Oil Price Shocks: Causes and Consequences

Lutz Kilian

Department of Economics, University of Michigan, Ann Arbor, Michigan 48109-1220; email: lkilian@umich.edu

Annu. Rev. Resour. Econ. 2014. 6:133–54

The *Annual Review of Resource Economics* is online at resource.annualreviews.org

This article's doi:
10.1146/annurev-resource-083013-114701

JEL code: Q43

Keywords

macroeconomy, demand, supply, channels of transmission, business cycle, asymmetries

Abstract

Research on oil markets conducted during the last decade has challenged long-held beliefs about the causes and consequences of oil price shocks. As the empirical and theoretical models used by economists have evolved, so has our understanding of the determinants of oil price shocks and of the interaction between oil markets and the global economy. Some of the key insights are that the real price of oil is endogenous with respect to economic fundamentals and that oil price shocks do not occur *ceteris paribus*. As a result, one must explicitly account for the demand and supply shocks underlying oil price shocks when studying their transmission to the domestic economy. Disentangling cause and effect in the relationship between oil prices and the economy requires structural models of the global economy including the oil market.

1. INTRODUCTION

There has been interest in understanding the causes and consequences of oil price shocks ever since the 1970s. Oil price shocks have been blamed for US recessions, for higher inflation, for a slowdown in US productivity in the 1970s, and for stagflation (a term coined to refer to the unprecedented coincidence of inflation and economic stagnation during the 1970s). They have also been held responsible for changes in monetary policy, for far-reaching labor market adjustments, and for changes in energy technologies. Although the interest in oil price shocks waned in the 1990s, the fluctuations in the real price of oil since 2003 have led to a resurgence of research on oil markets.

The research conducted during this last decade has challenged long-held beliefs about the causes and consequences of oil price shocks. As the empirical and theoretical models used by economists have evolved, so has our understanding of the determinants of oil price shocks and of the interaction between oil markets and the global economy. For example, whereas traditionally the real price of oil was thought to be determined primarily by political events in the Middle East that were outside of the confines of macroeconomic models and could simply be taken as given when conducting policy analysis, it is now widely accepted that the real price of oil is determined endogenously in global markets much like the price of other global commodities. This means that oil price fluctuations can be understood only with the help of structural models of the global oil market that allow feedback from US and global macroeconomic aggregates to the price of oil. In other words, understanding the effects of oil price shocks on the domestic economy requires global macroeconomic models that explain jointly the determination of the price of oil and of macroeconomic aggregates.

Moreover, whereas in the past oil price shocks were thought to cause recessions for reasons unrelated to the state of the domestic economy, we are now aware that sustained increases in the real price of oil are often merely symptoms of a booming world economy. This phenomenon became self-evident after 2005 but is by no means new. For example, the increases in the real price of oil in the early and late 1970s, although distinct in some dimensions, shared many of the features of the surge in the real price of oil in 2003–2008. It has become clear that economists for many years have tended to conflate the recessionary effects of oil price shocks with other causes of those earlier recessions. This fact does not mean that oil markets are not important for understanding the US business cycle. It means that the relationship between the real price of oil and US real economic activity is far more complex than suggested by earlier models. Unexpected oil price increases may sometimes be associated with strongly recessionary effects, yet may coexist with strong domestic economic growth at other times. The literature has made great strides over the last decade in describing and quantifying this relationship. This survey summarizes these developments.

Section 2 reviews the leading explanations of oil price fluctuations. Section 3 discusses economic theories of the transmission of oil price shocks, carefully distinguishing between the direct effects of an exogenous oil price shock and the indirect effects that may give rise to asymmetric responses of the economy, depending on whether the oil price shock is positive or negative. Section 4 reviews the empirical evidence on the transmission of unanticipated oil price changes under the assumption that these changes are predetermined with respect to the US domestic economy. Section 5 discusses how the analysis changes once we allow for the fact that oil prices are endogenously determined in global markets and reflect a variety of different oil demand and oil supply shocks. The conclusions are in Section 6.

