price elasticity of demand, and the failure of physical production to increase—are the same key elements of an explanation of the same phenomenon based only on fundamentals. I therefore conclude that these two factors, rather than speculation per se, should be construed as the primary cause of the oil shock of 2007–08. Certainly the casual conclusion that one might have drawn from glancing at figure 1 and hearing some of the accounts of speculation[10]—that it was all just a mistake, and the price should have stayed at \$50 a barrel throughout 2005–08—would be profoundly in error.

IV. Consequences of Past Oil Shocks

In essentially any theoretical model of the economic effects of a change in oil prices, a key parameter is the dollar value of energy purchases as a share of total costs, such as the consumption spending share plotted in figure 3. To see why this is a key parameter, consider a firm producing output Y_t with inputs of capital K_t, labor N_t, and energy E_t. Suppose that the firm is operating at a point where the marginal product of energy is equal to its relative price:

$$\frac{\partial F(K_t, N_t, E_t)}{\partial E_t} = P_t. \tag{7}$$

Multiplying both sides of equation 7 by $E_t/F(K_t, N_t, E_t)$ establishes that the elasticity of output with respect to energy is given by the value share,

$$\frac{\partial \ln F(K_t, N_t, E_t)}{\partial \ln E_t} = \alpha_t$$

for $\alpha_t = P_t E_t / F(K_t, N_t, E_t)$. Alternatively, consider a consumer facing a π percent increase in the relative price of energy. One short-run option for the

10. Such a conclusion is implied, for example, by the Obama campaign's website in June 2008, which included a number of quotes from analysts such as Shell President John Hofmeister that the proper range of crude oil is "somewhere between \$35 and \$65 a barrel." See the weblog Econbrowser, "How Big a Contribution Could Oil Speculation Be Making?" (www.econbrowser.com/archives/2008/06/how_big_a_contr.html) for details.

consumer (and indeed, given the empirical evidence reviewed above, not a bad approximation to what actually happens) is to continue to purchase the same quantity of energy as before. If income cannot be increased, this would require that the consumer either reduce saving or cut spending on other items. If α_t denotes the consumer's energy expenditure share, the requisite percentage cut in spending on other items is given by $\alpha_t\pi$.

A large literature has investigated the economic consequences of previous oil price shocks. Recent refinements include investigations of the following: nonlinearity in the relationship, with oil price increases possibly having a bigger effect than oil price decreases (for example, Hamilton 2003); differences in the causes, with price increases brought about by surging global demand possibly having less of a disruptive effect than those caused by losses in supply (for example, Kilian 2009); and a changing relationship over time, such that the economy is more resilient to an oil price shock today than in the past (for example, Blanchard and Galí 2008).

Although these issues are unquestionably quite important, it is useful to look first at some simple linear representations of the basic correlations in the historical data, with a minor automatic adjustment for one source of a possible changing impact over time, namely, changes in α_t. This is the approach taken by Edelstein and Kilian (2007), who estimate monthly bivariate autoregressions of the form

$$x_t = k_1 + \sum_{s=1}^{6} \phi_{11} x_{t-s} + \sum_{s=1}^{6} \phi_{12} y_{t-s} + \varepsilon_{1t}$$

$$y_t = k_2 + \sum_{s=1}^{6} \phi_{21} x_{t-s} + \sum_{s=1}^{6} \phi_{22} y_{t-s} + \varepsilon_{2t},$$

where y_t is a macroeconomic variable of interest and x_t is the change in relative energy prices weighted by the expenditure share,

$$x_t = \alpha_t \left(\ln P_t - \ln P_{t-1}\right),$$

α_t is the series plotted in figure 3, and P_t the ratio of the personal consumption expenditure (PCE) deflator for energy goods and services to the overall PCE deflator. Thus, for example, a unit shock to x_t would result from a monthly 20 percent increase in relative energy prices ($\ln P_t - \ln P_{t-1} = 0.20$) at a time when energy consumes 5 percent of household budgets ($\alpha_t = 5.0$). A unit shock to x_t means that households would suffer a 1 percent loss in their ability to purchase nonenergy items if they attempted to hold real energy consumption fixed following a shock of size $x_t = 1$.

I reestimated a number of the Edelstein-Kilian regressions for the sample period they used (with the dependent variable running from July 1970 through July 2006). I begin by reporting the results for $y_t = 100(\ln Y_t - \ln Y_{t-1})$, where Y_t is real PCE. The top two panels in figure 14 reproduce their orthogonalized impulse-response functions (with energy prices x_t ordered first) for the cumulative consequences for the cumulative price change ($X_t = \Sigma_{j=1}^{t} x_j$, in the left panel) and for real PCE (100 ln Y_t, in the right panel) of a unit shock to x_{t-s}.

The left panel shows that there is relatively little serial correlation in the energy price change series x_t. Almost all of the price consequences appear within the first two months: if x_t increases by one unit at time t, one would typically expect another 0.5-unit move up at $t + 1$, with very minor adjustments thereafter, resulting in an eventual 1.7 percent cumulative loss in purchasing power. Two aspects of the right panel differ from what one would have expected given the simple expenditure impact sketched above. The first is the magnitude of the real PCE response: following a decline that eventually would have reduced consumers' ability to purchase nonenergy items by 1.7 percent, consumers on average eventually cut their spending by 2.2 percent. Why should consumption spending fall by even more than the predicted upper bound? The second surprising aspect concerns timing: although the change in price immediately reduces purchasing power, the biggest declines in total spending do not appear until six months or more after the initial shock.

One way that Edelstein and Kilian seek to explain these anomalies is by breaking down the responses into the various components of consumption. The three middle panels of figure 14 reproduce their findings for the services, nondurables, and durables components of real PCE. The first two responses are in line with the simple expenditure share effects, but the response of expenditure on durable goods is five times as large.

The bottom left panel of figure 14 focuses on the motor vehicles component of durables consumption. In contrast to the gradual response seen in the broader consumption categories, here the response is immediate and huge: for example, a 20 percent increase in energy prices, given an energy expenditure share of 5 percent, would result in a 10 percent decrease in spending on motor vehicles within two months. That a direct link might exist between such spending and energy prices is quite plausible, although not simply through the budget constraint effect. For this category of spending a number of other factors are much more important, such as consumers postponing purchases of new vehicles until better information about future

Figure 14. Responses to an Increase in Energy Prices That Reduces Disposable Income by 1 Percent[a]

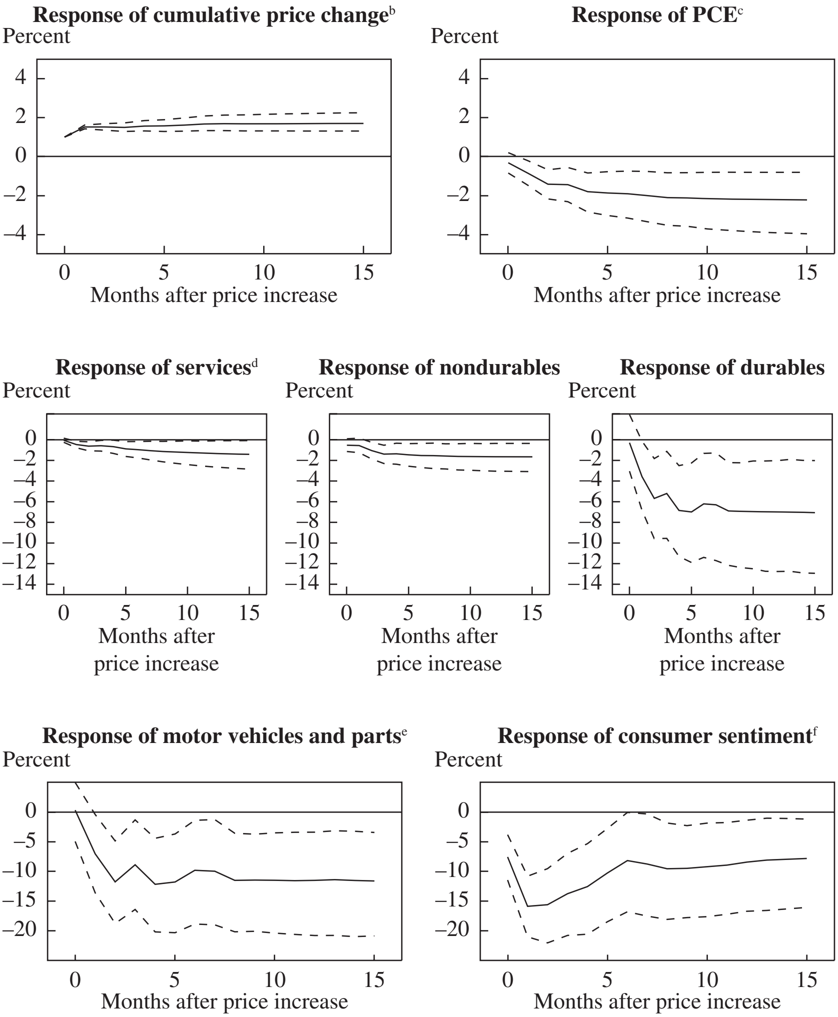

Source: Author's calculations.

a. Estimates from bivariate vector autoregressions based on data for July 1970 to July 2006. Dashed lines indicate 95 percent confidence intervals.

b. Response of $X_t = \Sigma_{j=1}^{t} x_j$ to a one-unit shock to x_{t-s} plotted as a function of s.

c. Response of 100 × the logarithm of real personal consumption expenditure at time t to a one-unit shock to x_{t-s} plotted as a function of s. Reproduces (with renormalization) figure 8a in Edelstein and Kilian (2007).

d. Panels in this row reproduce (with renormalization) figure 8b-d in Edelstein and Kilian (2007).

e. Panels in this row reproduce (with renormalization) figures 8e and 11a in Edelstein and Kilian (2007). See also footnote 10 in the text.

f. Unlike in the other panels, in which the second variable in the vector autoregression, $y_t = 100(\ln Y_t - \ln Y_{t-1})$, represents a rate of change (with the impulse response graphs subsequently translated back into implications for the levels $100 \ln Y_t$), here the variable y_t is the index of consumer sentiment itself (from the University of Michigan/Reuters), and the graph shows the consequences for y_{t+s} following a unit shock to x_t.

gas prices is available and shifting purchases to more-fuel-efficient (and perhaps less expensive) vehicles.

If one takes as given that there are large and immediate effects of oil price changes on purchases of certain items such as motor vehicles, both the delayed response and the multiplier effect on other categories of spending can be better understood. The shift in spending means a reduction in income for those employed in manufacturing and selling cars. Given the significant technological frictions in relocating the now-underutilized labor and capital to other sectors, the result is a decline in aggregate income and a loss in purchasing power over and above that caused by the initial price increase itself (Hamilton 1988).

The bottom right panel of figure 14 presents a second effect identified by Edelstein and Kilian that is also huge and immediate, namely, a drop in consumer sentiment. For whatever reason, consumers found the oil shocks of earlier decades to be very troubling events, with a 20 percent increase in relative energy prices (assuming again a base case value share of $\alpha_t = 5$) producing on average a 15-point drop in the index of consumer sentiment. One can argue whether a response of this magnitude is rational given the size of the shock. However, the budget consequences of spiking gasoline prices are something consumers experience immediately, and such a spike represents an aggregate event that forces everybody to make changes at the same time. Certainly for those whose jobs are related to the automobile industry (and for those who perceive that what happens in that industry will have eventual implications for their own job security), it is quite rational to view these events as carrying pessimistic implications beyond the immediate loss in spending power. In any case, the changes in sentiment observed in these data could easily have made a significant contribution to the subsequent path of both consumption and investment spending.

The bottom left panel of figure 14 showed the narrowest effect of the energy price shock, namely, the change in spending on motor vehicles and parts. How big a contribution would this alone have made to the subsequent economic downturns, ignoring any possible multiplier effects? The first column of the top panel of table 3 reports the actual average growth rate of real GDP over the five quarters following each of the four historical oil shocks discussed above. GDP fell over each of these four periods, and all four are included in the list of U.S. recessions. The second column reports results of a very simple calculation: what would average GDP growth have been if there had been zero change in the motor vehicles and parts component of GDP over these five quarters, with all other components of GDP

Table 3. Growth of Real GDP in Oil Shock Episodes under Alternative Scenarios
Percent a year[a]

			Without oil shock[c]	
Period	*Actual*	*Without automobiles*[b]	*Blanchard-Galí (2008)*	*Hamilton (2003)*
Pre-2000 recessions				
1974Q1–1975Q1	−2.5	−2.0	−0.1	+2.3
1979Q2–1980Q2	−0.4	+0.4	+0.4	+2.5
1981Q2–1982Q2	−1.5	−1.3	−2.0	+2.0
1990Q3–1991Q3	−0.1	+0.2	+0.5	+3.6
2007–08 recession				
2007Q4–2008Q3	+0.7	+1.2	+1.4	+4.2
2007Q4–2008Q4	−0.7	−0.0	−0.2	+3.2

Sources: Blanchard and Galí (2008); Hamilton (2003); Bureau of Economic Analysis data; author's calculations.

a. All data are annual averages.

b. Growth in real GDP excluding the contribution from motor vehicles and parts.

c. See the text for the details of the computation.

staying the same as reported?[11] Although motor vehicles and parts contributed modestly to the change in GDP (0.8 percentage point or less in any episode), it was enough to move the average from negative to positive territory in the cases of the 1979–80 and 1990–91 episodes. This offers some basis for thinking that without the significant downturn in the auto sector each of these two episodes might have been regarded as periods of sluggish growth rather than clear recessions. By contrast, in the more severe 1973–75 and 1981–82 recessions, clearly something more significant than just automobiles was bringing down the economy.

