Oxford Review of Economic Policy, Volume 27, Number 1, 2011, pp. 169–185

Oil price shocks and the macroeconomy

Paul Segal*

Abstract This paper examines the impact of oil price shocks and attempts to explain why the rise in oil prices up to 2008 had little impact on the world economy. It makes three main arguments. First, that oil prices have never been as important as is popularly thought. Second, that the most important route through which oil prices affect output is monetary policy: when oil prices pass through to core inflation, monetary authorities raise interest rates, slowing growth. Based on the second argument, the third argument is that high oil prices have not reduced growth in recent years because they no longer pass through to core inflation, so the monetary tightening previously seen in response to high oil prices is absent. It also argues that oil prices had little impact on the global recession of 2008–9.

Key words: oil prices, business cycles, monetary policy, inflation

JEL classification: E23, E31, E32, E58, Q41, Q43

I. Introduction

Oil prices and economic cycles have been firmly linked in the public imagination since the oil shocks of the 1970s and the global recessions that followed. Spurred by these events, economists in the 1980s analysed the relationship in a number of econometric studies, demonstrating a negative correlation in the US and other industrial countries between oil prices and macroeconomic performance. Yet, while the association is unambiguous at least in the US up to the early 1980s, at some point thereafter the relationship appeared to attenuate. More recently, in the lead-up to the global crisis that started in 2008, oil and other commodity prices rose to historic levels, then to collapse as the crisis struck. What impact have oil prices had on the macroeconomy, and what was their role in the recent crisis? Have they been more cause or consequence of macroeconomic cycles?

This paper suggests that the empirical evidence supports three arguments. First, oil shocks have never been as important as is commonly thought. Oil prices are just one more macroeconomic variable, and the view that they are ever the main determining factor of an economic downturn is not consistent with the evidence. While they have never been decisive, they have, however, played a role, and the second argument is that the most important component of the impact that oil prices have on short-run output runs through monetary policy. If oil prices raise inflation, then monetary authorities raise interest rates, slowing

*University of Sussex and Oxford Institute for Energy Studies, e-mail: paul.segal@oxfordenergy.org
doi: 10.1093/oxrep/grr001

activity. This argument complements the first argument because monetary authorities make policy in the light of the full range of economic news.

The third argument uses the first two arguments to answer the following question: why is the impact of high oil prices so much smaller today than in the 1970s? Given the importance of monetary policy in the causal chain, I argue that a key difference is that oil prices no longer feed through to core inflation, so that monetary policy no longer has to tighten in response to high oil prices. Indeed, as I discuss later, high oil prices can have a deflationary impact, requiring monetary policy to loosen. Precisely why there is less inflation pass-through remains unclear, but it is probably due to more flexible real wages, and to more credible monetary policy.

History illustrates this view. In the 1970s, the major economies of the world were suffering from high inflation and low growth independently of the rise in the oil price. Monetary policy responded to the high level of inflation, and the oil price contributed to, but was not the sole cause of, this inflation. On the other hand, the dramatic rise in oil prices up to mid-2008 had little inflationary impact and therefore did not lead to monetary tightening. Consistent with the above arguments, the evidence suggests that the oil price rise had no more than a marginal role in the major global downturn that followed.

This view of the role of oil prices contrasts with several arguments regarding their microeconomic impact, and I present a sceptical view of the microeconomic mechanisms described in the literature.

I also discuss the genesis of oil price rises and the popular argument that the minimal impact of high oil prices in recent years is due to their being 'demand driven', as opposed to the 'supply driven' price rises of the 1970s. I argue that the marginal effect of the oil price rise *per se* should not be affected by the source of the rise—although the state of other macroeconomic variables will obviously affect macroeconomic outcomes.

The focus of this paper is macroeconomic cycles in industrial countries, so I do not discuss the effects of oil price rises on developing countries. They may be very deleterious for a developing country that relies substantially on oil imports, or very positive for an oil exporter. But these effects, and the mechanisms through which they operate, are somewhat different from those discussed here.

Section II surveys the main findings regarding the associations between oil prices and macroeconomic indicators. In section III I discuss a variety of microeconomic mechanisms between oil prices and macroeconomic outcomes that have been suggested in the literature. Section IV discusses the role of monetary policy, considering both empirically based and dynamic stochastic general equilibrium (DSGE) model-based analyses of the nexus of oil prices, monetary policy, and output. Section V then uses this conclusion to explain the decline in the impact of oil price rises since the 1980s, including the oil price rise of the last few years. Section VI concludes.

II. The key relationships

The key finding that established the case that oil shocks cause recessions was Hamilton's (1983) result that, up to 1980, 'all but one of the US recessions since World War II have been preceded, typically with a lag of around three-fourths of a year, by a dramatic increase in the price of crude petroleum' (p. 228). A flurry of research then confirmed this relationship across a number of rich countries. Burbidge and Harrison (1984) used vector autoregressions (VARs) to estimate the impact of oil price rises in Canada, Japan, West Germany, the UK, and the USA, using data from 1961 to 1982. Estimating the impact of a temporary one-standard-deviation spike in the oil price, they find that 'oil-price shocks increase wages and prices in all countries, albeit with appreciable variation in the size of the effect' (p. 468). GDP is not included in the VAR, but industrial production declines substantially in the US and Japan, and much less so in the other countries. Also using a VAR, Blanchard and Galí (2007) similarly find that oil price rises over 1970–83 had a strong effect on most rich countries, but very little effect on inflation in West Germany, or on inflation or output in Japan. Mork (1994) observed (without recourse to any formal statistical analysis) that, in addition to the USA, all of Japan, West Germany, and the UK suffered declines in real GDP after the oil shock of 1973–4, while of these countries only Japan avoided a decline after the 1979–80 shock. A broad consensus emerges, therefore, that only Japan avoided some strong negative effects of the oil price rises of the 1970s.

Quantifying these effects into elasticities, Rotemberg and Woodford (1996) estimated that a 1 per cent rise in the price of oil is associated with a reduction in US output of about 0.25 per cent after five to seven quarters, while Leduc and Sill (2004) find that a doubling of the oil price is associated with a 4.5 per cent drop in US output. IMF (2000), using a calibrated model of the global economy (MULTIMOD), estimated that an unexpected $5 rise in the price of oil in 2001 would lead to a drop in global output of 0.3 per cent in each of the first 2 years, with the US and the Euro area down by 0.4 per cent. More recently IMF (2007, p. 17), using its Global Economic Model (GEM), found that a supply-induced doubling of the oil price leads to a slow-down in world GDP of 1.4 per cent at its trough, a little over 1 year after the shock, while world inflation rises by 1.5 per cent at its peak, after two quarters. This wide range of papers supports Hamilton's core finding that oil price rises reduce GDP.