2. ALTERNATIVE EXPLANATIONS OF OIL PRICE FLUCTUATIONS

Figure 1 plots the price of West Texas Intermediate (WTI) crude oil for 1948–1973. This price is nearly identical to the US producer price index for crude oil. During this period, there was no global market for crude oil. In fact, the United States was self-sufficient in crude oil. The price series in

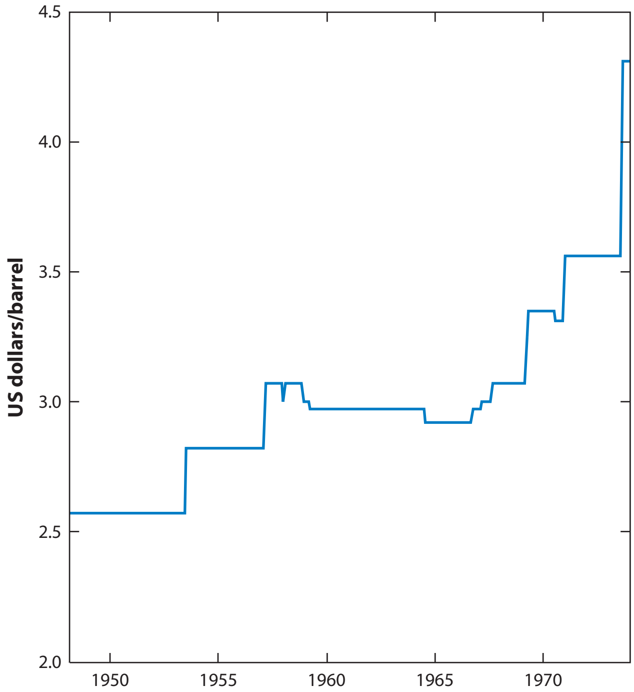

Figure 1

Price of West Texas Intermediate crude oil, 1948.1–1973.12.

Figure 1 differs from typical commodity prices in that the oil price often remained unchanged for extended periods, followed by occasional large, discrete adjustments. This pattern reflects the fact that the price of oil in the United States was regulated by the government. The regulation of the price of crude oil dates back to the 1920s, when, following the Texas oil boom, government agencies such as the Texas Railroad Commission were charged with regulating the oil industry. As discussed in Hamilton (1983), they did so by forecasting US oil consumption and by setting production targets accordingly, thus ensuring stable oil prices as a rule. Regulators were unable or unwilling to accommodate occasional unanticipated oil supply disruptions, however, preferring to exploit these events for sudden large oil price increases.

This system came to an end in 1971, when the United States ceased to be a net exporter of crude oil. By late 1973, the United States, like most other industrialized economies, was heavily dependent on imported crude oil from the Middle East. The market for crude oil had become global. This fact made it impractical for US government agencies to continue to regulate the price of crude oil in the United States. Although the price of domestically produced crude oil remained regulated to some extent as late as the early 1980s, regulators allowed the price to gradually converge toward the global price of crude oil. The transition toward a global market for crude oil is illustrated in **Figure 2**. The left panel of **Figure 2** shows the real (or inflation-adjusted) percent rate of growth in the WTI price. The low volatility of the series is punctuated by occasional spikes, reflecting discrete adjustments of the nominal price. Whereas there was only one monthly oil price series for all practical purposes prior to 1974, starting in 1974 the US Energy Information Administration began to collect a number of additional oil price series, reflecting the oil price paid by US refiners acquiring crude oil. The right panel of **Figure 2** plots the percent change in the real US refiners'