I next examine the implications of two earlier studies of the overall effects of oil prices. The first is that by Olivier Blanchard and Jordi Galí (2008), who conclude that oil shocks made a relatively modest contribution to the downturns of the 1970s and are even less important today. Their analysis is based on a vector autoregression (VAR) that has three nominal shocks in addition to oil prices (as captured by the consumer price index, the GDP deflator, and wages) and two output indicators (GDP and total hours worked). All variables are quarterly percentage changes, with the oil price measured as the average price of West Texas Intermediate (WTI) crude oil over the quarter, and a quadratic time trend is included. The

11. This is calculated by subtracting from the growth rate of real GDP the contribution of motor vehicles and parts as reported by the Bureau of Economic Analysis (National Income and Product Accounts, table 1.5.2). Note that this contribution is a negative number in each episode, so that subtracting it raises the GDP growth rate.

authors estimate two separate versions of the VAR, the first using data only from 1960Q1 to 1983Q1 and the second from 1984Q1 to 2007Q3.

I used the VAR coefficients as estimated from the separate subsamples to perform the following calculation. First, I constructed from the estimated coefficients of the VAR a dynamic forecast for each of the six variables for each episode based on information available (that is, the observed values of the six variables) as of the quarter preceding the episode. I then compared this forecast with the ex post realized values of these variables to generate an implied set of forecast errors for each variable one to five quarters ahead. Finally, I decomposed these observed errors into contemporaneously orthogonal components, based on the variance-covariance matrix used by Blanchard and Galí, to arrive at an answer to the following question: what would be the prediction error for each of the variables up to five quarters ahead if one could condition on the ex post realizations of the innovations in oil prices but did not know anything else?[12]

On the basis of these numbers, I calculated what average GDP growth over 1974Q1–1975Q1 would have been had there been no oil price shock but the other five shocks—to the consumer price index, the GDP deflator, wages, GDP, and hours worked—had been identical to the realized historical residuals. The answer to that "what if" question is reported in the third column of table 3. The Blanchard-Galí estimates imply that had there been no oil shock, the severe downturn of 1974–75 would have been only a very mild recession. Interestingly, although their estimated post-1984 effects of oil prices are much smaller than those for their earlier sample, and although the authors did not single out the aftermath of the 1990–91 Gulf War as a separate oil shock, their estimates also imply that had the price of oil not spiked following Iraq's invasion of Kuwait, the United States might have avoided the 1990–91 recession.

12. The estimated VAR coefficients imply a set of moving-average matrices $\hat{\psi}_s$ (as in equation 10.1.19 in Hamilton 1994), and the Cholesky factor of the residual variance matrix can be obtained as $\hat{\mathbf{\Omega}} = \hat{\mathbf{P}}\hat{\mathbf{P}}'$. The s-step-ahead forecast error can then be written $\mathbf{y}_{t+s} - \hat{\mathbf{y}}_{t+s|t-1} = \hat{\varepsilon}_{t+s} + \hat{\psi}_1\hat{\varepsilon}_{t+s-1} + \ldots + \hat{\psi}_s\hat{\varepsilon}_t$ for $\hat{\varepsilon}_t$, the implied one-step-ahead forecast errors. Define the orthogonalized residuals by $\hat{\mathbf{v}}_t = \hat{\mathbf{P}}^{-1}\hat{\varepsilon}_t$, and let $\hat{\mathbf{p}}_1$ denote the first column of $\hat{\mathbf{P}}$. Then the contribution of $\{\hat{v}_{1t}, \hat{v}_{1,t+1}, \ldots, \hat{v}_{1,t+s}\}$ to this forecast error is calculated as $\sum_{k=0}^{s} \hat{\psi}_k\hat{\mathbf{p}}_1\hat{v}_{1,t+s-k}$, and what $\mathbf{y}_{t+s}$ would have been in the absence of the oil shocks is calculated as $\mathbf{y}_{t+s} - \sum_{k=0}^{s}\hat{\psi}_k\hat{\mathbf{p}}_1\hat{v}_{1,t+s-k}$. Note that although the VAR shocks to oil prices and the consumer price index (CPI) are correlated in the data ($\hat{\varepsilon}_{1t}$ is correlated with $\hat{\varepsilon}_{2t}$), the shocks $\hat{v}_{1t}$ and $\hat{v}_{2t}$ are orthogonal in the sample by construction. Thus, when asking what would have happened if $\hat{v}_{1t}$ had been zero but $\hat{v}_{2t}$ had been as observed historically, one is implicitly subtracting out that movement in the CPI that is correlated statistically with the oil price and leaving in other, uncorrelated factors.

Surprisingly, the Blanchard-Galí estimates imply that the 1981–82 downturn would actually have been *more* severe in the absence of disturbances to oil prices. The explanation is that their measure for the price of oil, the price of WTI, actually *fell* between July 1980 and March 1981. Other indicators suggest a very different story: the EIA's series for the refiner acquisition cost (the series plotted in the middle panel of the third row in figure 5) shows a 27 percent (logarithmic) increase over this same period, the Bureau of Labor Statistics' producer price index (PPI) for crude petroleum (the series used in Hamilton 1983 and 2003) shows a 42 percent increase, the Bureau of Economic Analysis' implicit price deflator for consumption expenditure on energy goods and services (the series used by Edelstein and Kilian 2007) shows a 12 percent increase, and the Bureau of Labor Statistics' (BLS) seasonally adjusted consumer price index for gasoline shows a 14 percent increase. It therefore seems likely that despite the results implied by Blanchard and Galí's estimation, energy prices were a factor reducing GDP growth in this episode as in the others.

For another comparison I turned to the nonlinear specification investigated in Hamilton (2003), where the key result (equation 3.8) was obtained from a regression of quarterly real GDP growth on a constant, four of its own lags, and four lags of the "net oil price increase," defined as the percentage change in the crude oil PPI relative to its previous three-year high if oil prices made a new three-year high in quarter *t,* and zero if oil prices ended the quarter lower than at any point over the previous three years. The coefficients for that relationship, estimated over 1949Q2–2001Q3, were as follows:

$$(8) \qquad y_t = \underset{(0.13)}{0.98} + \underset{(0.07)}{0.22}\, y_{t-1} + \underset{(0.07)}{0.10}\, y_{t-2} - \underset{(0.07)}{0.08}\, y_{t-3} - \underset{(0.07)}{0.15}\, y_{t-4}$$
$$- \underset{(0.014)}{0.024}\, o^{\#}_{t-1} - \underset{(0.014)}{0.021}\, o^{\#}_{t-2} - \underset{(0.014)}{0.018}\, o^{\#}_{t-3} - \underset{(0.014)}{0.042}\, o^{\#}_{t-4}.$$

To get a sense of the magnitudes implied by these coefficients,[13] I calculated for each quarter in the episode the difference between the one-quarter-ahead forecast implied by equation 8 and what that one-quarter-ahead forecast would have been if the oil price measure $o^{\#}_{t-1}, \ldots, o^{\#}_{t-4}$ had instead been equal to zero, and I took this difference as a measure of the

13. One could in principle find the answer to an *s*-period-ahead forecasting equation as in the preceding footnote, although this would require a specification of the dynamic path followed by the net oil price increase variable. No such specification was proposed in Hamilton (2003), and it seems unlikely that spelling one out would change the results significantly from the simpler calculation reported here.

contribution of the oil shock to that quarter's real GDP growth. From the fourth column of table 3, it appears that this specification would attribute almost all of the deviation from trend in each of the four recessions to the oil shock alone.

To summarize, a range of estimates exist of the contribution of oil shocks to past U.S. recessions. But even the most modest estimates support the claim that the oil shocks made a significant contribution in at least some of these episodes. My conclusion is that had the oil shocks not occurred, GDP would have grown rather than fallen in at least some of these episodes.

V. Consequences of the Oil Shock of 2007–08

This section explores the contribution that oil prices may have made to the most recent recession. I begin by examining what happened to motor vehicle sales in response to the price changes depicted in figure 6. Figure 15 reports sales in the United States of domestically manufactured light vehicles, both cars (top panel) and light trucks (bottom panel). The latter include sport utility vehicles (SUVs), which through 2007 were outselling cars in the U.S. market. Beginning in 2008, sales of SUVs began to plunge, and in May, June, and July they were down more than 25 percent from the same months a year earlier. SUV sales rebounded somewhat when gas prices began to fall in August, only to suffer a second hit in September through December.

To what extent was the decline in SUV sales in the first half of 2008 caused by rising gasoline prices as opposed to falling income? One measure relevant for addressing this question is the contrast between sales of light trucks and those of cars. A general drop in income would affect both categories, whereas the effects of rising gasoline prices would hit light trucks much harder than cars. In the event, domestic car sales were down on average by only 7 percent in May, June, and July 2008 compared with the same months in 2007. Even more dramatic are the comparisons for imports. Imported cars were up 10 percent over these same three months (top panel of figure 16); sales of imported light trucks (bottom panel), by contrast, were down 22 percent. Thus, the dominant story in the first half of 2008 was one of American consumers switching from SUVs to smaller cars and more-fuel-efficient imports.

Although gasoline prices were likely a key factor behind plunging sales for U.S. automakers in the first half of 2008, falling income appears to have been the biggest factor driving sales back down in the fourth quarter of 2008. Then, in contrast to the first half of 2008, the sales decline was

Figure 15. Sales of Domestically Produced Cars and Light Trucks, by Month, United States, 2004–08[a]

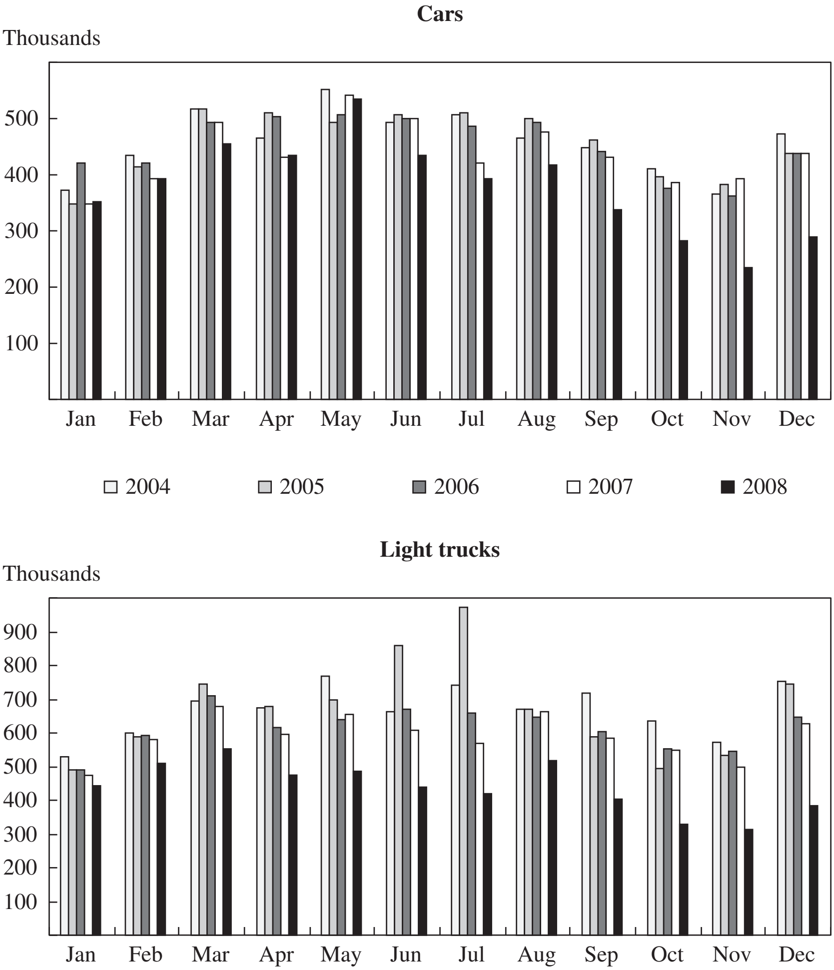

Source: WardsAuto.com, "Key Automotive Data" (wardsauto.com/keydata).
a. All vehicles produced within North America and sold in the United States.

across the board, affecting cars if anything more than SUVs, and imports along with domestically produced vehicles.