The first wrinkle to be documented in the relationship between oil prices and economic activity was Mork's (1989) finding of a statistical asymmetry in the impact of oil price rises and declines in the US. He found that oil price rises were followed by declines in GDP, but that declines in the oil price—including the drop of 1985–6—were not followed by rises in GDP. Mork *et al.* (1994) found a similar asymmetry in Japan, West Germany, France, Canada, the UK, and Norway over 1967–92.

The second wrinkle was rather more serious. Darrat *et al.* (1996) and Hooker (1996*a*) find that when more recent data are used, going up to the early 1990s, then the oil price no longer appears to cause declines in output. Darrat *et al.*, using a 6-variable VAR over 1960–93, find that industrial output Granger-causes[1] oil consumption, but that neither oil consumption nor oil prices Granger-cause industrial output, even allowing for asymmetric effects. In addition to having a later time period, their study differs from Hamilton's (1983) analysis by including the interest rate as a variable. They observe that their VAR implies that 'oil prices exert significant effects on the monetary base . . . and then the monetary base causes significant shifts in industrial production. . . . Thus, accommodative monetary policy appears to have weakened any economic consequences of oil price shocks' (p. 328). We return to this argument later, in section IV. Hooker (1996*a*) similarly finds that, allowing for asymmetry and other possible misspecifications, the oil price did not Granger-cause changes in US GDP for the sub-period 1973:4 to 1994:2. Hamilton's (1996) response argues for the use of his net oil price increase (NOPI) variable—which registers only those rises in the oil price that are

[1] For variable x to 'Granger-cause' variable y means that, when y is regressed on lagged values of itself and of x, then some lagged values of x are significant.

not just reversing declines of the preceding year—but even NOPI fails to Granger- cause GDP in the period in question (Hooker, 1996*b*).

Evidence also emerged that, even in the 1970s, the effect of oil price shocks was smaller than subsequently believed. Hunt (2005) considers the role of the oil price shocks in the US recession of 1974–5 using a version of the IMF's GEM, and finds that 'the simulation results do not suggest that the oil price shock alone can account for the extent of the slowing in real activity and the acceleration in inflation that occurred in the United States in 1974 and 1975'. Similarly, Bernanke *et al.* (1997) find that the 1974–5 recession is not well explained by the oil shock, even allowing for an endogenous monetary policy reaction. More important was a spike in general commodity prices including food, and the rise in interest rates caused by that spike. The decline in output over 1979–81, on the other hand, is well explained in their analysis by the oil shock and subsequent monetary policy response. But the continued decline after this point was, according to the analysis, due to an autonomous tightening of monetary policy in 1980–1 that was independent of oil prices.

One key component of the impact of oil prices on GDP is the effect of the latter on inflation. Consistent with Bernanke *et al.* (1997), Blinder and Rudd (2008) find that food price rises and other factors were more important than the oil price shocks in explaining the two large inflation spikes in the 1970s. Hooker (2002) finds that oil prices fed through to core inflation in the USA up to 1981, but not after that point. Blanchard and Galí (2007) run VARs for the sub-periods 1970–83 and 1984–2006 and find much lower impacts of oil prices on both inflation and output in the second period than in the first. As will be discussed later, the decline in feed-through to inflation appears to be an important part of the story.

While these studies find that oil price shocks no longer have a significant impact on the macroeconomy, Barsky and Kilian (2004) argue that they have never been a major driver of macroeconomic cycles. They point out that oil prices are themselves partly determined by global macroeconomic conditions. They also point out, following Mabro (1998), that the impact of exogenous events in the Middle East on the oil price depends a lot on the tightness of global demand for oil, so that embargo- or war-induced price spikes are partly due to global economic conditions, in addition to the disturbances themselves.

III. Mechanisms for the impact of the oil price on the macroeconomy

Up until the early 1980s, oil prices do seem to have had a significant impact on output. But it is surprisingly difficult for neoclassical economic models to account for this impact. For a neoclassical economist the most natural way to think of oil is as an input to the economy's production function. When an input gets more expensive, the profit-maximizing level of output declines. The standard way to present this argument (e. g. Hamilton, 2005) is with a simple model of a representative firm with the following production function:

$$Y = F(L, K, E)$$

where L is labour, K is capital, and E is energy input. With output price p, wage w, capital rent r, and the nominal price of energy p_E, profits are

$$pY - wL - rK - p_E E.$$

In a competitive market the firm buys energy up to the point where its price is equal to its marginal value product:

$$p_E = pF_E(L, K, E)$$

where F_E is the partial derivative of F with respect to E. Multiplying both sides of this equation by E and dividing by pY,

$$p_E E/pY = F_E(L, K, E)E/Y.$$

The right-hand side is the elasticity of output with respect to energy use, while the left-hand side is the share of energy expenditures in total output.

So how much of an impact can this model account for in the US data? In most of the late 1970s the share of energy expenditures in total output was 4–5 per cent, rising to 8 per cent at the peak of the oil price spike in 1979–80. It was much lower, at 1–2 per cent, during the 1990s and early 2000s, and in 2008 it peaked at 4.8 per cent. The largest drop in oil use in the US occurred from 1978, when 6.9 billion barrels were consumed, to 1983, when 5.6 billion barrels were consumed—a decline of 19 per cent. Multiplying 19 per cent times the peak expenditure share of 8 per cent (thereby overestimating the effect) yields a drop in GDP, over 5 years, of 1.5 per cent. The largest drop in oil consumption in a single year was from 1979 to 1980 with a decline of 7.4 per cent, which would yield a drop in GDP of 0.6 per cent. Negative shocks of these magnitudes are not nearly enough to turn a normal period of growth into a recession, yet, as we have seen, recessions in a number of major economies are what followed both the 1973 and the 1979 oil price shocks.

For this reason, several papers have suggested alternative routes through which the rise in the oil price can itself lead to greater declines in output than predicted by the basic model above. Hamilton (1988) develops a model, with flexible prices and wages, in which a rise in energy prices causes a decline in wages and employment in one sector relative to the rest of the economy. Assuming a delay attached to starting new employment, this leads to (*a*) frictional unemployment, and (*b*) the possibility that some workers do not take a job in another sector in the expectation that their old (and better-paying) job may reappear when energy prices decline again.