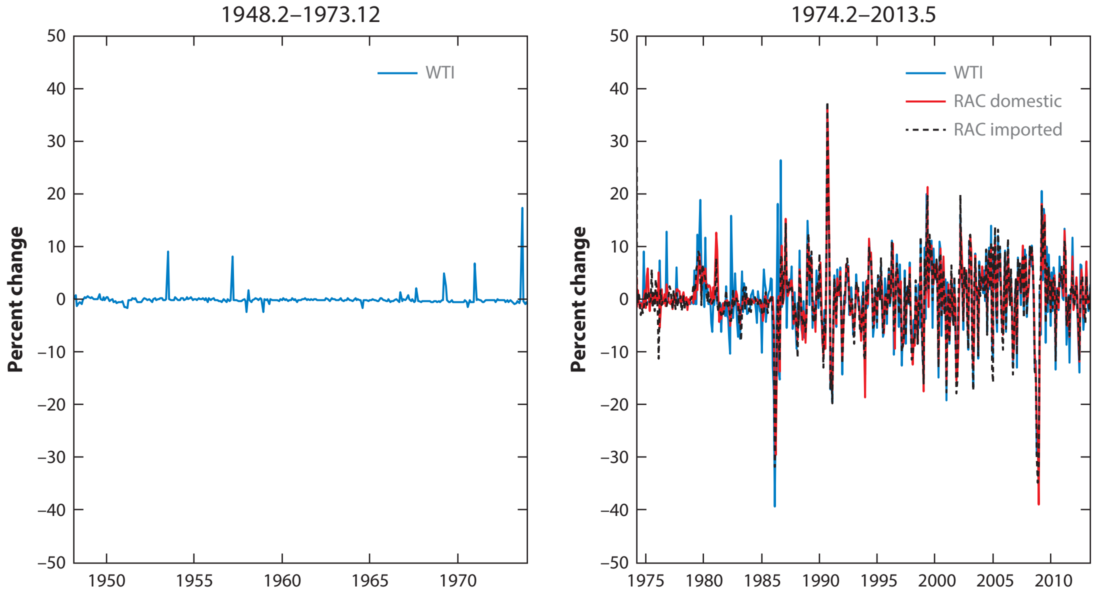

Figure 2

Percent change in the real price of oil. RAC denotes the US Refiners' Acquisition Cost as reported by the US Energy Information Administration, and WTI denotes West Texas Intermediate.

acquisition cost for imported crude oil and for domestically produced crude oil together with the WTI series. The figure makes immediately apparent that the time series process governing the real price of oil changed in late 1973. Starting in 1974, the volatility of the real price of oil began to resemble the volatility of other commodity prices that are determined in global markets.

Whereas the case can be made that the oil price spikes in the left panel of **Figure 2** are exogenous with respect to the global economy (which means that global macroeconomic aggregates cannot explain these price spikes), since 1974 the real price of oil, like the real price of other industrial commodities, has been endogenous with respect to global macroeconomic conditions (Alquist et al. 2013). On the one hand, changes in the real price of oil (all else equal) affect the macroeconomic performance of oil-importing economies such as the United States; on the other hand, changes in macroeconomic conditions in oil-importing economies (all else equal) affect the real price of oil. Thus, identifying cause and effect in the relationship between the real price of oil and macroeconomic conditions in oil-importing economies requires a structural model of the global economy including the global oil market.

2.1. Flow Supply Shocks

Hamilton (2003) suggests that this identification problem may be ignored for all practical purposes to the extent that all major fluctuations in the price of oil can be attributed to disruptions to the flow of oil production triggered by political events in the Middle East that are exogenous with respect to the US economy. Potential examples of such political events include the 1973 Yom Kippur War followed by the Arab oil embargo of 1973–1974, the Iranian Revolution of 1978–1979, the Iran-Iraq War of 1980–1988, the Persian Gulf War of 1990–1991, the Venezuelan crisis of 2002 and the Iraq War of 2003, and the Libyan uprising of 2011. Exogeneity with respect to the US economy here means that these events did not occur in response to the current or past state of the US economy.