The result was a significant shock to the auto industry in the first half of 2008, quite comparable in magnitude to what was observed in the wake of the oil shock of 1990. The contribution of motor vehicles and parts to real GDP (measured in 2000 dollars at an annual rate) was $30 billion smaller

Figure 16. Sales of Imported Cars and Light Trucks, by Month, United States, 2004–08[a]

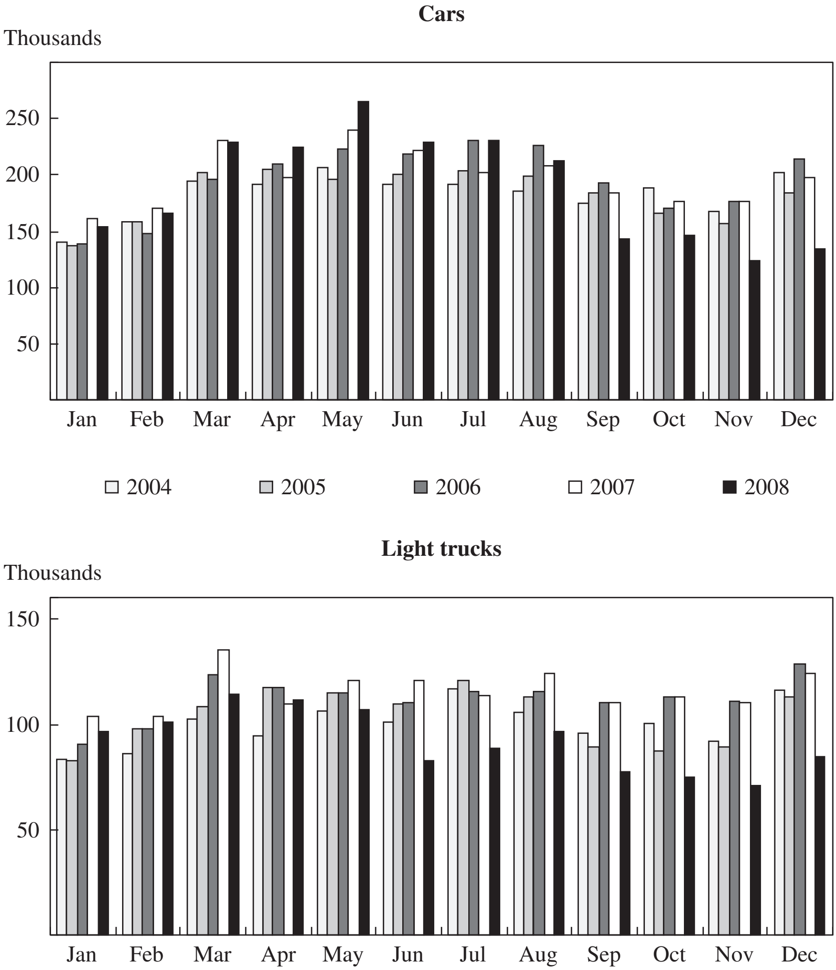

Source: WardsAuto.com, "Key Automotive Data" (wardsauto.com/keydata).
a. All vehicles produced outside of North America and sold in the United States.

in 1991Q1 than it had been in 1990Q3, similar to the $35 billion decline in this sector between 2007Q4 and 2008Q2.[14] Granted, that $35 billion in 2007–08 represents a smaller share of GDP than did the lost auto production in 1990–91, but it is still a sizable number, and without that contribution it would be hard to defend the claim that a recession began in 2007Q4.

14. Bureau of Economic Analysis, National Income and Product Accounts, table 1.5.6.

The last two rows of table 3 include details on this breakdown, looking ahead either four or five quarters beginning with 2007Q4. Average real GDP growth in the four quarters 2007Q4–2008Q3 was actually +0.75 percent at an annual rate. Had there been no decline in automobile output, that number would have been nearly half a percentage point higher. It would be very hard to characterize 2007Q4–2008Q3 as a full year of recession had average growth indeed been +1.2 percent. The Business Cycle Dating Committee of the National Bureau of Economic Research (NBER) reports that it was looking not just at GDP (which even with the decline in automobiles showed clearly positive growth), but also at gross domestic income (GDI), which differs from GDP by only a statistical discrepancy (see Nalewaik 2007). Real GDI growth averaged –0.4 percent over this period,[15] offering more justification for the NBER's recession call. But without the decline in the auto sector, this number, too, would have been positive, albeit very small.

The 2007–08 shock was also comparable to that of 1990–91 in terms of the effect on employment in the auto industry. Seasonally adjusted manufacturing employment in motor vehicles and parts fell by 94,000 between July 1990 and March 1991, compared with 132,000 between July 2007 and August 2008.[16] Again the latter is relative to a larger economy, but again it is not an inconsequential number. A year-over-year drop in total employment is viewed by some as a defining characteristic of a U.S. recession. This threshold was crossed in May 2008, when 87,000 fewer workers were employed than in May 2007.[17] Again without the contribution of automobiles, it would not be at all clear that the economy should have been characterized as being in recession during 2007Q4–2008Q2.

Of course, the first half of 2008 saw not just a big decline in auto purchases but also a slowdown in overall consumer spending and a big drop in consumer sentiment, again very much consistent with what was observed after earlier oil shocks. Like SUV sales, consumer sentiment spiked back up dramatically in an initial response to falling gasoline prices at the end of the summer, but, like SUV sales, it then plunged back down as broader economic malaise developed in the fall of 2008.

15. Following Nalewaik (2007), I calculated real GDI by dividing nominal GDI (Bureau of Economic Analysis, National Income and Product Accounts, table 1.10) by the implicit GDP deflator.

16. Bureau of Labor Statistics, "Employment, Hours, and Earnings from the Current Employment Statistics Survey (National)" (series CES3133600101).

17. This is based on the BLS seasonally unadjusted establishment survey data.

For some more formal statistical evidence and quantification, I turn to several of the studies described in the previous section. I first examine in table 4 how well the models proposed in previous studies perform against data that have become available since those papers were written. To evaluate the Edelstein-Kilian bivariate VARs, I used the parameter values for the relationships as estimated over July 1970 through July 2006 to form forecasts over the postsample interval August 2006 to September 2008. I then compared the postsample one-month-ahead mean squared errors (MSEs) with those that would have been obtained by a univariate autoregression (excluding energy prices) estimated over the same original sample (July 1970 through July 2006). As reported in the last column of table 4, for each of the six Edelstein-Kilian relationships used here, energy prices made a useful contribution to the postsample forecasts. This allows some confidence in using those estimates to measure the contribution that energy prices may have made to the economy's response to the oil shock of 2007–08.

I also used the Edelstein-Kilian relationships as estimated over July 1970 to July 2006 to form a 1- to 12-month-ahead forecast of how these variables might have behaved from September 2007 through September 2008 had there been no oil shock. The top panel of figure 17 presents the results for real PCE. In the absence of any new shocks, the Edelstein-Kilian bivariate VAR would have predicted continued growth in PCE at the rate it had grown over the previous half year. In the event, consumption grew much more slowly than predicted through May and then started to decline. The figure also shows the portion of the forecast error at any date attributable to the cumulative surprises in energy prices between September 2007 and the indicated date, calculated as described above. Energy prices can account for about half of the gap between predicted and actual consumption spending over this period. The second and third panels repeat the exercise for the big drops in spending on motor vehicles and in consumer sentiment. Most of the declines in these two series through the beginning of 2008 and about half the decline through the summer of 2008 would be attributed to energy prices, according to the Edelstein-Kilian regressions.

I also examined the postsample performance of the Blanchard-Galí VAR as estimated over their second subsample, 1984Q1–2007Q3. In this case I compared their six-variable VAR with a five-variable VAR that leaves out oil prices. Their model with oil prices in fact does somewhat worse at predicting GDP growth rates for data that became available after their study than would a similar VAR without oil prices (third row from the

Table 4. Improvements in Postsample Mean Squared Errors Achieved by Alternative Models

Study	*Sample period*	*Evaluation period*	*Dependent variable*	*Comparison model*	*Improvement (percent)*
Edelstein and Kilian (2007)	July 1970–July 2006	Aug. 2006–Sept. 2008	Real PCE	AR(6)	33
			PCE services	AR(6)	8
			PCE nondurables	AR(6)	23
			PCE durables	AR(6)	30
			PCE automobiles	AR(6)	26
			Consumer sentiment	AR(6)	9
Blanchard and Galí (2008)	1984Q1–2007Q3	2007Q4–2008Q4	Real GDP	VAR(5)	–11
Hamilton (2003)	1949Q2–2001Q3	2001Q4–2008Q4	Real GDP	AR(4)	0
	1949Q2–2001Q3	2007Q1–2008Q4	Real GDP	AR(4)	45

Source: Author’s calculations.

Figure 17. Contribution of Energy Prices and Other Factors to Consumption Spending and Consumer Sentiment, United States, 2007–08

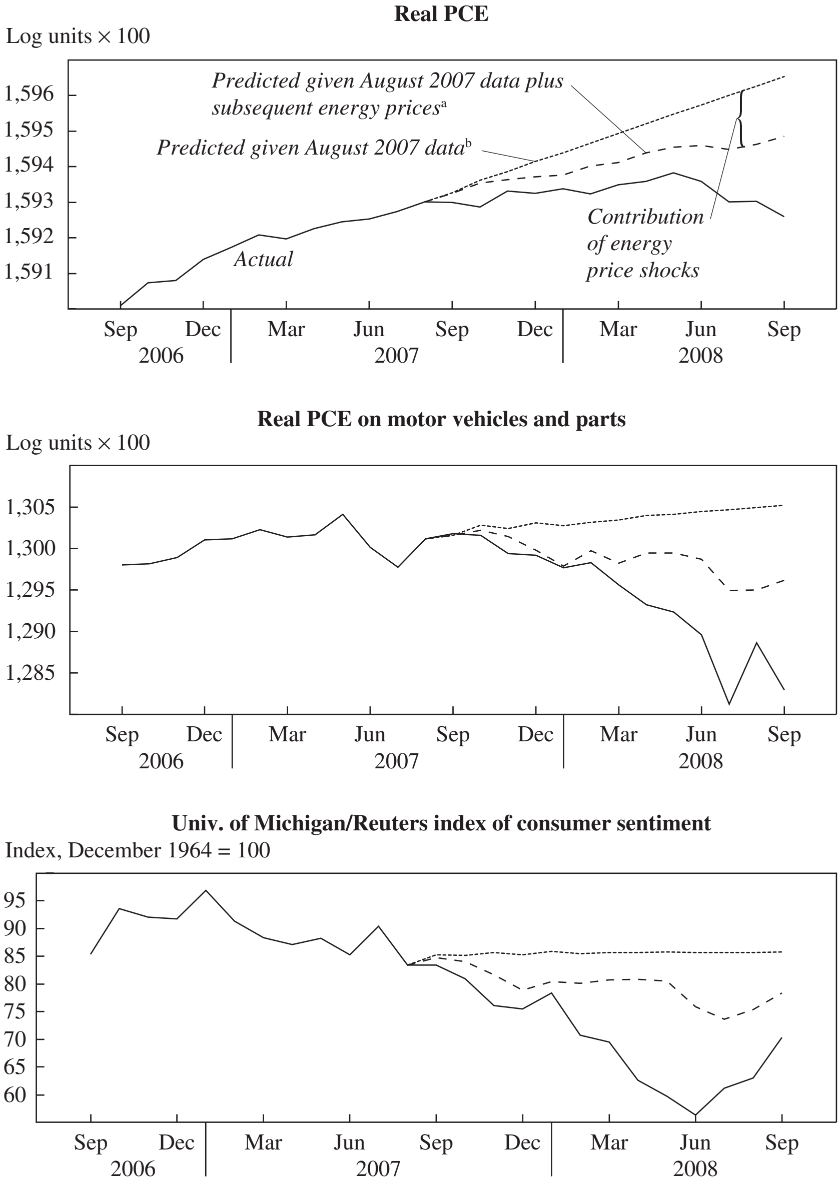

Source: Author's calculations.

a. Forecast made on the basis of information available as of August 2007 plus innovations in the energy price measure over September 2007 through September 2008.

b. Forecast made on the basis of information available as of August 2007.

bottom of table 4). I nevertheless examined how much of the downturn of 2007–08 their coefficients would attribute to oil prices, looking at the errors made in forecasting GDP growth over 2007Q4–2008Q3 or −Q4 on the basis of information available as of 2007Q3, and at the contribution of oil price surprises to these forecast errors. The result of this calculation (third column of table 3) suggests that real GDP growth would have been 0.7 percentage point higher on average in the absence of the oil shock. Thus, the Blanchard-Galí calculations also support the conclusion that the period 2007Q4–2008Q3 would not reasonably be considered the beginning of a recession had there been no contribution from the oil shock.