Davis and Haltiwanger (2001) describe these mechanisms as ‘allocative channels’ or ‘the aspects of oil price changes that alter the closeness of the match between the desired and the actual distributions of labor and capital inputs’ (p. 468). They empirically test these channels using the example of the car industry, observing that:

> the oil price shock of 1973 increased the demand for small, fuel-efficient cars and simultaneously reduced the demand for larger cars. American automobile companies were poorly situated to respond to this shock, because their capital stock and work force were primarily directed toward the production of large cars. Consequently, capacity utilization and output fell in the wake of the oil price shock, even though a handful of plants equipped to produce small cars operated at peak capacity. (pp. 466–7)

The contrast with Hamilton (1988) is that Hamilton considered the possibility of unemployment owing to simple frictions and to a reluctance by workers to take worse-paid jobs when a better-paid job might re-appear. Davis and Haltiwanger, on the other hand, assume that the sector in which demand rises—the production of smaller cars—cannot employ all the workers discharged by the disfavoured sector because of the nature of the physical (and perhaps human) capital stock.

This allocative channel contrasts with the 'aggregate channels' of movements in aggregate supply or demand curves, such as the decline in supply illustrated by the simple neoclassical model, and the demand impact of monetary policy discussed below in section IV. Davis and Haltiwanger (2001) use their plant-level industry data to test the relative importance of the aggregate and allocative channels. The allocative channel should lead to a rise in both job destruction and job creation when oil prices rise, as labour demand rises in relatively favoured sectors (even if capacity constraints prevent job creation from keeping up with job destruction). The aggregate channels, on the other hand, should lead to a rise in job destruction and a *fall* in job creation. Moreover, the allocative and aggregate channels should respond differently to declines in oil prices: a decline should have much the same effect from the allocative point of view, probably producing a net reduction in employment, whereas from the aggregate point of view, they argue, it could be expected to increase employment.

Using plant-level data they run a VAR across industries, using oil prices, a measure of credit to represent monetary conditions, and other sector-specific and common variables. They find that job destruction rises sharply in response to an oil price rise, peaking at the fourth quarter and then dropping to slightly below its baseline level after about eight quarters. Job creation declines, reaching its trough in the fourth quarter and then rising to somewhat above baseline from the eighth quarter. The net effect after 16 quarters is approximately zero, with 0.63 per cent of all manufacturing jobs having been relocated by this point.

The authors state on this basis that 'the response pattern fits the profile of an ''allocative disturbance''' (p. 488). This seems wrong. The initial response of falling job creation is what one would expect from the aggregate channel, not the allocative channel. It is only after 2 full years that job creation becomes positive in their estimate. After such a length of time an economy can have fallen into a recession and then climbed back out of it, with jobs that were destroyed on the way down being replaced on the way back up. Indeed, on the allocative interpretation there is no reason why job destruction would decline to below baseline at the same time as job creation rises above baseline, as in their estimates. This is explicable as a result of aggregate economic recovery, but there is no 'allocative' reason why the increased creation of new jobs in relatively favoured industries should coincide with the reduced destruction of jobs.

The possibility that non-competitive behaviour in markets may explain a large oil price effect is explored by Rotemberg and Woodford (1996). They specify a model with three different forms of non-competitive behaviour, two involving endogenous mark-ups, and find that these models more closely replicate their empirical estimate than a competitive model.

In response to Rotemberg and Woodford, Finn (2000) presents a model with a relatively large oil price effect, but which is competitive and frictionless. The key novelty in her neoclassical model is that the rate of capital depreciation is increasing and convex in the level of capital utilization. Hence depreciation costs provide an additional reason not to utilize fully the capital stock when other inputs (such as energy) are more expensive.

Finn argues that her calibrated model plausibly replicates Rotenberg and Woodford's empirical estimates of the impact of an exogenous oil price rise. However, the peak response in her capital utilization model, around the fifth to sixth quarter, is less than two-thirds of the empirical estimate (about 63 per cent, reading off her figure). In the absence of a formal test of similarity the accuracy of this estimation has to remain a matter of opinion, but I would suggest that 63 per cent is not a very close fit.

Moreover, in both Rotemberg and Woodford's and Finn's cases, the theoretical estimates produce U-shaped response curves, while Rotemberg and Woodford's empirically estimated response curve has the shape of a wave, declining first very slowly and then more rapidly, before recovering. That is, in the data the oil price rise has only a very small impact in the first two quarters, with output dropping more rapidly from the third and fourth quarters. None of the calibrated models produced by these authors are able to replicate this pattern.

One motivation for the above studies was an important paper by Bohi (1991), who considered industrial output data in Germany, Japan, the UK, and the USA at the 3-digit level. Looking at the two oil shocks in the 1970s, he found that there was no significant correlation between the energy intensity of a sector and the extent of its downturn in the first oil shock, while in the second shock there was a significant correlation only in Germany and the UK, not in the US or Japan. Jones *et al.* (2004, p. 5) argue that 'much of [the] costly reallocation occurs at the 4-digit SIC level and consequently remains invisible at the more aggregated, 3-digit data such as those used by Bohi'. But if these re-allocation mechanisms were all that costly, then we would certainly see the results of those costs in terms of reduced output at the 3-digit level. One way to look at Bohi's finding in the light of Davis and Haltiwanger's plant-level findings, therefore, is to note that the churning found by the latter study is just not very costly. It also casts doubt on Finn's (2000) mechanism: if capital utilization rates are heavily affected by oil prices, then the impact would certainly be felt more strongly in more energy-intensive sectors.

These papers have attempted to explain a large oil price effect by positing non-standard mechanisms. While these mechanisms could plausibly play a role in the oil price-macroeconomy relationship, none was convincingly shown to be capable of explaining substantial movements in GDP. Most importantly, Bohi's finding that more energy-intensive industries generally do not suffer worse downturns casts further doubt on their salience. We now turn to an alternative form of explanation for the impact of oil prices.

IV. Oil shocks and monetary policy

Monetary factors and aggregate demand provide an alternative explanation for the historical impact of oil prices. To understand the impact of an oil price shock from the macroeconomic point of view, it can be decomposed into two stages. In the first stage the price of petroleum products rises. Assuming that other prices are sticky downwards, this external inflationary shock raises the domestic price level. This implies lower real money balances, higher interest rates, lower demand, and hence lower output. If oil imports comprise 2 per cent of expenditure and the oil price doubles, then, assuming a low elasticity of oil demand, expenditure on other goods will decline by about 2 per cent. Such a shock will require monetary loosening. In Solow's (1980, p. 263) analysis of the 1973–4 oil price shock, he notes that 'between the end of 1973 and the end of 1975, real M2 fell by 3 or 4 per cent; if potential output rose by a routine 6 or 7 per cent during those two years, then the equivalent reduction of the real money supply more like 10%'.[2] He argues that this tightening of the real money supply had a strong contractionary effect.[3] The impact on output could, therefore, have been ameliorated by relaxing monetary policy.