There are three problems with this explanation. First, it frequently does not fit the data. For example, although the Arab-Israeli War of 1973 was clearly exogenous with respect to the US economy, it did not constitute a shock to the flow of crude oil supplies, because the war was not fought on the territory of oil-producing economies and no oil production facilities were damaged. The Arab oil embargo of 1973–1974, in contrast, did affect the flow of oil supplies, but the embargo decision was taken explicitly with reference to the state of the US economy (Kilian 2008b). In the case of the Iranian Revolution, the timing is off in that the oil price surged starting only in May 1979 after the exogenous oil supply disruption in Iran was over. Finally, following the flow supply disruptions associated with the outbreak of the Iran-Iraq War in late 1980, the Venezuelan crisis of late 2002, and the Iraq War of early 2003, the price of crude oil hardly increased. Second, more formal regression analysis confirms that quantitative measures of exogenous oil supply shocks associated with political events in the Middle East invariably have little predictive power for the percent change in the real price of oil (Kilian 2008a,b). Third, numerous subsequent empirical studies show that most major oil price increases since late 1973 have had an important endogenous component associated with the global business cycle.

A case in point is the 1973–1974 episode, which constituted perhaps the most dramatic surge in the real price of crude oil in modern history. The nominal price of oil quadrupled over the course of half a year. Measures of exogenous oil supply shocks as proposed in Hamilton (2003) and Kilian (2008b), however, explain at most 25% of the observed oil price increase in 1973–1974. This raises the question of what explains the remaining oil price increase. The answer is that, by construction, at least 75% of that oil price increase must be attributed to shifts in the demand for oil. This answer is quite different from that proposed in Hamilton (2003) but is supported by independent evidence. As observed in Barsky & Kilian (2002), there was a global demand boom in the early 1970s in all industrial commodity markets across the board, reflecting the fact that, for the first time in postwar history, there was a simultaneous peak in the business cycle in the United States, in Europe, and in Japan. Between November 1971 and February 1974, the price of industrial raw materials and metals increased in real terms by approximately 95%, whereas the real price of oil increased by 125%. Thus, real nonoil industrial commodity prices in the absence of supply shocks cumulatively increased by approximately 75% as much as the real price of oil, which corroborates the 75% increase attributed to oil demand shocks according to the direct evidence based on exogenous oil supply shocks. Moreover, because much of the nonoil commodity price increase occurred while institutional constraints kept the price of crude oil low, we can be sure that the commodity price increases of the early 1970s were not driven by higher oil prices.

2.2. Flow Demand Shocks

There is evidence that flow demand shocks associated with the global business cycle were a primary determinant not only of the 1973–1974 oil price increase, but of most major oil price increases. Flow demand refers to the demand for oil to be consumed immediately in the process of producing refined products such as gasoline, diesel, heating oil, kerosene, and jet fuel. As the global economy expands, flow demand for oil increases because oil is a necessary ingredient for the modern economy. Hence, it is not surprising that the real price of oil, all else equal, depends on the state of the global economy. Nevertheless, the role of flow demand for the real price of oil was not appreciated for a long time. Flow demand was first identified as a major determinant of the real price of oil by Barsky & Kilian (2002). Barsky & Kilian rely on indirect evidence, including the striking comovement in oil and other commodity prices, to show that the major oil price fluctuations in the 1970s and early 1980s appear to be associated in large part with fluctuations in the global business cycle.

Only since Kilian (2009b) have researchers been able to directly quantify the importance of flow demand shocks for the evolution of the real price of oil. This study proposes a structural vector autoregressive (VAR) model of the global market for crude oil since 1973. This model allows one to decompose the evolution of the real price of oil into distinct components associated with demand and supply shocks. Kilian (2009b) shows that most large and persistent fluctuations in the real price of oil since the 1970s have been associated with the cumulative effects of oil demand rather than oil supply shocks, complementing the results in Kilian (2008a,b) on the limited role of exogenous oil supply shocks. The methodological approach introduced in Kilian (2009b) has been refined in recent years, but the substance of the findings has remained intact. For example, Kilian & Hicks (2013) show the pattern of flow demand shocks during 2003–2008 to be consistent with the pattern of forecast errors made by professional real-GDP forecasters. Kilian & Murphy (2012) examine the robustness of Kilian's (2009b) findings to alternative identifying assumptions based on the signs of responses to shocks. Baumeister & Peersman (2013a) in addition allow for time-varying parameters in a similar VAR model. Kilian & Murphy (2014) and Kilian & Lee (2014) refine earlier structural oil market models by allowing for an explicit role for speculation in oil markets, again confirming the substance of the earlier results.