Finally, I looked at the postsample performance of the GDP forecasting regression (equation 8 above) estimated in Hamilton (2003). As seen in the next-to-last row of table 4, this relationship has about the same MSE over the period 2001Q4–2008Q4 as does a univariate autoregression fit to the 1949Q2–2001Q3 data. In part this lack of improvement is due to the fact that the oil-based relationship predicts slower GDP growth than was observed for 2005 and 2006, when the price of oil rose but the economy seemed little affected.

It is interesting to note that the historical relationship (equation 8) significantly outperforms a univariate specification when evaluated on the same postsample intervals used to evaluate the Edelstein-Kilian and Blanchard-Galí relationships in table 4. In that case equation 8 achieves a 45 percent improvement in terms of the postsample MSE over the period 2007Q1–2008Q4 compared with a univariate autoregression. Indeed, the relationship could account for the entire downturn of 2007–08 (see the last column of table 3). If one could have known in advance what would happen to oil prices during 2007–08, and if one had used the historically estimated relationship to form a one- to five-quarter-ahead forecast of real GDP, looking ahead to 2007Q4–2008Q4 from 2007Q3, one could have predicted real GDP for both 2008Q3 and 2008Q4 quite accurately (figure 18).

That last claim seems hard to believe, since Blanchard and Galí are doubtless correct that the effects of changes in oil prices have lessened somewhat as the economy has become less manufacturing based and more flexible, and given that the housing downturn surely made a critical contribution to the recession that began in 2007. Nevertheless, a few points about the separate contributions of housing and the oil shock deserve mentioning. First, housing had been exerting a significant drag on the economy before the oil shock, despite which economic growth continued. Residential fixed investment subtracted an average of 0.94 percentage point from

Figure 18. Dynamic Forecasts of Real U.S. GDP from Information Available in 2007Q3 with and without Oil Price Data Available in 2007Q4–2008Q4

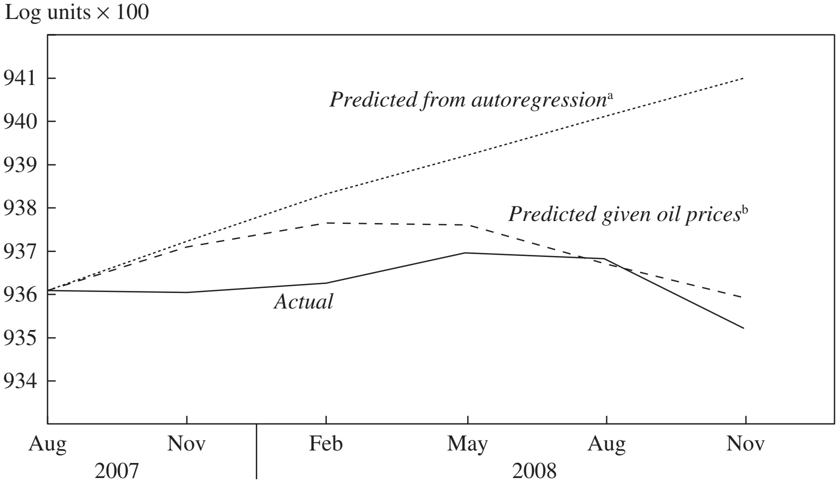

Source: Author's calculations.

a. Dynamic forecast one to five quarters ahead based on coefficients from a univariate AR(4) regression estimated on data from 1949Q2 to 2001Q3 and applied to GDP data through 2007Q3.

b. Dynamic conditional forecast one to five quarters ahead based on coefficients reported in equation 8 (which was estimated on data from 1949Q2 to 2001Q3) applied to GDP data through 2007Q3 and conditioning on the ex post realizations of the net oil price increase measure $o^{\#}_{t+s}$ for $t + s$ = 2007Q4–2008Q3.

annual real GDP growth over 2006Q4–2007Q3, when the economy was not in a recession, but only 0.89 percentage point over 2007Q4–2008Q3, which includes the beginning of the recession. At a minimum it is clear that something other than housing must have deteriorated to turn slow growth into a recession. That something, in my mind, includes the collapse in auto purchases, the slowdown in overall consumption spending, and deteriorating consumer sentiment, to which the oil shock indisputably contributed.

Second, there is an interaction effect between the oil shock and the problems in housing. Joe Cortright (2008) notes that in the Chicago, Los Angeles, Pittsburgh, Portland (Oregon), and Tampa metropolitan statistical areas, home prices in 2007 were likely to rise slightly in the zip codes closest to the central urban areas but to fall significantly in zip codes with longer average commuting distances. Foreclosure rates also rose with distance from the center. And certainly to the extent that the oil shock made a direct contribution to lower income and higher unemployment, that would also depress housing demand. For example, the estimates in Hamilton

(2008) imply that a 1-percentage-point reduction in real GDP growth translates into a 2.6 percent reduction in the demand for new homes.

Eventually, the declines in income and home prices raised mortgage delinquency rates beyond a threshold at which the overall solvency of the financial system itself came into question, and the modest recession of 2007Q4–2008Q3 turned into a ferocious downturn in 2008Q4. Whether those events would have been avoided had the economy not gone into recession, or instead merely postponed, is a matter of conjecture. Regardless of how one answers that question, the evidence to me is persuasive that had there been no oil shock, economists today would be describing the economy in 2007Q4–2008Q3 as growing slowly but not in a recession.

A final question worth examining is why the oil price increases before 2007Q4 failed to have a bigger effect on the economy. Why did consumers respond so little when the price of oil went from $41 a barrel in July 2004 to $65 in August 2005 (a 59 percent increase), yet so strongly to the 86 percent increase from $72 a barrel in August 2007 to $134 in June 2008?[18] Equations posed in terms of percentage changes, such as equation 8, would predict that the 2004–05 price increases should also have had significant effects on output. Yet the dollar impact on household budgets of the $62-a-barrel price increase in 2007–08 is considerably more than twice that of the $24-a-barrel increase in 2004–05.

To explore this possibility more concretely, I looked at the consequences of modifying equation 8 to take into account the changes in the energy budget share over time, replacing $o_t^{\#}$ with the product $o_t^{\#}\alpha_{t-1}$ for α_t, the energy share plotted in figure 3.[19] This results in a slight improvement in fit for the original sample period (t = 1949Q2–2001Q3), raising the log likelihood from −281.78 for the original specification to −281.47 for the new. The share-weighted regression performs significantly better post-sample, producing a 10.8 percent improvement in the MSE over the period 2001Q4–2008Q4 relative to an autoregression with no role for oil prices. For the specific period 2005Q1–2006Q4, the modified specification as estimated over 1949Q2 to 2001Q3 would have predicted an average annual real GDP growth rate of 1.9 percent (calculated logarithmically), a bit below the sluggish 2.5 percent actually observed.

Oil prices thus appear to have exerted a moderate drag on real GDP in 2005–06 and made a more significant negative contribution in 2007–08.

18. Oil prices quoted here are monthly averages of daily West Texas Intermediate prices.

19. The monthly series was converted to quarterly data by using only the third month of each quarter. Values for α_t for quarters before 1959Q1 were simply set equal to the January 1959 value (7.354).

The principal reason that Americans ignored the earlier price increases would seem to be that they could afford to do so. By 2007Q4 they no longer could, and the sharp spike in oil prices led to an observed economic response similar to that seen in earlier episodes.

VI. Policy Implications

I have raised the possibility that miscalculation of the long-run price elasticity of oil demand by market participants was one factor in the oil shock of 2007–08, and that speculative investing in oil futures contracts may have contributed to that miscalculation. Were any policies available to mitigate the problems that this produced for the economy? One option would have been for the federal government to sell some oil directly out of the Strategic Petroleum Reserve (SPR) in the spring of 2008, perhaps timing the sales to coincide with expiry dates on NYMEX crude oil contracts. If there was speculative momentum buying, such a step might have succeeded in reversing it. At worst, the government would have made a profit on its SPR investment by buying low and selling high.

A more conventional policy tool would have been monetary policy. A number of observers (for example, Frankel 2008) have suggested that the very rapid decline in short-term interest rates in 2008Q1 fanned the flames of commodity speculation, with negative real interest rates encouraging investments in physical commodities. In January 2009 Federal Reserve chairman Ben Bernanke offered the following retrospective on that debate:

> The [Federal Open Market] Committee's aggressive monetary easing was not without risks. During the early phase of rate reductions, some observers expressed concern that these policy actions would stoke inflation. These concerns intensified as inflation reached high levels in mid-2008, mostly reflecting a surge in the prices of oil and other commodities. The Committee takes its responsibility to ensure price stability extremely seriously, and throughout this period it remained closely attuned to developments in inflation and inflation expectations. However, the Committee also maintained the view that the rapid rise in commodity prices in 2008 primarily reflected sharply increased demand for raw materials in emerging market economies, in combination with constraints on the supply of these materials, rather than general inflationary pressures. Committee members expected that, at some point, global economic growth would moderate, resulting in slower increases in the demand for commodities and a leveling out in their prices—as reflected, for example, in the pattern of futures market prices. As you know, commodity prices peaked during the summer and, rather than leveling out, have actually fallen dramatically with the weakening in global economic activity. As a consequence, overall inflation has already declined significantly and appears likely to moderate further. (Bernanke 2009)

Bernanke seems here to be taking the position that since the Federal Reserve had been correct about the long run—that ultimately there would be a significant downturn in both the economy and commodity prices, creating strong disinflationary pressure—the short-run consequences (booming commodity prices in the first half of 2008) were less relevant. On the other hand, if indeed the spike in oil prices was one causal factor contributing to the downturn itself, the Federal Reserve can hardly afford to ignore those short-run implications. The evidence examined here suggests that the Federal Reserve needs to give careful consideration to the possible consequences of its actions for relative commodity prices.

But although the question of the possible contributions of speculators and the Federal Reserve is an interesting one, it should not distract from the broader truth: some degree of significant oil price appreciation during 2007–08 was an inevitable consequence of booming demand and stagnant production. It is worth emphasizing that this is fundamentally a long-run problem, which has been resolved rather spectacularly for the time being by a collapse in the world economy. However, one may hope that this collapse will prove to be a short-run cure for the problem of excess energy demand. If growth in the newly industrialized countries resumes at its former pace, it will not be too many more years before we find ourselves back in the kind of calculus that was the driving factor behind the problem in the first place. Policymakers would be wise to focus on real options for addressing those long-run challenges, rather than blame what happened in 2008 entirely on a market aberration.

ACKNOWLEDGMENTS I am grateful to Alan Blinder, Lutz Kilian, participants at the Brookings Papers conference, and an anonymous referee for helpful comments on an earlier draft of this paper, and to Davide Debortoli for supplying the Blanchard-Galí data and code.

References

Al-Sowayegh, Abdulaziz H. 1984. *Arab Petropolitics.* London: Croom Helm.

Barsky, Robert B., and Lutz Kilian. 2002. "Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative." *NBER Macroeconomics Annual 2001,* vol. 16, pp. 137–83.

Bernanke, Ben S. 2009. "The Crisis and the Policy Response." Stamp Lecture, London School of Economics, January 13 (www.federalreserve.gov/newsevents/speech/bernanke20090113a.htm).

Blanchard, Olivier J., and Jordi Galí. 2008. "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?" In *International Dimensions of Monetary Policy,* edited by Jordi Galí and Mark J. Gertler. University of Chicago Press.

Blinder, Alan S., and Jeremy B. Rudd. 2008. "The Supply-Shock Explanation of the Great Stagflation Revisited." Working Paper 14563. Cambridge, Mass.: National Bureau of Economic Research.

Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld. 2008. "A Meta-Analysis of the Price Elasticity of Gasoline Demand: A SUR Approach." *Energy Economics* 30, no. 5: 2105–22.

Cooper, John C. B. 2003. "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries." *OPEC Review* 27, no. 1: 1–8.

Cortright, Joe. 2008. "Driven to the Brink: How the Gas Price Spike Popped the Housing Bubble and Devalued the Suburbs." Discussion paper. Chicago: CEOs for Cities. www.ceosforcities.org/files/Driven%20to%20the%20Brink%20FINAL.pdf.

Dahl, Carol A. 1993. "A Survey of Oil Demand Elasticities for Developing Countries." *OPEC Review* 17, no. 4(Winter): 399–419.

Dahl, Carol, and Thomas Sterner. 1991. "Analysing Gasoline Demand Elasticities: A Survey." *Energy Economics* 13, no. 3: 203–10.

Edelstein, Paul, and Lutz Kilian. 2007. "Retail Energy Prices and Consumer Expenditures." Working paper. University of Michigan.

Energy Information Administration. 2006. "A Primer on Gasoline Prices." Washington (www.eia.doe.gov/pub/oil_gas/petroleum/analysis_publications/primer_on_gasoline_prices/html/petbro.html).

Espey, Molly. 1998. "Gasoline Demand Revisited: An International Meta-Analysis of Elasticities." *Energy Economics* 20, no. 3: 273–95.