[2] He also argues that supply management through lower taxes would have been required, in addition to demand stimulus, to minimize the negative impact of the oil price shock.

[3] Contrary to this, Blanchard (2002) points out that despite rising nominal interest rates, rising inflation (and inflation expectations) implied that real interest rates declined to very low levels in 1974–5, the short-term real interest rate falling below zero.

Another way to look at the impact of this first-stage effect and its contractionary implications is to note the analogy between a rise in the price of oil and a rise in consumption taxes (Allsopp, 2006). As the Organization of the Petroleum Exporting Countries (OPEC) pointed out in the 1970s, oil importers could have countered the rise in the international wholesale price of oil by reducing their own domestic taxes on oil products. Private actors in the importing countries would thereby have experienced no change in prices and, instead, the shock would have been absorbed by government deficits. Similarly, from the point of view of the private sector a rise in the oil price is equivalent to a rise in taxes on petroleum products, with the difference being the impact on government debt. A rise in taxes is contractionary and, all else equal, would demand an offsetting cut in interest rates.

The second-round effects of the rise in the oil price follow when domestic actors are unwilling to accept the decline in real income caused by the first-round effect. If workers and producers demand to be compensated for higher fuel prices, then a wage–price spiral can develop. It is only in this second round that core inflation, excluding the prices of petroleum products, is significantly affected. One important point made by Mork (1994) is that this can explain the observed asymmetry in upwards and downwards movements in the oil price if prices (and wages) are sticky downwards but not upwards.[4]

The role of monetary policy can be seen in terms of a simple Taylor rule, under which the interest rate is rising in the output gap y and in the expected rate of inflation π:

$$i = \alpha y + \beta \pi.$$

A rise in the oil price can lead to an interest-rate response through both of these variables. The first-round effect is contractionary in the same manner as a rise in taxes, lowering y and encouraging the monetary authority to lower the interest rate. The second-round effect raises expected inflation in anticipation of a wage–price spiral, encouraging the monetary authority to raise the interest rate. The response would, therefore, depend on the relative strengths of these two effects.

The above model is very standard, but the question remains whether it is the right model for the current purpose. So why would one believe that monetary policy is the most important route through which oil prices affect the economy? There are two parts to the claim: first, that if one controls for monetary policy, then oil prices should lose their historical impact. Second, that oil prices do, indeed, affect monetary policy.

The first claim is supported by the VARs estimated by Darrat *et al.* (1996) and Hooker (1996*a*), which find that the impact of oil prices on US output disappears when interest rates are controlled for. Hooker (1999) then followed up by re-testing the relationship using both Hamilton's (1996) 'net oil price increase' measure and Lee *et al.*'s (1995) oil price transformation. The former measures only oil price rises that surpass the level of the previous year, while the latter divides oil price rises by their conditional variance, and both set price declines to zero. Hooker finds that one data point, 1957:1, is extremely influential and, without it, oil prices are generally no longer significant even in the pre-1970s data. The only way that oil price variables can be made significant is by using annual as opposed to quarterly data and excluding standard control variables, including the interest and inflation rates, from

[4] This is also consistent with Balke *et al.*'s (2002) finding that monetary policy reacts asymmetrically to oil price rises: a rise in the oil price combined with price stickiness causes monetary tightening which the central bank has to loosen, while a fall in the oil price leads to a rise in other prices, keeping the real money supply steady and requiring no extra response from monetary policy.

the equation. Interest rates therefore appear to be an essential route through which oil prices affect output. Kilian (2006, p. 24) also finds suggestive evidence: in his regressions short-term interest rates rise in response to exogenous oil supply shocks in all countries but Japan, and Japan is the country with the lowest GDP response to the shocks of the G7 countries.

A more detailed analysis that supports both of the above claims is given by Bernanke, Gertler, and Watson (1997) (henceforth BGW), who use a structured VAR to investigate the role of macroeconomic policy rules in determining the macroeconomic response to an oil price shock. Standard empirical estimates, including VARs, suffer from the Lucas critique: the structure of the economy is partly determined by agents' expectations about variables in the economy, including the interest rate, and the parameters estimated in the VAR reflect these expectations. Hence if the monetary policy rule (or reaction function) changes, and agents in the economy know that it has changed, then the parameters estimated under the original reaction function will no longer be appropriate. The counterfactual of a different monetary policy cannot, therefore, be tested simply by changing the interest rate variable in the estimated VAR. Put another way, a change in the monetary policy rule does not just change the behaviour of the interest rate, but it also changes the behaviour of economic actors in response to a given state of the world. An empirical simulation that merely changes the interest rate will, therefore, not capture the full effect of a change in the monetary policy rule.

To reduce the sensitivity of their analysis to the Lucas critique, BGW add some structure to their VAR by making two assumptions. First, that the federal funds (FF) rate and its expected future values determine market interest rates, and that market interest rates in turn determine variations in output. Second, they assume that there is no further mechanism through which the FF rate affects the economy.

With this model they run three different types of simulation in response to a rise in the oil price. First, they run a standard VAR simulation allowing the FF rate to change in response to the oil price, as in the data. Second, they simply fix the FF rate at its base value throughout the simulation. It is this exercise that is particularly vulnerable to the Lucas critique. Third, described as 'anticipated policy', they apply the methodology described above: they fix the FF rate at its base value, and in addition set market interest rates (and hence market expectations) to be consistent with this future path of the FF rate. This should not be vulnerable to the Lucas critique because expectations in this case have adapted to the monetary authority's behaviour.

In the second and third cases both prices and output are higher, as expected, with the third simulation slightly higher than the second. Thus 'a nonresponsive monetary policy suffices to eliminate most of the output effect of an oil price shock, particularly after the first eight to ten months' (BGW, p. 118). Their Figure 4, with the estimated responses, is repeated here (see Figure 1).

Their final experiment is to compare the impact of an oil price shock and accompanying endogenous monetary policy response with the impact of the same monetary policy change with oil prices kept constant. They find that the impact on output is virtually identical, although the impact on prices is not the same: prices first rise and then fall under an oil shock, whereas they just fall under exogenous monetary policy. The rise in prices under the oil price spike, they suggest, may be what explains the concomitant rise in interest rates—the 'second round' effect on inflation discussed above. The authors conclude that 'the majority of the impact of an oil price shock on the real economy is attributable to the central bank's response to the inflationary pressures engendered by the shock' (p. 122).