These and other studies provide overwhelming evidence that oil demand shocks collectively explain most major oil price fluctuations since 1973, with a central role played by flow demand shocks. Only for the 1990 episode is there evidence that flow supply shocks played a nonnegligible role (Kilian & Murphy 2014). Whereas the idea that flow demand shocks can explain major oil price fluctuations was met with great skepticism when it was first proposed by Barsky & Kilian (2002), this view is the conventional view today.

2.3. The Role of Expectations in the Physical Market for Crude Oil

The fact that flow supply disruptions have historically had small effects on the real price of oil does not mean that political events in the Middle East do not matter, because such events may also affect the real price of oil by shifting expectations about future shortfalls of oil supply relative to oil demand. This point has received increasing recognition only recently (e.g., Kilian 2009b, Alquist & Kilian 2010, Dvir & Rogoff 2010, Knittel & Pindyck 2013, Kilian & Lee 2014). Whereas shifts in expectations about the future scarcity of oil are not observable, their implications for inventory holdings are. The central idea is that anyone expecting the real price of oil to increase in the future has an incentive to store oil for future use, which in turn provides incentives to curb current oil consumption and stimulates additional oil production. Thus, expectation shifts, rather than shifting the flow demand for oil, represent shocks to the demand for oil stocks (or oil inventories). Kilian & Murphy (2014) show that the implications of unobservable shifts in expectations for the real price of oil may be inferred from fully identified structural VAR oil market models that include data on changes in crude oil inventories held aboveground.

It is worth emphasizing that expectation-driven shifts in oil inventory demand in the physical market for crude oil may equivalently be labeled speculative demand shocks. A speculator in the physical market for crude oil is someone who buys crude oil not for current consumption, but for future use. What typically triggers speculative purchases of crude oil is an expectation of rising oil prices, which is the same as expecting a shortfall of oil supply relative to oil demand. This fact allows one to assess the presence of speculative pressures in the physical market for oil by using structural models designed to accommodate unobserved shifts in expectations. Kilian & Murphy (2014) demonstrate that shifts in expectations (or, equivalently, speculative demand pressures) played an important role during many episodes of interest, including, for example, the oil price surge in the second half of 1979, the oil price collapse of 1986, and the oil price spike of 1990. They do not find any support for speculative demand shocks explaining the surge in the real price of oil between 2003 and mid-2008, however.

2.4. Financial Speculation in Oil Futures Markets

There has been much debate in recent years about the possibility that trades of financial investors in oil futures markets may have caused the surge in the real price of oil between 2003 and mid-2008. The oil futures market is a public exchange for oil futures contracts that was created in the 1980s. Futures contracts allow investors to agree in advance on the price at which they will exchange a given quantity of crude oil at some future date. This "paper market" for oil allows participants to speculate on the future price of oil without having to store oil physically. It thus provides an alternative venue for trading on the basis of expectations. Many of the traders on the exchange have some physical exposure to oil price risks that they wish to hedge. Others may simply wish to make money in exchange for assuming risk. Since approximately 2003, the fraction of the latter traders has increased dramatically. This fact recently gave rise to the perception that the positions taken by financial investors after 2003 exogenously created speculative pressures on oil futures prices that then spilled over to the physical market for crude oil.

If this perception were correct, one would expect a corresponding shift in inventory demand driven by speculative motives after 2003, as the physical market and the futures market are linked by an arbitrage condition (e.g., Alquist & Kilian 2010). In fact, this arbitrage is essential for explaining how an exogenous increase in oil futures prices may raise the price of oil in the physical market. Thus, the hypothesis that financial traders caused the surge in the real price of oil after 2003 can be tested by using the same model that sheds light on the presence of speculative demand shocks in the physical market for oil. Kilian & Murphy (2014) and Kilian & Lee (2014) demonstrate that there is no evidence of such speculative demand pressures in the physical oil market after 2003. This result, under the maintained assumption of arbitrage, also rules out speculative pressures in the oil futures market. This conclusion is consistent with a large body of literature on the role of speculation (e.g., Fattouh & Mahadeva 2013, Fattouh et al. 2013, Knittel & Pindyck 2013).