Frankel, Jeffrey A. 2008. "The Effect of Monetary Policy on Real Commodity Prices." In *Asset Prices and Monetary Policy,* edited by John Y. Campbell. University of Chicago Press.

Gately, Dermot, and Hillard G. Huntington. 2002. "The Asymmetric Effects of Changes in Price and Income on Energy and Oil Demand." *Energy Journal* 23, no. 1: 19–55.

Graham, Daniel J., and Stephen Glaister. 2004. "Road Traffic Demand Elasticity Estimates: A Review." *Transport Reviews* 24, no. 3: 261–74.

Hamilton, James D. 1983. “Oil and the Macroeconomy since World War II.” *Journal of Political Economy* 91, no. 2: 228–48.

———. 1988. “A Neoclassical Model of Unemployment and the Business Cycle.” *Journal of Political Economy* 96, no. 3: 593–617.

———. 1994. *Time Series Analysis.* Princeton University Press.

———. 2003. “What Is an Oil Shock?” *Journal of Econometrics* 113, no. 2: 363–98.

———. 2008. “Daily Monetary Policy Shocks and New Home Sales.” *Journal of Monetary Economics* 55, no. 7: 1171–90.

———. 2009. “Understanding Crude Oil Prices.” *Energy Journal* 30, no. 2: 179–206.

Hotelling, Harold. 1931. “The Economics of Exhaustible Resources.” *Journal of Political Economy* 39, no. 2: 137–75.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand.” *Energy Journal* 29, no. 1: 113–34.

Jovanovic, Boyan. 2007. “Bubbles in Prices of Exhaustible Resources.” Working Paper 13320. Cambridge, Mass.: National Bureau of Economic Research.

Kilian, Lutz. 2008. “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” *Review of Economics and Statistics* 90, no. 2: 216–40.

———. 2009. “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market.” *American Economic Review* 99, no. 3: 1053–69.

Leamer, Edward E. 2008. “What’s a Recession, Anyway?” Working Paper 14221. Cambridge, Mass.: National Bureau of Economic Research.

Masters, Michael W. 2008. “Testimony of Michael W. Masters, Managing Member / Portfolio Manager, Masters Capital Management, LLC.” Testimony before the U.S. Senate Committee on Homeland Security and Governmental Affairs (hsgac.senate.gov/public/_files/052008Masters.pdf).

Nalewaik, Jeremy J. 2007. “Estimating Probabilities of Recession in Real Time Using GDP and GDI.” Finance and Economics Discussion Series Working Paper 2007–07. Board of Governors of the Federal Reserve System, Division of Research and Statistics and Monetary Affairs.

Comments and Discussion

COMMENT BY

ALAN S. BLINDER James Hamilton's fine paper poses a question that seems especially appropriate, having been presented at this conference just a few days before Passover: Why was this oil price shock different from all other oil price shocks (if indeed it was)?

For readers without a working knowledge of the Haggadah, the famous Four Questions posed at each seder are what we would now call (although the ancient rabbis did not) "bullet points" under one umbrella question: "Why is this night different from all other nights?" Analogously, I would phrase Hamilton's umbrella question as, "Why was the 2002–08 oil shock so different from OPEC I and OPEC II?" Notice that I make the start date 2002, not 2007 as Hamilton does, for that is when the long escalation of the oil price began. I will discuss five subquestions under this umbrella question:

—Was it because this shock was demand-driven rather than supply-driven?

—Was it because this shock was gradual rather than sudden?

—Was it because the U.S. economy is less energy-intensive now than it was then?

—Was it because this shock was not accompanied by a large food price shock? or (most conjecturally)

—Was it because the U.S. labor market has more wage flexibility now than it did then?

As I proceed through this list, it will be apparent that I agree with Hamilton much more than I disagree with him.

DEMAND- VERSUS SUPPLY-DRIVEN OIL SHOCKS. Many observers have emphasized the following apparently stark difference between the oil shocks of the 1970s and the oil shock of this decade:[1] OPEC I and II were initiated by supply interruptions, in each case related to war in the Middle East, whereas the oil shock of 2002–08 was driven mainly by rapid growth of world demand, especially from China. This sounds like a fundamental difference. After all, demand shocks and supply shocks induce the opposite covariance between price and quantity. But Hamilton's useful narrative prompts a reconsideration of this point. As he tells the story, OPEC I and II were supply disruptions greatly magnified by hoarding (which is a demand surge), whereas the 2002–08 run-up resulted from surging demand colliding with stagnant supply. Phrased that way, the two do not sound all that different. Besides, a demand-induced increase in world oil prices is a supply shock to an oil-importing country. So, if this oil shock really was different from the others, we need to look elsewhere for the reasons.

Before doing so, one other point about the supply-versus-demand distinction should be made, largely for future reference. Many previous scholars have noted that the apparent macroeconomic impacts of oil shocks are far larger than the predicted neoclassical supply-side effects, which are based on factor substitution along a factor price frontier. The clear implication is that these oil shocks must also have had powerful Keynesian demand-side effects, for example through the "oil tax" draining consumer purchasing power. As Kilian (2008) points out, when strong world GDP growth pushes up the price of oil, oil-importing countries like the United States should also experience rapid export growth, which should offset part of the demand-reducing effects of the oil shock. The table below, which comes straight from the National Income and Product Accounts (table 1.1.2), shows that this did indeed happen to the United States in the present decade:

	Contribution of net exports to real GDP growth (percentage points)
2004	–0.7
2005	–0.2
2006	0
2007	+0.6
2008	+1.4.

GRADUAL VERSUS SUDDEN. I believe economists started using the term supply *shock* back in late 1973 because of the extreme suddenness of the oil price rise. It is well known, and repeated by Hamilton, that price elasticities

1. I will refer to OPEC I and II as the "oil shocks of the 1970s" for convenience even though OPEC II continued into 1980.

on both the demand and the supply sides of the oil market are much larger in the long run than in the short run, presumably because substitution (whether in consumption or production) is far more difficult in the short run. It follows that a sudden price shock, such as in 1973–74, is more disruptive than a gradual one, such as in 2002–08.

In a Brookings Paper written while the 2002–08 shock was still in progress, William Nordhaus (2007, p. 227) calculated the size of the "oil tax" in that shock and three previous ones on an *annualized* flow basis:

	"Oil tax" as a percent of GDP (annualized)
1973Q3–1975Q4	1.9
1978Q4–1981Q1	2.1
1990Q3–1990Q4	3.0
2002Q4–2006Q2	0.7

Looked at in this way, the most recent oil shock was only about one-third as severe as OPEC I and OPEC II because it developed so gradually. (The 1990 oil shock was the sharpest, but it lasted only one quarter.) That is, indeed, a big difference. No other long, gradual oil price shock like this can be found in the data.

DECLINING ENERGY INTENSITY. It is well known that the U.S. economy (like others) uses far less energy per dollar of GDP now than it did before OPEC I. Figure 1, taken from my recent paper with Jeremy Rudd (2008), is one concrete depiction of that fact. The sharp decline in U.S. energy intensity (measured as thousands of BTUs consumed per dollar of real GDP) since the early 1970s is evident, as is the slowdown in that downtrend that began when oil became cheap again in 1985–86. Energy intensity by 2007 was roughly half what it had been in 1973. On this basis alone, the 2002–08 oil shock should have packed only half the punch of the 1970s oil shocks. This may be the only issue on which I disagree with Hamilton, who seems skeptical that "delayed consequences of increased energy conservation following the 1970s oil shocks" played a major role in reducing the demand for oil. The figure, it seems to me, highlights that this was an important difference between now and then.

THE FORGOTTEN FOOD PRICE SHOCKS. I have noticed over the years that many economists have dismayingly selective memories when it comes to things that interfere with their favored story lines. In fact, there was another sort of supply shock in the 1970s. Large and durable food price shocks accompanied both OPEC I and OPEC II, but have since been mostly forgotten. Figure 2, which again comes from Blinder and Rudd (2008), shows the rate of change of the food component of the consumer

Figure 1. Energy Intensity of the Economy, 1950–2007

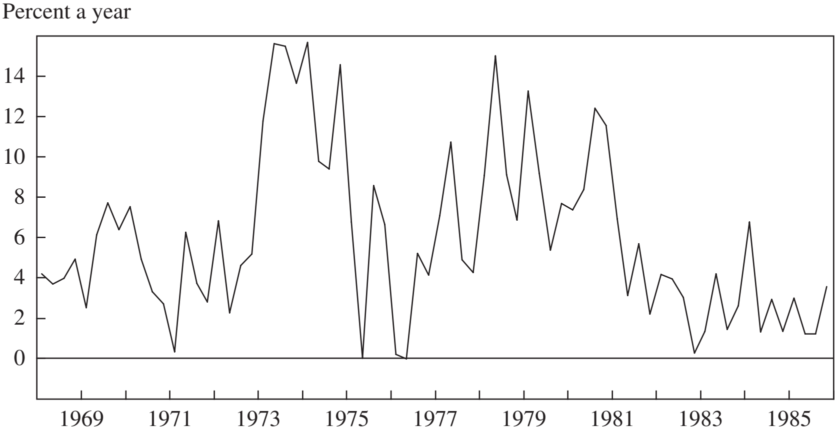

Source: Blinder and Rudd (2008).

a. Annualized quarterly log differences of the food and beverages component of the personal consumption expenditure deflator.

price index (CPI) from 1968 through 1985. It is clear, as I emphasized in Blinder (1979, 1982), that major food price shocks—featuring sustained double-digit rates of increase—buffeted the economy in 1972–74 and again in 1978–80. The price run-ups were smaller than those for oil, but the weight of food in the CPI is much larger.

Figure 2. Consumer Food Price Inflation, 1968–85[a]

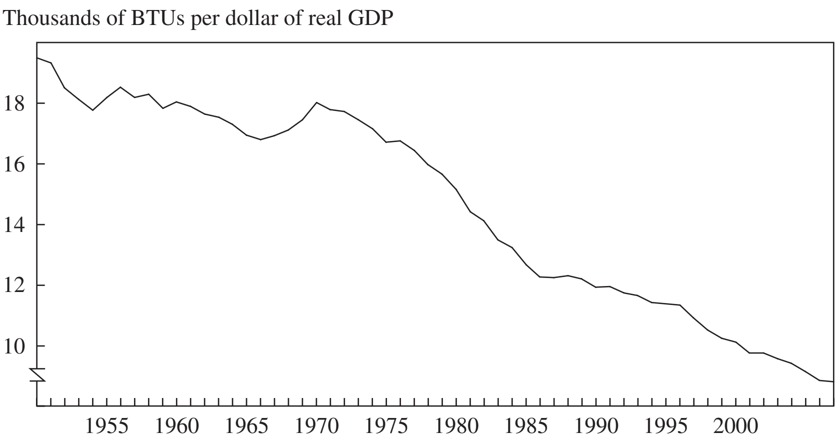

Source: Energy Information Administration, *Annual Energy Review 2007,* Table 1.5.

Much popular discussion at the time, and much econometric work since, failed to disentangle the inflationary effects of the food shocks from those of the oil shocks. In retrospect, I wonder whether economists attributed too much of the economy's travails in 1973–74 and 1978–80 to the oil shocks, and not enough to the food shocks.[2] This is a very important difference between the 1970s and the recent episode, because the recent run-up in oil prices was *not* accompanied by major food price inflation. In 2007 and 2008, the food component of the CPI rose 5 and 6 percent, respectively, on a December-to-December basis.

MORE WAGE FLEXIBILITY TODAY? I come now to my final, and most speculative, possible reason why the most recent oil shock may have been different from all previous ones. Is there greater (real) wage flexibility today than in the 1970s? With flexible real wages, the economy will react to an oil price hike in a more neoclassical and less Keynesian fashion. And, as noted earlier, the neoclassical supply-side effects appear to be smaller than the Keynesian demand-side effects. So if wages have become more flexible, this should have muted the impact of more recent oil shocks on both output and inflation. Both Nordhaus (2007) and Olivier Blanchard and Jordi Galí (2007) offer shreds of evidence to suggest that this is so.

DATING THE RECESSION. I conclude with one last, interesting question raised by Hamilton, even though it is not part of the traditional seder inquiries. Hamilton asks whether the four quarters 2007Q4–2008Q3 would have been considered recession quarters—as the NBER Business Cycle Dating Committee does—in the absence of the depressing effects of the oil shock. The overwhelming weight of the evidence that he adduces points to the answer being no, and I agree wholeheartedly.

In fact, I would like to push the question one step further. Should the three quarters 2007Q4–2008Q2 be considered recession quarters *at all?* Real GDP growth rates for these three quarters were (pending further revisions) –0.2, +0.9, and +2.8 percent, respectively. The average is +1.2 percent, which is certainly slow growth but does not look like a recession. By contrast, the quarterly growth rates for the subsequent three quarters (2008Q3–2009Q1) were –0.5, –6.3, and –5.7 percent, respectively, averaging –4.2 percent. Now, that's a recession. Furthermore, as is well known, virtually every macroeconomic time series took a sharp turn for the worse immediately following the failure of Lehman Brothers on September 15,

2. According to Blinder and Rudd's (2008, table 4) estimates, the food shock had a larger effect on inflation than the oil shock did in 1973–74, and the two shocks had roughly equal impacts on inflation in 1978–80.