Figure 1: Responses to a Hamilton oil price shock

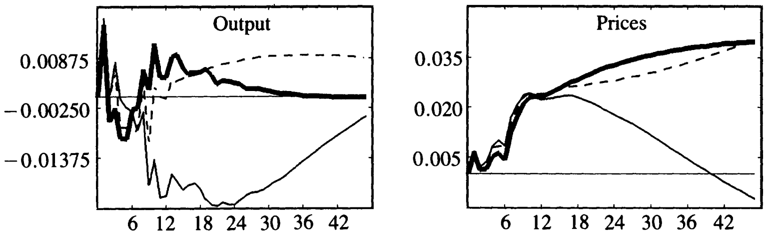

Notes: Thin line—standard VAR; dotted line—fixed FF rate, no expectations; thick line—fixed FF rate with consistent expectations.
Source: Bernanke *et al.* (1997, Figure 4, p. 117), reproduced with permission from the Brookings Institution.

Hamilton and Herrera (2004) criticize BGW on two grounds. First, that they allow oil price shocks to affect the economy up to a maximum lag of just 7 months (chosen by BGW on the basis of the Akaike information criterion). Second, they argue that BGW fall foul of the Lucas critique: that keeping the FF rate constant would require such heroic efforts on the part of the Fed that it would radically change expectations.

Are Hamilton and Herrera right to claim that BGW falls foul of the Lucas critique? Only if one applies a far more exacting standard of plausibility than anyone applies to the calibrated DSGE models that were designed to overcome the Lucas critique. This is because if the two assumptions that BGW make are true, then the model is expectations-consistent and the Lucas critique does not apply. But DSGE models normally make the same two assumptions, and many more besides; if one rejects the assumptions then one rejects a lot of macroeconomics.

Despite this, Bernanke *et al.* (2004) address both of these criticisms simultaneously by using quarterly data with six lags, and freezing monetary policy for only 1 year, allowing it then to follow its usual estimated path. In this case they find that monetary policy accounts for about half of the drop in output following an oil price spike. Unfortunately they do not make the adjustments one at a time, so we do not know the effect of responding to the valid criticism over lags, and ignoring the incorrect criticism from the Lucas critique. Clearly, the quarterly simulation with monetary policy persisting longer than a year would have resulted in a larger weight on monetary policy.

Leduc and Sill (2004) address the same question as BGW using a calibrated DSGE model instead of empirical VAR estimates. By using a fully specified model they are able to build in alternative monetary policy rules, with consistent expectations, in order to estimate how they affect the impact of an oil price shock. Wages and prices are both sticky, with quadratic costs to adjusting them at a rate different from the growth of the money supply. They follow Finn (2000) in assuming that the utilization rate of capital is an increasing function of energy use per unit of capital stock, and that capital depreciation is an increasing function of capital utilization.

Using a simple Taylor rule based on Fed policy over 1979–95 (estimated by Orphanides, 2001) for the monetary policy reaction function, they calibrate the model to replicate the finding of a 4.5 per cent drop in GDP in response to a doubling of the oil price. In order to estimate the impact of an oil price shock, they start by estimating the impact under what they call a ‘k-percent rule’: the money supply is increased by an exogenous (and unspecified) k per cent each period. They interpret this as ‘neutral’ monetary policy and describe the reaction of the economy to a doubling of the oil price under this monetary policy as the pure effect of the oil price shock. This is, therefore, a different interpretation of ‘neutral’ monetary policy from the constant FF rate simulated by BGW. When the impact on output under the k-percent rule is compared with the impact on output under the Taylor rule, monetary policy under the Taylor rule 'accounts for' the difference between the two estimated output responses. The difference is 37 per cent of the total response: that is, output falls by 37 per cent less under the k-percent rule than under the Taylor rule that the Fed is estimated to have actually followed. Other monetary policy rules, including a constant interest rate, are found to have less of an impact. Their conclusion is that monetary policy can be considered to make up 37 per cent of the impact on an oil price shock on output.

These two different analyses of monetary policy—the empirical estimates of BGW and the DSGE estimates of Leduc and Sill—both imply that monetary policy is at least an important factor in determining the impact of oil prices. The interaction of oil prices, inflation, and monetary policy is decomposed more explicitly by Hunt *et al.* (2001). These authors use a version of the IMF's MULTIMOD model. In their calibrations a 50 per cent rise in the price of oil has a direct contemporaneous effect on the consumer price index (CPI)—through gasoline and other direct energy costs—of 1.3 percentage points in both the US and the Euro area, 0.6 points in both Japan and the UK, and 0.8 points in Canada. The impact on core inflation then works through two routes in the model: the CPI appears directly in the core inflation equation, on the basis that workers demand wage rises in response to rises in the CPI, and that this leads to rises in costs and hence core prices; and expectations of future CPI rises also enter the core inflation equation. The core inflation equation has two coefficients for each of these two routes, which are estimated across the different countries.

When these parameters are set to their average levels across the industrialized countries, the CPI jumps up in response to an oil price spike and then is maintained at an elevated level. Core inflation rises more gradually, but then also remains steady at a higher level. When the parameters are set to zero the CPI experiences the same initial jump, but then settles back to base level, while there is no impact on core inflation. The model thereby captures the idea of oil price pass-through. Varying values for the two price pass-through parameters across countries also lead to different impacts on output, with higher pass-through generally leading to larger drops in output.

These studies all suggest that the inflationary effect of oil price rises and its resultant effect on monetary policy is an important route through which oil prices affect output. According to the empirical estimates of BGW, it appears to be the dominant route. One finding of particular significance is that a rise in the interest rate has the same impact on output regardless of whether it is accompanied by a rise in the price of oil. We now use these findings to help to explain the reduced impact of oil prices today.

V. The declines in the impact of oil shocks

We saw that oil prices seem to have lost their ability to shock macroeconomies since the 1980s. But how are we to interpret the massive rise in oil prices up to mid-2008 and the subsequent global recession?

Petroleum expenditure as a share of GDP for the world, the OECD, and the US is plotted in Figure 2. Taking the average of the three years leading up the shock and comparing it with the level at the peak of the shock, the rise in petroleum expenditure as a share of GDP was 3.2 per cent to 1974 and 3.6 per cent to 1979. The rise from the average over 2001–3 compared to the peak in 2008 is 3 per cent (2.8 per cent for the world).