One unresolved question in this literature is the extent to which the presence of financial investors in oil futures markets may have affected the risk premium in oil futures markets (defined as the financial reward for assuming risk), which has traditionally been treated as constant over time. Although this possibility has been discussed at some length among financial economists in recent years, few studies quantify the evolution of the time-varying risk premium in a satisfactory manner. On the basis of conventional tests in the tradition of Fama & French (1987) applied to data ending in 2006, Alquist & Kilian (2010) conclude that there is no evidence of a time-varying risk premium in oil futures markets. More recent evidence in Hamilton & Wu (2014) based on a term structure model suggests that the risk premium was indeed stable and close to zero from 1990 until 2005 but has been highly volatile about a mean close to zero since then. What this new evidence means is not yet clear. Understanding the implications of time variation in the risk premium for the real price of oil in physical markets requires new models of the link between physical and financial oil markets, such as the model recently developed by Fattouh & Mahadeva (2013). Fattouh & Mahadeva conclude that fluctuations in the risk premium cannot explain the evolution of the real price of oil between 2003 and mid-2008. This question remains an active area of research.

Alternative theoretical models of the effects of the financialization of oil futures markets, such as the model of Sockin & Xiong (2013), are based on the premise that consumers of oil mistake exogenous increases in the oil futures price as a predictor of future increases in real economic activity. This argument is difficult to reconcile with full-fledged structural oil market models. Moreover, recent work by Hu & Xiong (2013) shows that there is no empirical support for this premise. In related work, it has also been suggested that an exogenous increase in oil futures prices may be transmitted directly to the real price of oil in the physical oil market without affecting crude oil inventories, to the extent that refiners use oil futures prices as reference prices in writing contracts for the delivery of crude oil. This argument implies that oil futures prices should be helpful in predicting the US refiners' acquisition cost at the relevant horizons, but it can be shown that no such predictive relationship exists in the data, contradicting this argument (e.g., Alquist et al. 2013).

The central premise in much of the literature on financial speculation has been that the influx of financial investors after 2003 was exogenous with respect to the physical market for oil. This premise is highly questionable. Clearly, investors did not decide to enter the oil futures market for no good reason. Commodity index fund investors in particular must have been attracted to the oil market by the perception that recent high returns in commodity markets triggered by additional flow demand from emerging Asia would persist in the future. Hence, it makes more sense to think of their activities as part of the propagation of earlier flow demand shocks rather than as a new and independent shock. This is indeed the view underlying more recent theoretical models such as the model of Basak & Pavlova (2013). Such models imply that the presence of financial investors in the oil futures market amplified the response of the real price of oil to flow demand shocks, at least in 2007–2008. How quantitatively important these amplification mechanisms are remains an open question.

2.5. The Role of OPEC

A final potential explanation of oil price fluctuations is based on the premise that the Organization of Petroleum-Exporting Countries (OPEC) has been effectively controlling the price of oil since late 1973. The assertion is that OPEC is a cartel that controls the price of oil either directly or by coordinating its oil production decisions. The literature has not been kind to this explanation (e.g., Smith 2005). There is no evidence of OPEC ever having been able to raise the price of oil on its own, nor is there evidence that OPEC has been able to prevent the price of oil from falling to record lows in 1986 and in 1999.