2008. With full knowledge that several members of the NBER dating committee are participants at this conference, I will close by suggesting that the recession began on that fateful day.

REFERENCES FOR THE BLINDER COMMENT

Blanchard, Olivier J., and Jordi Galí. 2007. "The Macroeconomic Effects of Oil Shocks: Why Are the 2000s So Different from the 1970s?" Working Paper 13368. Cambridge, Mass.: National Bureau of Economic Research (September).

Blinder, Alan S. 1979. *Economic Policy and the Great Stagflation.* New York: Academic Press.

———. 1982. "The Anatomy of Double-Digit Inflation in the 1970s." In *Inflation: Causes and Effects,* edited by Robert E. Hall. University of Chicago Press.

Blinder, Alan S., and Jeremy B. Rudd. 2008. "The Supply-Shock Explanation of the Great Stagflation Revisited." Working Paper 14563. Cambridge, Mass.: National Bureau of Economic Research (December).

Kilian, Lutz. 2008. "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature* 46, no. 4: 871–909.

Nordhaus, William D. 2007. "Who's Afraid of a Big Bad Oil Shock?" *BPEA,* no. 2: 219–38.

COMMENT BY

LUTZ KILIAN James Hamilton has provided an insightful analysis of the latest oil price shock. He makes the case for viewing this episode not merely as a market aberration, but as a systemic and long-term problem that is likely to resurface once the global economy recovers from the current recession. After reviewing time-series plots of the relevant price and quantity data and conventional estimates of the price elasticity of oil demand, he concludes that constraints on the production of crude oil after 2005 and growing demand for crude oil driven by the recent boom in the world economy are the primary explanation of the 2007–08 oil price shock.

Having outlined an explanation based purely on economic fundamentals, Hamilton also gives serious consideration to the view that speculation in oil markets may have worsened the oil price spike of 2007–08. One possible view is that speculators in oil-importing countries, anticipating future oil shortages, caused the spot price to increase; however, the fact that oil inventories did not increase substantially in recent years cautions against that interpretation. Another view is that oil-producing countries were misled by rising oil futures prices into reducing current production. Although the analysis in the paper shows that data on the oil futures spread do not support that view, it is conceivable that oil producers nevertheless withheld oil supplies in anticipation of even higher oil prices. That conjecture is hard to prove or disprove. Either way, the recent surge in the price of oil was ultimately driven by excess demand for crude oil.

The paper highlights the roles of consumer sentiment, of the automobile sector, and of the housing sector in the transmission of the 2007–08 oil price shock. Evidence is presented that this shock was a major factor in causing the current recession and that its impact was magnified by the rising energy share in expenditure. The paper also includes a discussion of two policy tools that might have been used to slow the surge in the price of oil: one is the release of strategic oil reserves to calm speculators; the other is a slower easing of interest rates in 2008. The paper concludes that it would be wise for policymakers to address the long-run policy challenges of booming oil demand and stagnant global oil production, as the recent economic collapse is likely to prove only a short-run cure for the problem of excess demand.

I agree with many of the points in this paper, and in only a few instances would I have favored a more subtle interpretation. At the risk of downplaying the many areas of agreement, in this comment I will focus on two main themes that strike me as especially worthy of discussion.

One of these is that whereas earlier oil price shocks were primarily caused by exogenous physical disruptions of supply, the price run-up of 2007–08 was caused by strong demand confronting stagnating world production. Although I agree with this analysis of 2007–08 and with the proposition that this latest episode has been different from earlier ones, a growing body of evidence argues against the notion that the earlier oil price shocks were driven primarily by unexpected disruptions of the global supply of crude oil.

The paper acknowledges that demand pressures arising from increased global real activity made some contribution to oil price increases during several earlier episodes, but this alternative explanation is never fully investigated. For example, although the paper briefly mentions (and dismisses as implausible) the effect of inflation (and of the devaluation of the dollar) on OPEC supply decisions in 1973, it completely abstracts from shifts in the demand for oil associated with fluctuations in global real activity, except for the analysis of 2007–08. As recent experience has demonstrated, however, such shifts have the potential to cause large fluctuations in the real price of oil. Hence, even if, for the sake of argument, all of the observed oil supply cutbacks in late 1973 or in 1979–80 were exogenous, it would not be self-evident that these supply disruptions, rather than fluctuations in the global business cycle, were the driving force behind the 1973–74 and 1979 increases in the real price of oil.

Figure 1. Global Real Activity and the Real Price of Oil, 1973–2008

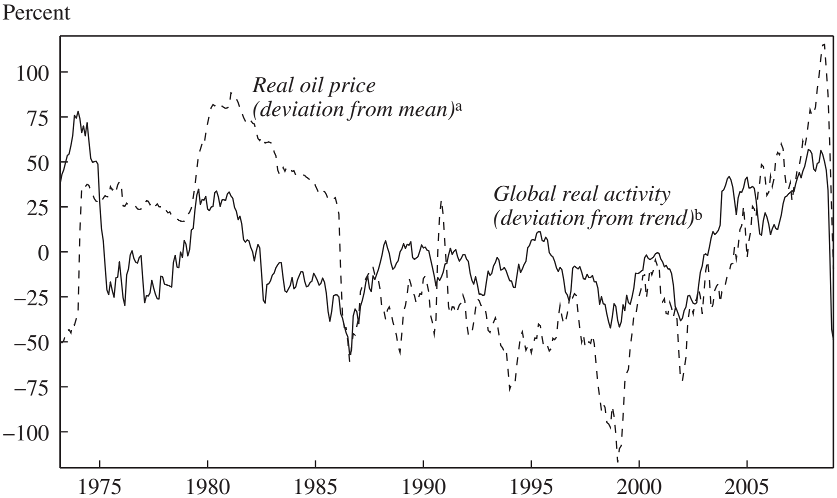

Source: Authorís calculations.
a. Price is deflated by the U.S. consumer price index.
b. Update of the measure described in Kilian (2009).

Moreover, there is good reason to be skeptical of the assertion that oil supply shocks were the primary explanation of all oil price shocks before 2007–08. Not only does Hamilton's figure 5 show considerable variation in the time-series patterns across oil price shock episodes, arguing against a common explanation, but no mention is made of the crucial point that commonly used measures of exogenous oil supply disruptions explain at most about 20 percent of the observed increase in the real price of oil in 1973–74. Alternative measures that I have proposed (Kilian 2008) imply even lower estimates of the predictive power of exogenous oil supply shocks. This raises the question of what explains the remaining 80 percent of the observed oil price increase.

By construction, the answer to this question must have to do with shifts in the demand for oil. Arguably the most important driving force behind the demand for oil is global real activity. Figure 1 plots the real price of crude oil and a measure of global real economic activity for 1973 through the end of 2008. As expected, not all movements in the real price of oil were associated with swings in global real activity, but the three major oil price shock episodes of 1973–74, 1979–80, and 2002–08 all coincided with major surges in global real activity. The attentive reader will notice that the increase in real activity in 1973 predated the increase in the real

Table 1. Growth in Inflation-Adjusted Prices of Selected Commodities[a]
Percent

Commodity	*November 1971– February 1974*	*August 1977– February 1980*	*June 2001– June 2008*
Crude oil	125.3	70.7	331.5
Industrial raw materials	92.6	24.2	67.0
Metals	95.9	27.6	235.1

Source: Author's calculations using data from the Commodity Research Bureau and Kilian (2009).
a. Cumulative changes over the indicated period relative to the U.S. consumer price index.

price of oil; in fact, it started in late 1971. The reason for this asynchronicity, as discussed in Robert Barsky and Kilian (2002) and Kilian (2008), is that the price of crude oil before late 1973 was not determined by market forces and remained below its market-clearing level. Had the price of oil been free to move, it would have risen much earlier, in line with other industrial commodity prices. The second major upswing in oil prices coincided with a somewhat smaller surge in global real activity starting in 1978. Finally, it is also evident that the latest oil price shock started in 2003 rather than 2007 and once again coincided with a very large swing in global real activity.

How much of an increase in the real price of industrial commodities such global demand swings may cause depends on how elastically the commodities in question can be supplied. It is instructive to contrast the increases in selected aggregate commodity price indices during the three episodes of interest. Table 1 shows that between late 1971 and early 1974, both industrial raw materials and metals prices increased by about 95 percent in real terms, despite a secular downward trend in these prices. Since contemporary sources indicate no important supply shocks in these markets at the time, and since most of the increases predate the oil price increase in late 1973, it is reasonably certain that all of these increases were driven by shifts in global demand (also see National Commission on Supplies and Shortages 1976). The observed increase in the real price of oil is only moderately higher, suggesting that stronger global demand is the explanation of the extra 80 percent increase in the real price of oil. Thus, 1973–74 appears much more similar to the current episode than this paper would have us believe.

Likewise, for 1979–80, table 1 suggests that demand pressures seem capable of explaining perhaps a 30 percent increase in the real price of oil. A leading candidate to explain the remainder is rising concern in 1979 about future oil supply shortfalls, since once again exogenous oil supply shocks fail to explain the timing and magnitude of these oil price increases. For the period between about 2002 and mid-2008, there is evidence that sustained demand pressures were associated with even larger real commodity price increases than in the 1970s. The reason that the real price of crude oil rose even faster than other industrial commodity prices in this episode—and here I fully agree with the points made in the paper—is that the supply of crude oil, having risen substantially between 2002 and 2005 in response to higher prices, stagnated after 2005.

Casual inspection of the data is a good starting point, but more formal regression analysis is required to identify unanticipated movements in global demand and oil supply and to account for their delayed effects on the real price of oil. Kilian (2009) shows that one can incorporate both global oil supply and global real economic activity into a regression framework that allows one to quantify the ability of unexpected physical shortfalls of oil production ("oil supply shocks") and of demand shocks driven by the global business cycle ("aggregate demand shocks") to explain the real price of oil. That model also includes a third shock, which may be viewed alternatively as an oil market–specific demand shock, reflecting, for example, shifts in uncertainty about future oil supply shortfalls, or as a measure of the difference between market expectations and econometric expectations of future oil supplies and global real activity. For the purpose of this discussion, I will focus mainly on the first two shocks to maintain consistency with Hamilton's analysis.

Figure 2 illustrates that the increase in the real price of oil from 2002 until mid-2008 was driven by a series of positive aggregate demand shocks associated with shifts in global economic activity. Oil supply shocks played no role. This analysis is very much consistent with this paper's interpretation of this episode, but it highlights again that this oil price shock really started in 2003 rather than 2007. Figure 2 also shows that the drop in the real price of oil after mid-2008 reflected only in part an unexpected reduction in global real activity. Other factors, presumably associated with the worsening financial crisis, also played some role, as shown in the bottom right panel.