Figure 2: Petroleum expenditures as a share of GDP (%)

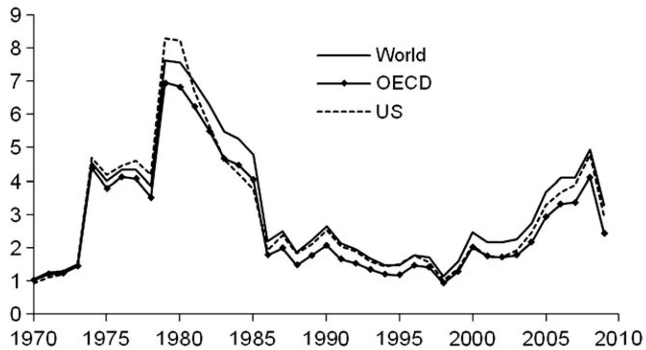

Source: Author's calculation using BP (2007), *World Development Indicators Online*, and *World Economic Outlook Online* data.

The rise up to 2008 is therefore comparable in magnitude to the two shocks of the 1970s. Given the global crisis that followed, has the macroeconomic effect of oil prices returned with a vengeance?

Probably not. The best, and best known, attempt made to pin the recession on oil prices is that of Hamilton (2009). Hamilton finds an oil-price-induced slowdown in consumer purchases, and auto purchases in particular, in 2007 and early 2008. This appears to have reduced US growth up to mid-2008. Hamilton concludes that 'if there had been no oil shock, we would have described the US economy in 2007:Q4-2008:Q3 as growing slowly, but not in a recession'. This is possible, but it is an extremely weak conclusion: the difference between 'growing slowly' and 'recession' is rather small. Also, he fails to control for monetary policy, so implicitly his estimate allows some output loss due to the monetary tightening that historically tends to accompany high oil prices.

Most importantly, the period he analyses was a mild recession during which output declined at an annual rate of 0.7 per cent. The massive drop in output that occurred as part of the global crisis took place from 2008:Q3 to 2009:Q1, when US GDP declined at an annual rate of 5 per cent. It is hard to argue that oil prices had much to do with this much more severe downturn.[5]

(i) Monetary policy and inflation

How does the role of monetary policy help to explain the decline in the impact of oil prices? Hooker (2002) finds that oil prices stopped feeding through to core inflation some time in the early 1980s, and that the decline in the energy intensity of the economy, and deregulation of prices, are not able to explain the break. Nordhaus (2007) confirms that over 2002–6 there was no pass-through of energy prices to other prices and wages. One explanation might be that more reactive monetary policy nipped any incipient wage–price spiral in the bud, avoiding the need for stricter monetary tightening later on. However, Hooker also finds that monetary policy became less rather than more strict in its response to oil price shocks. Tighter

[5] Segal (2007) wrote: 'As of October 2007 there is some doubt that the rapid growth of the last five years will continue, but if there is a slow down then it will probably be due to credit markets and global imbalances.' This seems consistent with what followed.

monetary policy, therefore, cannot explain the decline in pass-through. His suggestion, instead, is that the causality may in fact be the reverse: the reduced pass-through of oil prices to core inflation may explain the reduced monetary policy reaction in the later period.

The monetary policy component of the explanation for the decline in the impact of oil prices on output, therefore, seems to rest on the fact that oil prices have less of an impact on core inflation than in the past, obviating the need for monetary tightening. The decline in the second-round pass-through means that workers and firms simply accept the decline in real income caused by the rise in the price of petroleum products. In terms of the Taylor rule described earlier, the inflation effect is absent and interest rates do not need to rise. Indeed, the contractionary effect of the high oil price may be dominant, in which case the high oil prices up to 2008 may have contributed to the historically low interest rates seen over the same period (Allsopp, 2006).

The key question then becomes why oil prices no longer pass through to core inflation. This matter requires further research, but one can point to some plausible hypotheses. Labour's bargaining power is probably weaker today than in the 1970s, owing to both weaker unions and international capital mobility.[6] This reduces labour's ability to demand wage catch-up, consistent with Nordhaus's (2007) finding that over 2002–6 wages were unresponsive to price increases.

A second reason is 'globalization', or increased competition from imports. Firms sometimes prefer to maintain market share at lower margins rather than lose sales to imports such as the low-priced imports from Asia, and China in particular. Firms may therefore absorb losses rather than raise prices. A third reason affecting both workers and firms may be that monetary policy is now more credible.[7] Thus agents in the economy know that any wage–price spiral would be crushed by the interest rate response, and therefore do not make wage or price demands in response to a rise in the price of oil. Hooker (2002) refers to Taylor's (2000) argument that low levels of inflation and inflation persistence lead to lower pricing power on the part of firms, and therefore lower pass-through of increased costs (of all kinds) to prices. However, Taylor's argument applies only to temporary rises in costs, and therefore cannot explain the lack of impact of the current sustained oil price rise.

(ii) Supply shocks and demand shocks

One popular interpretation of the difference between the rise in oil prices up to 2008 and the oil shocks of the 1970s is that the former was 'demand driven', while the latter were 'supply driven'. This is argued, for instance, in the IMF's *World Economic Outlook* (WEO) (IMF, 2007) in a box entitled 'Understanding the Link Between Oil Prices and the World Economy'.

Their box compares the impact on the world economy of a supply-driven oil price shock with the impact of a demand-driven oil price shock using the IMF's GEM. In the model, 'monetary policy is specified in terms of a credible commitment to an interest rate rule that targets inflation',[8] so the pass-through of oil prices to core inflation induces a rise in interest

[6] This bargaining argument is made, for instance, by Rodrik (1997).

[7] As discussed in Walton (2006), for instance.

[8] Reported in Elekdag *et al.* (2007, p. 6), in which the authors describe the model used for the IMF estimations. Specifically (p. 9), 'In order to hit a target of 2 per cent (2.5 per cent in the United States) for core consumer price inflation four to six quarters in the future, the change in the interest rate must be twice as large as the deviation of core inflation from its target level', where core inflation excludes gasoline prices.

rates. The supply-driven oil price shock is modelled as a restriction in supply that leads to a rise in the oil price of 100 per cent at its peak. This very standard exercise results in a drop in world GDP of 1.4 per cent at its trough, a little over 1 year after the shock, while world inflation rises by 1.5 per cent at its peak, after two quarters.

The second simulation in the box describes the impact of an oil price shock caused by what the WEO describes as a 'demand shock'. This 'demand shock' is modelled as 'a significant increase in productivity growth in oil-importing countries that permanently raises global growth by ½ of a percentage point [which] generates a significant short-run surge in oil prices that is sustained over the medium term . . . [reflecting] the low short-term elasticity of supply' (pp. 18–19). So what impact does this 'demand shock'-driven oil price shock have on world GDP? The WEO reports that 'the short-run path for world GDP is opposite to that resulting from a supply-induced increase in the price of oil because higher prices are being caused by stronger growth' (p. 19).