Oil market historians have documented that OPEC behavior in the 1970s was chaotic, with no effective coordination among its members (Skeet 1988). Only in the early 1980s did OPEC attempt to act as a cartel in an attempt to prevent oil prices from falling. Given the lack of cooperation from other OPEC members, Saudi Arabia unilaterally reduced its oil production to compensate for higher oil production elsewhere in the world, including some OPEC member countries. This attempt was unsuccessful in that it slowed, but could not stop, the decline in the real price of oil. As both the price of crude oil and the quantity of oil exported by Saudi Arabia continued to fall, Saudi oil revenues shrank, and this policy proved unsustainable. In late 1985, the "cartel" collapsed, as predicted by procylical models of cartels (Green & Porter 1984). OPEC has made no serious attempt to control oil prices since then. This does not mean that the oil market has been competitive. Several studies model noncompetitive behavior in the oil market (e.g., Almoguera et al. 2011). One particularly interesting question in this context is how to model Saudi Arabia's traditional role in oil markets as a supplier of last resort (e.g., Nakov & Nuño 2013).

3. CONSEQUENCES OF EXOGENOUS OIL PRICE SHOCKS: ECONOMIC THEORY

Reflecting the preoccupation of many economists with exogenous OPEC supply shocks in oil markets, much of the early literature on the transmission of oil price shocks focuses on models with exogenous oil prices. Often the real price of oil is modeled as an exogenous autoregressive moving average (ARMA) process (e.g., Atkeson & Kehoe 1999, Leduc & Sill 2004). It is implicitly understood in these models that demand does not matter for the determination of the real price of oil and that all oil price shocks are the same and driven only by oil supply shocks. As discussed above, this traditional view has been overturned in the recent oil market literature. In this section, we nevertheless maintain the simplifying assumption of an exogenous real price of oil because it simplifies the exposition of the basic mechanisms of the transmission of oil price shocks. We abstract from other changes in the economy normally associated with an oil price shock driven by flow demand, as though oil price shocks were driven entirely by exogenous flow supply shocks. In Section 5, we discuss the fundamental changes in the analysis required to deal with the reality of fully endogenous oil prices.

3.1. The Direct Effects of an Exogenous Oil Price Shock on Real GDP

Most theoretical models of the transmission of oil price shocks focus on the implications of exogenous variation in the real price of imported crude oil. The transmission of such oil price shocks relies on two main channels. One immediate effect of an unexpected increase in the price of imported crude oil is a reduction in the purchasing power of domestic households as income is transferred abroad.[1] This first effect is akin to an adverse aggregate demand shock in a macroeconomic model of aggregate demand and aggregate supply. The other immediate effect is to increase the cost of producing domestic output to the extent that oil is a factor of production along with capital and labor, which is akin to an adverse aggregate supply shock. These direct effects of an exogenous increase in the real price of oil imports are symmetric in oil price increases and decreases. An unexpected increase in the real price of oil causes aggregate production and income to fall by as much as an unexpected decline in the real price of oil of the same magnitude causes aggregate income and production to increase.

3.1.1. Direct effects: supply channel of transmission. The intuition obtained from textbook models of aggregate demand and aggregate supply extends to modern dynamic stochastic general equilibrium (DSGE) models of oil-importing economies. The traditional approach in DSGE models of the transmission of exogenous oil price shocks has been to treat oil as an intermediate input in domestic production and to ignore the demand channel of transmission. Exogenous increases in the price of imported crude oil are viewed as terms-of-trade shocks (e.g., Kim & Loungani 1992). There are three well-known problems in explaining a decline in real GDP on the basis of this intermediate-input cost or supply channel. The first problem is that, if oil is an imported commodity, the interpretation of crude oil as an intermediate input in the value-added production function is questionable. Under standard assumptions, imported oil enters the production function of domestic gross output, but it does not enter the production function of domestic value added (e.g., Rotemberg & Woodford 1996). Because gross output is separable in value added and imported energy, holding capital and labor fixed, oil price shocks do not move value added. Hence, oil price shocks by definition cannot be interpreted as productivity shocks for real GDP (Barsky & Kilian 2004). The second problem is that, to the extent that oil prices affect domestic output, under standard assumptions their impact should be bounded by the cost share of

[1]It is important for the argument that it is the price of imported crude oil that increases, because an increase in the price of domestically produced crude oil by itself would merely cause a redistribution of income rather than a reduction in domestic income in the aggregate.

oil in GDP, which is very small. For example, Kilian & Vigfusson (2013) show that the share of domestically produced and imported crude oil in US GDP has fluctuated between 1% and 8% since 1973.