Figure 2 also shows that the 1979–80 oil price shock actually reflected a composite of oil supply shocks (in 1980 rather than 1979), global aggregate demand shocks affecting all industrial commodity markets (starting in 1978), and other shocks, especially in 1979 (see Kilian 2009). This evidence is at odds with the view that all earlier oil shock episodes were driven primarily by oil supply disruptions. Likewise it has been shown that there is no evidence that the 1990–91 oil price shock was driven primarily

Figure 2. Historical Decomposition of Fluctuations in the Real Price of Oil, 1978–84 and 2002–08[a]

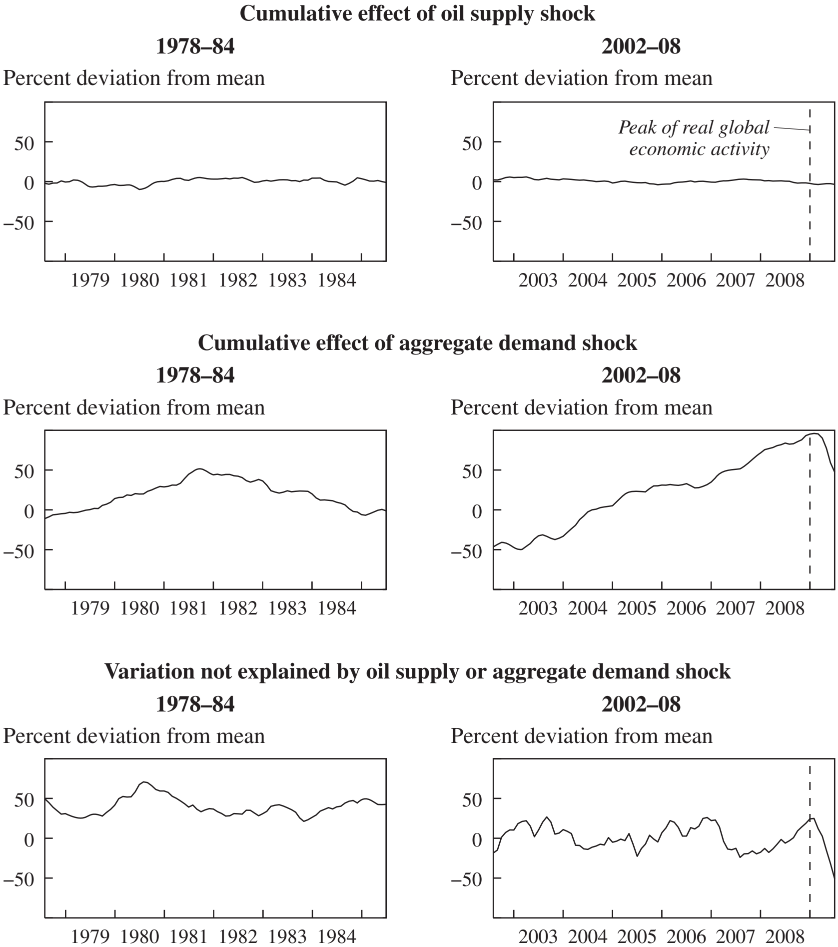

Source: Author's estimates.

a. Estimates from the structural VAR model of Kilian (2009) based on data for February 1973–December 2008.

by oil supply shocks, nor did the Venezuelan crisis of late 2002 and the Iraq War of 2003, which jointly triggered an oil supply disruption not unlike those of the 1970s, have much of an effect on the real price of oil (see Kilian 2008, 2009). Thus, none of the major oil price shocks since the 1970s appear primarily supply-driven. What has changed relative to earlier episodes is the composition of demand and supply shocks, with repeated positive global aggregate demand shocks alone explaining most of the run-up in oil prices since 2003. One may question how market participants could have been surprised again and again over the course of several years by strong global real economic activity. Hicks and Kilian (2009) provide evidence from data on professional GDP forecast revisions that this was indeed the case and that the most persistent forecast errors were associated with unexpectedly rapid growth in Asia. Moreover, estimated responses to such forecast errors show a pattern similar to the estimated response to aggregate demand shocks.

The other main theme of the paper is that the effect of the latest oil price shock on the U.S. economy has been quite similar to that of earlier ones. This argument is based on recursively identified vector autoregressions in which the oil price is ordered prior to the macroeconomic aggregate of interest. The global oil market model of Kilian (2009) takes the analysis a step further and expresses the VAR oil price innovation as a linear combination of oil demand and oil supply shocks, each of which is predetermined with respect to U.S. macroeconomic aggregates. This highlights two implicit assumptions that Hamilton makes in assessing the effects of oil price shocks. One assumption is that oil price innovations are homogeneous over time. This assumption would be innocuous if all oil price shocks were driven by exogenous oil supply disruptions, but, as has already been shown, oil price innovations reflect both oil demand shocks and oil supply shocks, the composition of which differs from one episode to the next, violating that assumption. The other assumption is that an oil price innovation is not associated with contemporaneous movement in any other macroeconomic variable. This presumption is violated if the oil price innovation is driven by global aggregate demand shocks. In that case, not only will the oil price innovation be correlated with innovations to the price of other industrial commodities, but the demand shock will also have a direct effect on the U.S. economy, for example, through the trade and external finance channel (see Kilian, Rebucci, and Spatafora 2009).

This does not mean that one cannot estimate the responses associated with an oil price innovation. Indeed, I have done so in my own work. One does, however, have to be clear that these responses do not represent the *causal* effect of an innovation to the price of oil, because the ceteris paribus condition is violated. Moreover, one has to keep in mind that these estimates represent the response to a shock of average composition over the sample period. They may be misleading when the composition of the oil price shock in question is atypical by historical standards, as is the case for the 2003–08 episode. Figure 3 illustrates this point. Since the latest oil

Figure 3. Responses of U.S. Real GDP to Oil Market Shocks[a]

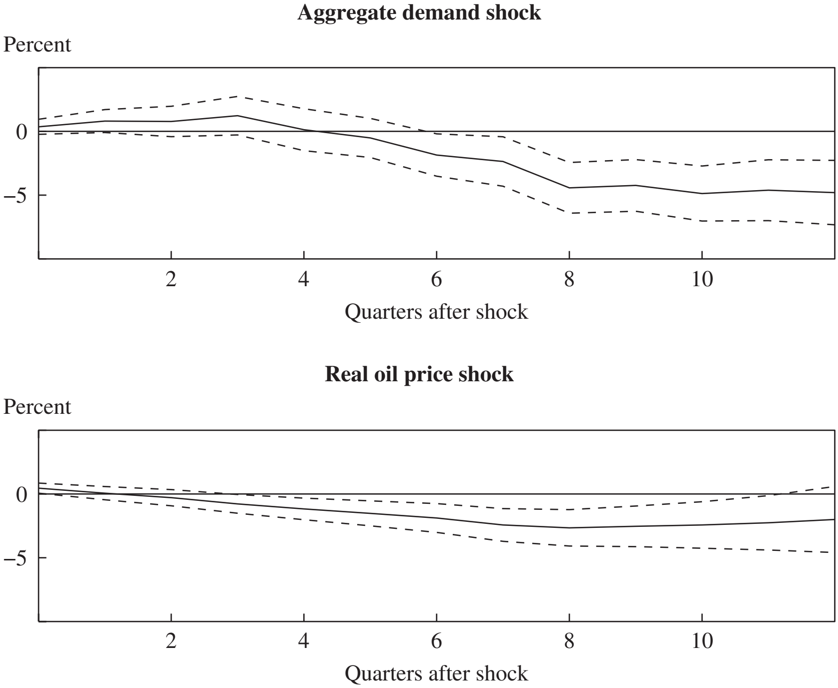

Source: Author's estimates.

a. The top panel is based on the methodology of Kilian (2009), and the bottom panel on a VAR in the real price of oil and real GDP growth for the same sample period of 1975Q1–2008Q4. Responses are to a 1-standard-deviation shock and are expressed as cumulative quarterly annualized growth rates. Dashed lines are 1-standard-error bands.

price surge was driven primarily by global aggregate demand shocks (as opposed to a more typical mixture of demand and supply shocks), the upper panel focuses on the response of U.S. real GDP to a positive aggregate demand shock. The lower panel shows the response of U.S. real GDP to a real oil price innovation estimated on the same sample. That response reflects the average composition of demand and supply shocks over the entire sample period. Although broadly similar, the exact timing, the magnitude, and at times even the sign of the response estimates differ. A positive global aggregate demand shock has positive, if statistically insignificant, effects on real GDP within the first year, reflecting the sluggish response of industrial commodity and oil prices and the economic stimulus from abroad. Only starting in the second year does the response turn negative, as the stimulus fades and higher oil and industrial commodity prices stifle

Figure 4. Historical Decomposition of U.S. Real GDP Growth, 1978–84 and 2002–08[a]

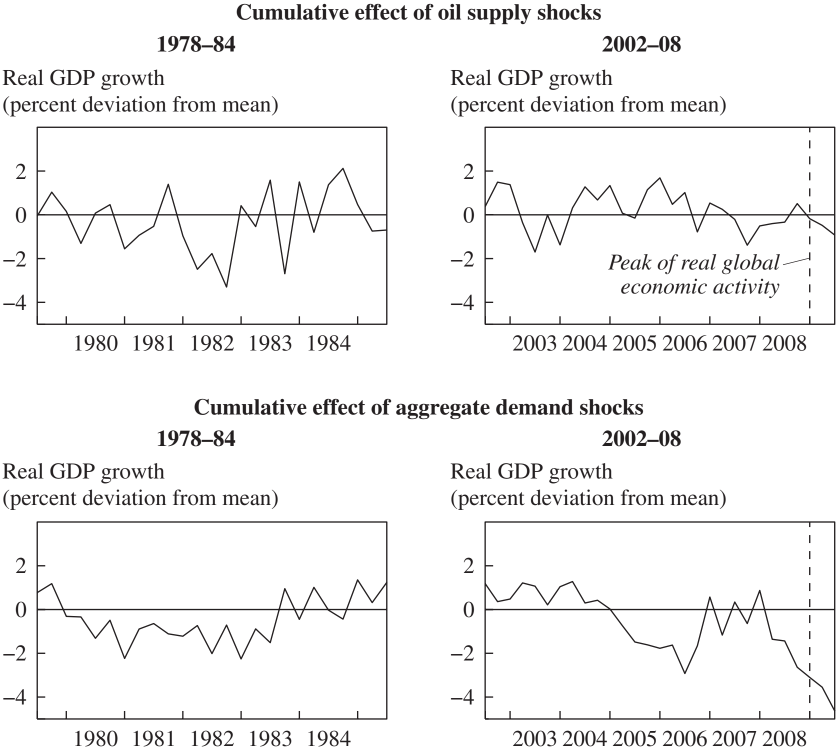

Source: Author's estimates.

a. Estimates based on the methodology of Kilian (2009) using data for February 1973–December 2008. Data are cumulative quarterly annualized growth rates.

economic growth. A shock of average composition implies a decline in real GDP starting in the second quarter, in contrast. The response is significant in the second and third years. An immediate implication is that the effects associated with the 2007–08 oil price increase could not possibly be the same as those associated with earlier oil price shocks, even if it were the case that the earlier shocks were driven by oil supply disruptions. Just how different the implied effects on real GDP are is documented below.

In discussing the impact of oil demand and oil supply shocks over extended periods, it is essential to consider the cumulative effect of all of these shocks over time rather than the response to a one-time shock. Figure 4 shows the contribution of aggregate demand and oil supply shocks to the observed variation in U.S. real GDP growth, relative to average growth, for selected periods. All of the results are based on the methodology of Kilian (2009). Compared with previous estimates from similar models but shorter time spans, figure 4 reveals a somewhat larger impact of oil supply disruptions on U.S. growth in the early 1980s, coupled with a persistent reduction in growth associated with aggregate demand shocks. In contrast, the primary explanation of below-average U.S. real GDP growth after 2004 is the unexpected increase in global real economic activity that started in 2002. Consistent with the impulse response estimate in figure 3, the initial effect of positive aggregate demand shocks on U.S. real GDP growth was largely positive (see figure 4). Only in 2004 do industrial commodity price and oil price increases start taking their toll. The top panel of figure 4 also suggests that these effects were offset—for some time—by the growth-enhancing effect of positive oil supply shocks. Given the unprecedented drop in global real activity of close to 95 percent from the peak in June 2008, shown in figure 1, the fact that the estimated effect is increasing sharply at the end of the sample does not come as a complete surprise.

How do these estimates compare with conventional estimates based on VAR models for real GDP growth and the real price of oil? Figure 5 shows that the fully structural VAR model predicts somewhat larger economic contractions in the early 1980s and in 2008 than the VAR model involving real oil price innovations. Interestingly, the overall effect of the demand and supply shocks on U.S. real GDP growth in 2005 through 2007 proved small by the standards of the early 1980s. The negative effect of earlier aggregate demand shocks was initially offset in part by the positive effects of other shocks, including positive oil supply shocks in 2004–06. Thus, only in late 2007 and in 2008 did the full effect of the continued unexpected global expansion make itself felt. This result corroborates the interpretation of the 2007–08 data in the paper. It is also consistent with the observation that the higher oil and industrial commodity prices triggered by repeated positive aggregate demand shocks, as in earlier episodes, caused a reduction in consumer spending, mainly in the residential housing sector and in the automobile sector. What is interesting about the results in figure 4 is that they suggest a somewhat smaller role for the financial crisis in late 2008 than one might have suspected. In that sense I agree with Hamilton that developments in global oil and other commodity markets appear to have played an important role in the latest U.S. recession.