This analysis is highly problematic. I place scare quotes around 'demand shock' because the shock that they model is a demand shock only from the point of view of the oil market: from the point of view of every other market, a rise in productivity growth is a positive supply shock. But worse than this, the construction of the shock tells us nothing about the impact of oil prices. Why? Because the rise in productivity growth has apparently been *calibrated* precisely to raise global growth by ½ a percentage point per year within the model. The WEO's Figure 1 in the box duly shows that under the 'demand shock', world GDP rises above baseline at a rate of ½ a percentage point a year. To observe then that the 'demand shock'-induced oil price shock did not prevent world GDP growth from rising by ½ a percentage point a year is simply to observe that the calibrators successfully performed the proposed calibration.

The calibrations presented in the main text of the WEO therefore tell us nothing about the differential impacts of different kinds of oil price shocks. The box reports a more meaningful exercise, however, in footnote 4 (p. 19) where it states that 'if the same increase in productivity is considered in a version of the model that does not include oil, world GDP expands by slightly more in the short and medium term than in the model with oil'. This is an appropriate comparison as it tells us whether the rise in oil prices affects GDP growth for a given rate of productivity growth. The fact that world GDP expands by only 'slightly more' without oil leads to the conclusion in the footnote that 'this suggests that while high oil prices have resulted in a drag on world growth, these effects are relatively minor'.

While this comparison is meaningful, it is not comparable to the exercise with an oil-supply-induced oil price shock, because the two different shocks involve very different sizes of increase in the price of oil. While the oil supply shock is calibrated to induce a peak rise of 100 per cent in the price of oil, the 'demand shock' induces a rise in the price of oil which spikes at only 35 per cent in the first year. Since the induced oil price shocks are of such different magnitudes, the exercises reported in the box tell us nothing about the difference made by the source of the oil price shock.

A more sophisticated analysis is provided by Kilian (2009), who decomposes oil price movements into three components: changes in oil supply (under the assumption that supply is price-inelastic in the short run); changes in aggregate global demand (measured using an index of dry-cargo single-voyage freight rates); and changes in oil-specific demand. The last includes 'shifts in the price of oil driven by higher precautionary demand associated with fears about the availability of future oil supplies' (pp. 1–2).

Using a structured VAR for the global economy, he finds that oil supply shocks have a much smaller impact on oil prices than the other two types of shock. Using a similar model for the US specifically, he finds that both oil supply disruptions and oil market-specific demand shocks significantly lower US GDP growth. Aggregate demand-driven oil price increases are only partially significant at the 10 per cent level. On the basis of the estimated impacts of the shocks on both output and inflation, he states that 'the risk of stagflationary responses depends very much on the origin of the oil price increase and is much more pronounced for oil demand shocks than for oil supply shocks' (p. 22).

This analysis is informative about the impact of the three types of shock as Kilian defines them. But each is defined as a one-standard-deviation structural innovation, and each therefore has a different effect on the oil price (as shown in Kilian's Figure 5). The oil-specific demand shock, in particular, has a much larger impact on the price of oil in the first 2 years than the other shocks. Thus it is not surprising that it has a stronger impact on other macroeconomic variables. Like the IMF exercise, it fails to show that the marginal impact of a given oil price rise differs according to the cause of the rise.

Beyond these problems with specific analyses, the opposition of supply- and demand-induced oil price shocks has conceptual problems as well. Put simply, from the point of view of any individual country, increased world demand for oil is equivalent to reduced supply. For the US, a rise in oil consumption by China of 500,000 barrels of oil per day is exactly equivalent to a decline in global supply of the same quantity. The only difference is that, in the former case, the rise in demand for oil in China is probably accompanied by a rise in demand and supply of other goods from China. The economic environment will, therefore, be different, so economic outcomes will be different, but this is no reason to claim that the marginal impact of the oil price rise itself will be different.

VI. Conclusion

The literature on the role of oil prices in the macroeconomy falls into two broad camps. One analyses the effect of a rise in the oil price from a microeconomic point of view, explaining the link through various market frictions, non-competitive behaviour, and the complementarity between energy and capital. The other takes a more macroeconomic point of view, analysing the impact of oil prices on aggregate demand through inflation, and the corresponding impact of monetary policy responses. I argued that the microeconomic mechanisms, while plausible in their own right, probably were not able to explain the magnitude of the estimated impact of the oil price on the macroeconomy through the 1970s and early 1980s, and that monetary policy was the more promising explanation.

We also saw, however, that high oil prices have had little impact on the macroeconomy over the last few years, and that in the data oil price rises already stopped having an impact some time in the 1980s. At the same time, oil price rises stopped passing through to inflation, and this may hold the explanation: if oil price rises do not raise prices, then interest rates do not need to respond to them, and the impact on aggregate activity may therefore be minimal. Some hypotheses regarding why oil prices have less pass-through to core inflation were suggested, but more research on this area would seem to be required. The view that the effect of oil prices is lower when they are 'demand driven' as opposed to 'supply driven' is implausible a priori, and is not in fact supported by the evidence adduced in its favour.

The fundamental point is that under most circumstances monetary policy-makers have wide discretion in choosing the path of output. The binding constraint is not on output, but rather the impact on inflation, i.e. the Philips curve. If an oil price spike looked set to reduce output significantly, the monetary authority could choose to loosen monetary policy to offset it. The only reason not to do so is a fear of inflation, which oil prices may exacerbate if there are second-round price effects. The oil price has never been more than one of many significant economic variables, and since the disappearance of these second-round effects from the 1980s, its independent influence on macroeconomic cycles has been lower still.

References

Allsopp, C. J. (2006), 'Why is the Macroeconomic Impact of Oil Prices Different this Time?', *Oxford Energy Forum*, Issue 66, August.

Balke, N. S., Brown, S. P. A., and Yücel, M. K. (2002), 'Oil Price Shocks and the US Economy: Where Does the Asymmetry Originate?', *The Energy Journal*, **23**(3), 27–52.

Barsky, R. B., and Kilian, L. (2004), 'Oil and the Macroeconomy since the 1970s', *Journal of Economic Perspectives*, **18**(4), 115–34.

Bernanke, B. S., Gertler, M., and Watson, M. (1997), 'Systematic Monetary Policy and the Effects of Oil Price Shocks', *Brookings Papers on Economic Activity*, **1**, 91–142.

— — — (2004), 'Reply: Oil Shocks and Aggregate Macroeconomic Behaviour: The Role of Monetary Policy', *Journal of Money, Credit and Banking*, **36**(2), 287–91.

Blanchard, O. (2002), *'Comment', NBER Macroeconomics Annual 2001*, Cambridge, MA, MIT Press, 183–92.

— Galí, J. (2007), 'The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?', NBER Working Paper No. 13368, September.

Blinder, A. S., and Rudd, J. B. (2008), 'The Supply-shock Explanation of the Great Stagflation', NBER Working Paper 14563.

Bohi, D. (1991), 'On the Macroeconomic Effects of Energy Price Shocks', *Resources and Energy*, **13**(2), 145–62.

BP (2007), *Statistical Review of World Energy 2007*, downloaded from www.bp.com

Burbidge, J., and Harrison, A. (1984), 'Testing for the Effects of Oil-Price Rises using Vector Autoregressions', *International Economic Review*, **25**(2), 459–84.

Darrat, A. F., Gilley, O. W., and Meyer, D. J. (1996), 'US Oil Consumption, Oil Prices, and the Macroeconomy', *Empirical Economics*, **21**, 317–34.

Davis, S. J., and Haltiwanger, J. (2001), 'Sectoral Job Creation and Destruction Responses to Oil Price Changes', *Journal of Monetary Economics*, **48**, 465–512.

Elekdag, S., Lalonde, R., Laxton, D., Muir, D., and Pesenti, P. (2007), 'Oil Price Movements and the Global Economy: A Model-based Assessment', Bank of Canada, Working Paper 2007–34.

Finn, M. G. (2000), 'Perfect Competition and the Effects of Energy Price Increases on Economic Activity', *Journal of Money, Credit and Banking*, **32**(3), 400–16.

Hamilton, J. D. (1983), 'Oil and the Macroeconomy since World War II', *Journal of Political Economy*, **91**(2), 228–48.

— (1988), 'A Neoclassical Model of Unemployment and the Business Cycle', *Journal of Political Economy*, **96**(3), 593–617.

— (1996), 'This is What Happened to the Oil Price–Macroeconomy Relationship', *Journal of Monetary Economics*, **38**, 215–20.

— (2005), 'Oil and the Macroeconomy', Department of Economics, University of California, San Diego, mimeo, prepared for *Palgrave Dictionary of Economics*.

— (2009), 'Causes and Consequences of the Oil Shock of 2007–08', NBER Working Paper 15002.

— Herrera, A. M. (2004), 'Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy', *Journal of Money, Credit, and Banking*, **36**(2), 265–86.

Hooker, M. A. (1996*a*), 'What Happened to the Oil Price–Macroeconomy Relationship?', *Journal of Monetary Economics*, **38**, 195–213, with reply by Hamilton and final response from Hooker.
— (1996*b*), 'This is What Happened to the Oil Price–Macroeconomy Relationship: Reply', *Journal of Monetary Economics*, **38**, 221–2.
— (1999), 'Oil and the Macroeconomy Revisited', Working Paper, Federal Reserve Board, August.
— (2002), 'Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime', *Journal of Money, Credit and Banking*, **34**(2), 540–61.
Hunt, B. (2005), 'Oil Price Shocks: Can They Account for the Stagflation in the 1970s?', IMF Working Paper WP/05/215, Washington, DC, International Monetary Fund.
— Isard, P., and Laxton, D. (2001), 'The Macroeconomic Effects of Higher Oil Prices', IMF Working Paper WP/01/04.
IMF (2000), 'The Impact of Higher Oil Prices on the Global Economy', paper prepared by the IMF Research Department, 8 December, Washington, DC, International Monetary Fund.
— (2007), *World Economic Outlook,* April 2007, Washington, DC, International Monetary Fund.
Jones, D. W., Leiby, P. N., and Paik, I. K. (2004), 'Oil Price Shocks and the Macroeconomy: What Has Been Learned Since 1996', *The Energy Journal*, **25**(2), 1–32.
Kilian, L. (2006), 'A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries', Department of Economics, University of Michigan and CEPR, mimeo, available at http://www-personal.umich.edu/~lkilian/kilianjeearev2.pdf
— (2009), 'Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market', *American Economic Review*, **99**(3).
Leduc, S., and Sill, K. (2004), 'A Quantitative Analysis of Oil-price Shocks, Systematic Monetary Policy, and Economic Downturns', *Journal of Monetary Economics*, **51**, 781–808.
Lee, K., Ni, S., and Ratti, R. (1995), 'Oil Shocks and the Macroeconomy: The Role of Price Variability', *The Energy Journal*, **16**(4), 39–56.
Mabro, R. (1998), 'OPEC Behaviour 1960–1998: A Review of the Literature', *Journal of Energy Literature*, **4**(1), 3–27.
Mork, K. A. (1989), 'Oil and the Macroeconomy When Prices go Up and Down: An Extension of Hamilton's Results', *Journal of Political Economy*, **97**(3), 740–4.
— (1994), 'Business Cycles and the Oil Market', *The Energy Journal*, **15**, Special Issue, 15–38.
— Mysen, H. T., and Olsen, O. (1994), 'Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries', *The Energy Journal*, **15**(4), 19–35.
Nordhaus, W. D. (2007), 'Who's Afraid of a Big Bad Oil Shock?', *Brookings Papers on Economic Activity*, **2**.
Orphanides, A. (2001), 'Monetary Policy Rules Based on Real-time Data', *American Economic Review*, **91**, 964–85.
Rodrik, D. (1997), *Has Globalization Gone Too Far?* Institute for International Economics, Washington, DC.
Rotemberg, J. J., and Woodford, M. (1996), 'Imperfect Competition and the Effects of Energy Price Increases on Economic Activity', *Journal of Money, Credit and Banking*, **28**(4), Part 1, 549–77.
Segal, P. (2007), 'Why Do Oil Price Shocks No Longer Shock?', Oxford Institute for Energy Studies Working Paper WPM 35.
Solow, R. M. (1980), 'What to Do (Macroeconomically) When OPEC Comes', in S. Fischer (ed.), *Rational Expectations and Economic Policy*, Chicago, IL, University of Chicago Press, 249–67.
Taylor, J. B. (2000), 'Low Inflation, Pass-through, and the Pricing Power of Firms', *European Economic Review*, **44**(7), 1389–408.
Walton, D. (2006), 'Has Oil Lost the Capacity to Shock?', *Bank of England Quarterly Bulletin*, Spring, 105–14.