In closing, although this paper presents an impressive body of evidence that sheds light on the mechanics of how oil price shocks are propagated, and although I agree with many of its substantive conclusions, my concern is that the narrow causal interpretation of oil price shocks in this paper is

Figure 5. Explanatory Power of Oil Demand and Supply Shocks Combined and of Real Oil Price Shocks, 1979–85 and 2002–08[a]

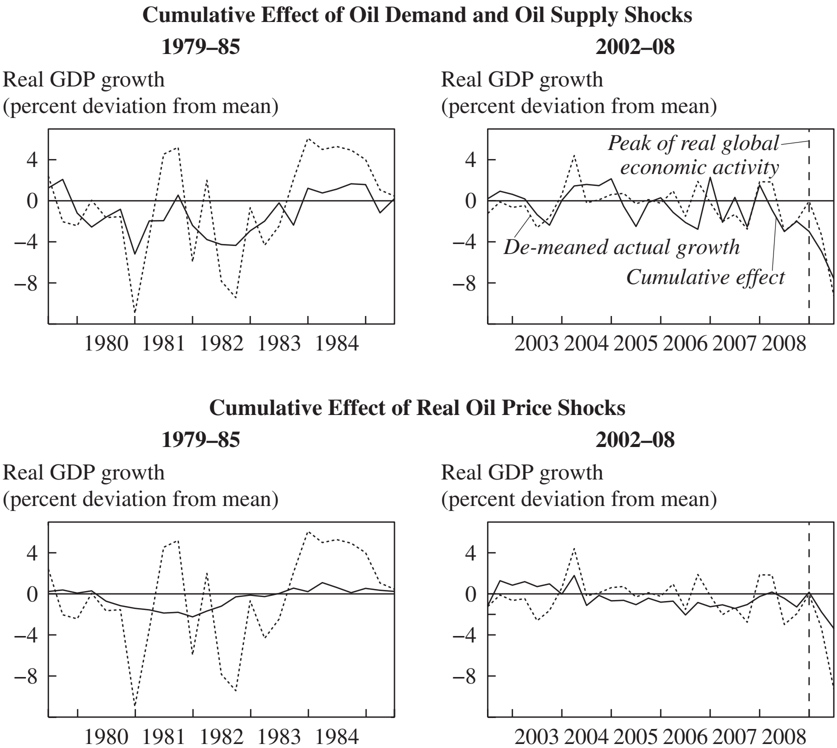

Source: Author's estimates.

a. Estimates based on the methodology of Kilian (2009) using data for February 1973–December 2008. Data are cumulative quarterly annualized growth rates.

misleading. This is not merely an issue of how to interpret the resulting responses, but one that affects the magnitude of the estimates. For example, although the direction of the estimated effects is broadly similar, the evidence in figure 5 above suggests that the aggregate demand shocks driving the 2007–08 oil price increase may have had greater effects on U.S. real GDP than suggested by models that ignore changes in the composition of oil price shocks.

I have also provided evidence that, contrary to the assertion in this paper, demand shocks have always played an important role in oil markets. What is different about the latest episode is that the oil price increase was driven almost exclusively by one specific type of demand shock, reflecting continued unexpected increases in global real activity during 2002–08, primarily associated with unexpected growth in emerging Asia (see Hicks and Kilian 2009). From a policy point of view the central question is how much of that unexpected growth reflected an exogenous economic transformation in emerging Asia. An alternative explanation is that the Federal Reserve sustained growth in the United States longer than appropriate, easing monetary policy too early and too much, thus enabling the export-based Chinese economy and the world economy more generally to thrive, and fueling the commodity and oil price boom that contributed to the current collapse of the real economy. I agree with the author that this possibility deserves careful study. A third explanation is that the sustained prosperity in the United States between 2002 and mid-2008 was not directly linked to monetary policy, but to the failure of the Federal Reserve and other regulators to rein in financial and housing markets. It seems unlikely that one could unravel the relative contribution of each explanation without the help of a fully specified multicountry open-economy model.

REFERENCES FOR THE KILIAN COMMENT

Barsky, Robert B., and Lutz Kilian. 2002. "Do We Really Know That Oil Caused the Great Stagflation?" *NBER Macroeconomics Annual,* vol. 16, pp. 137–83.

Hicks, Bruce, and Lutz Kilian. 2009. "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003–2008?" Discussion Paper DP7265. London: Centre for Economic Policy Research.

Kilian, Lutz. 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics* 90, no. 2: 216–40.

_________. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99, no. 3: 1053–69.

Kilian, Lutz, Alessandro Rebucci, and Nikola Spatafora. 2009. "Oil Shocks and External Balances." *Journal of International Economics* 77, no. 2: 181–94.

National Commission on Supplies and Shortages. 1976. *The Commodity Shortages of 1973–1974: Case Studies.* Washington: U.S. Government Printing Office.

GENERAL DISCUSSION Robert Gordon pointed out what he saw as four omissions in the paper. The first was the lack of a systematic macroeconomic theory of commodity supply shocks, such as that developed by Edmund Phelps and himself in the mid-1970s. A low price elasticity of demand for oil means that the energy share in GDP must rise, with the nonenergy share falling by the same amount. A complete theory would take into account both the direct effect of higher oil prices in the production function and the indirect effect of crowding out private nonoil real spending. The second was a lack of systematic treatment of inflation. Over the past three decades the response of aggregate inflation to higher oil prices has become more muted. This may have occurred through a smaller response of wages, possibly because of the declining power of labor unions and the decreasing prevalence of cost-of-living escalators in wage contracts. The third was that the paper omitted the housing bubble as a cause of the 2007–08 downturn, instead fully attributing the collapse to the oil price shock. Finally, four additional, interacting shocks occurred during the recession of the mid-1970s and were not considered in the paper: food price shocks, the depreciation of the dollar, the end of the Nixon price controls, and a general slowdown in productivity growth.

Robert Hall commended the author on yet again convincingly showing that every recent U.S. recession has been caused by an oil price spike. He recommended, however, that the data be continued beyond the moment of the spike itself. In particular, durable spending, dominated by falling automobile sales, declined steeply until October 2008 and has since then been essentially constant. If financial stress were the most important cause of the current recession, he posited, spending would not have stabilized at that point. Hall also discussed the paper's demonstration of an asymmetry in the economic response to oil price changes. In 1986 oil prices dropped sharply but did not cause an economic boom; Hall conjectured that if prices had instead risen by the same amount, there would have been a deep recession in 1987. In the current recession, fuel prices at the end of 2008 were significantly lower than at the beginning of the year, yet the economy did not improve and automobile sales remained low.

George Perry disputed the claim that oil prices were the prime cause of all recent U.S. recessions. First of all, no important rise in oil prices preceded the 2001–02 recession, and the price spike in 1991 lasted only a few months. A more interesting difference was that the Federal Reserve's response to the oil price shocks differed across time, and the responses, in turn, were informed by how oil prices were expected to impact core inflation. In the mid- and late 1970s' oil price shocks, wages in important industries were closely indexed to the consumer price index, leading to significant risk of a price-wage-price spiral in core inflation. The Federal Reserve responded aggressively. Indeed, it deliberately caused the 1980 and 1982 recessions in response to worsening inflation. By contrast, its response to the most recent oil price shock was very mild because inflation has been low in recent years and wages have not moved in response to price shocks.

Justin Wolfers, citing the discussion of the Romer and Romer paper in this volume, noted that four events may not be enough to prove that oil price shocks are the sole cause of recessions; after all, the chance of four consecutive coin tosses being heads is 1 in 16, whereas statistical significance conventionally requires a probability of 1 in 20 or better. He suggested that if professional recognition could be achieved by being correct four times out of four, many economists would choose that route rather than, say, the much more arduous one of developing a macroeconomic model.

Frederic Mishkin seconded George Perry's comments about the Federal Reserve's differing reactions to increases in oil prices. The better grounding of inflation expectations in the recent period has allowed the Federal Reserve to keep interest rates lower than would have been possible before, given the same rise in oil prices. Better-grounded inflation expectations also give the Federal Reserve much more flexibility in dealing with negative demand shocks.

David Romer conjectured that if, before reading the paper or hearing the presentation, all the conference participants had been asked to state their view on the primary cause of the current recession, oil prices would likely not have garnered many votes. Yet the paper had completely convinced the two discussants, at least, of its remarkable assertion that oil was the primary cause. Romer disagreed, arguing that the paper significantly overstated the role of oil in several ways. First, the effect of oil prices is known to be nonlinear, making it possible for an equation to overfit and overestimate the effect. Second, the regression in the paper that produced the strongest result did not correct for the fact that the energy intensity of the economy has decreased over time. Romer also cited the reasons already given as to why oil is less of a factor in macroeconomic stability today than in the past: wages are more flexible, inflation behaves differently, and the Federal Reserve knows better how to deal with oil price changes.

Caroline Hoxby discussed what she saw as an important implication of the paper, namely, that automakers deserve more of the blame for the current recession than has been attributed to them. Their failure to diversify into smaller, more fuel-efficient automobiles suggested that one way in which each oil shock has contributed to causing a recession was through decreased demand for U.S.-made cars. In fact, a large movement in American buying preferences from big cars to smaller ones has followed every oil shock. If U.S. manufacturers had invested in small cars in addition to SUVs, less income would have been lost to other countries such as Japan.

Richard Cooper commented on the paper's use of world growth rates as though they were actually known. He contended that the actual growth rate of gross world product is not known, and the rate most commonly used in analysis, that published by the International Monetary Fund, weights countries by purchasing power parity, which overweights countries like China and underweights countries like Japan; the result is a 1-percentage-point overstatement in the overall rate. An accurate measure of demand growth, he proposed, would use market exchange rates. Cooper also wondered about the role of speculation interacting with the fundamental determinants of demand in causing the sharp increase in oil prices. Much as demand for a medical school education might spuriously appear to have risen because the average applicant applies to more schools, so total demand for oil might appear to have risen because of the increase in speculative demand. During the second OPEC shock, Iranian oil production declined steeply, from 4 million barrels a day to 500,000 barrels, significantly hurting BP, which obtained much of its supply from that country. BP's customers, such as New Zealand, scrambled to order as much oil as they could from anyone who would take orders. This kind of behavior caused the apparent demand for oil to exceed actual final demand, resulting in rising inventories in 1979. In 2008, hedging purchases (for example, by airlines) plus speculative purchases (for example, by hedge funds) in the forward market drove up spot prices—which are determined by near-term forward prices—even when the buyers did not intend to take delivery on all of their purchases.

Matthew Shapiro questioned the vertical supply curves used in the paper. In markets for natural resources such as oil, the supply curve in any one period should be approximately horizontal. That is, the flow supply is highly elastic because of intertemporal arbitrage. He agreed that vertical curves might be appropriate in the case of a short-run supply restriction, such as an OPEC embargo. The short-run supply curve can also be steep if there are adjustment costs. Indeed, evidence from a comparison of the asset value of oil in the ground with the spot price of oil during 2008 suggests that short-run supply considerations affected both the increase in oil in the first half of the year and its collapse after the onset of the financial crisis. The slope of the short-run supply curve cannot, however, have a major role in explaining a steady increase in prices over a seven-year period.

Philippe Aghion considered the implications of the different types of oil shocks on technological innovation. He noted that in the 1970s the shock had been on the supply side, which had led to research and development in order to deal with a supply shortfall. The current shock was on the demand side but also involved the issue of global warming. He acknowledged that a much longer time series would be necessary to study these effects.

Benjamin Friedman observed that the paper's results embedded two interesting propositions. The first was that recessions, and output reactions generally, are invariant to monetary policy. The conventional wisdom holds that a backward shift in the aggregate supply curve presents the Federal Reserve with the unpleasant choice of how much of the shift to absorb in inflation and how much to let a backward shift in aggregate demand cover. The paper insinuated that this notion is wrong, that the actions of the central bank are nothing more than a sideshow. The second proposition was that output appears invariant to the conduct of monetary policy because it was constant across all episodes discussed. Friedman found this proposition quite interesting because Republicans were in charge during some episodes and Democrats in others, inflation was high at the beginning of some episodes and low at the beginning of others, and inflation expectations were well anchored in some but not anchored at all during others. If monetary policy was truly invariant across the entire sample, he noted, that would be an interesting finding in and of itself.

Lutz Kilian responded to a number of issues raised in the discussion. He clarified that his view was that unexpected oil price increases are not the one factor that explains recessions but rather are a symptom of deeper shocks to the underlying oil supply and oil demand that affect the economy both directly and indirectly through the prices of oil and other imported commodities. Hence, attempts to distinguish between alternative monocausal explanations of recessions such as "oil" or "money" are inherently misleading. He also observed that, historically, residential housing expenditure has been particularly sensitive to oil price shocks, consistent with Hamilton's analysis. Regarding the role of asymmetries, he noted that contrary to the conventional wisdom, significant asymmetry is not found in the overall economic response to oil price shocks but may appear in the response of nonresidential investment as a consequence of identification problems and of composition effects. Addressing the comments about the Federal Reserve's reaction to oil price shocks, he noted that recent research shows that there is no theoretical basis for a mechanical response of monetary policy to oil price shocks, nor is there compelling empirical evidence that the Federal Reserve's response to oil price shocks has had substantial effects on macroeconomic outcomes in the past.

John Campbell remarked that physical activity has to occur in the spot market for oil in order to enforce arbitrage in the futures market, but such activity may never occur in equilibrium because enforcement is not needed. He compared trade in the oil market to that in financial markets, where there is both trade in the assets themselves and trade in futures and other derivatives of those assets; in a well-functioning economy, transactions across these market segments do not have to occur to establish the law of one price. Specifically, he pointed out that an exogenous increase in the future oil price allows spot prices to rise without inducing a strong supply response, by contrast with the situation in which high current prices and lower expected future prices generate a large supply reaction. He wondered why oil market participants had formed expectations of persistently high oil prices in the future; a possible explanation was some combination of rational speculation, irrational speculation, and speculation that confused other participants into forming beliefs about future prices that turned out to be unreasonable.