There are three proposals in the literature for dealing with these problems. All three involve major modifications of the baseline DSGE model of an oil-importing economy to generate quantitatively important effects of exogenous oil price shocks on domestic real GDP. The first proposal is due to Rotemberg & Woodford (1996) and relies on large and time-varying markups to generate large effects of oil price shocks on real GDP. The second proposal is Atkeson & Kehoe's (1999) putty-clay model, which appeals to capital-energy complementarities in production. The third proposal is due to Finn (2000) and relies on a perfectly competitive model, in which energy is essential to obtaining a service flow from capital. In all three models, the supply channel of the transmission of energy price shocks may be quantitatively important, yet there is no consensus as to which, if any, of these models has empirical support. Moreover, whether any of these models can account for a large share of business cycle fluctuations in US real GDP remains to be shown.

3.1.2. Direct effects: demand channel of transmission. In part in response to these challenges, another branch of the literature instead focuses on the reduction in the demand for goods and services triggered by oil price shocks. In this alternative view, the primary channel of transmission is on the demand side of the economy. For example, in a recent survey on the effects of energy price shocks, Hamilton (2009) stresses that a key mechanism whereby energy price shocks affect the economy is through a disruption in consumers' and firms' spending on goods and services other than energy. Increasingly, modern dynamic optimizing macroeconomic models allow for this possibility by incorporating the oil consumption of households (e.g., Dhawan & Jeske 2008, Bodenstein & Guerrieri 2011).

The central idea is that higher energy prices are expected to reduce households' discretionary income, as consumers have less money to spend after paying their energy bills. All else equal, the less elastic the demand for energy, the larger this discretionary income effect is, but even with perfectly inelastic energy demand, the magnitude of the effect of a unit change in energy prices is bounded by the energy share in consumption. The average share of all energy expenditures in US consumer expenditures is approximately 6%. Edelstein & Kilian (2009) conclude that the implied reduction in consumer expenditures is too small to account for large reductions in real GDP.

Hamilton (1988) shows that the reduction in expenditures may be further amplified by increases in the operating cost of energy-using durables. As the dollar value of such purchases may be large relative to the value of the energy they use, even relatively small changes in energy prices can have large effects on output and employment. Empirical evidence in Edelstein & Kilian (2007, 2009) confirms the presence of this effect, but only for automobile purchases. There is no such evidence for other energy-using durables such as refrigerators and other appliances. Given the small share of auto purchases in total expenditures, the overall effect on aggregate real GDP that can be explained in this way remains modest.

3.2. The Direct Effects of an Exogenous Oil Price Shock on Inflation

Recent empirical evidence suggests that the domestic supply channel of transmission is weak and that the domestic demand channel of transmission dominates in practice. For example, Kilian & Park (2009) provide evidence that stock returns of industries that depend on final consumer demand (such as retail sales and tourism) are more sensitive to exogenous oil price shocks than stock returns of industries that are intensive in the use of energy (such as the chemical industry). This view is also consistent with narrative evidence of how oil price shocks affect US industries (Lee

& Ni 2002) and with the views of many policy makers that an exogenous increase in energy prices slows economic growth primarily through its effects on consumer spending.

On this basis, one would expect an exogenous oil price shock, if it occurs in isolation, to be recessionary and deflationary. Of course, an oil price shock can shift both the aggregate demand and aggregate supply curves. Such a shift would reinforce the decline in real GDP but would result in at least partially offsetting effects of the oil price shock on the price level. Neither interpretation matches the common perception among macroeconomists that exogenous oil price shocks are recessionary and strongly inflationary. That perception is much harder to explain. Even if an exogenous oil price shock resulted in an adverse shift of the aggregate supply curve, causing the price level to increase, it would not be expected to cause sustained inflation in the absence of real-wage rigidities (Bruno & Sachs 1982). Although the possibility of such a wage-price spiral—in which unions insist on wage increases to offset price increases caused by oil price shocks—is often discussed in policy circles, there is no obvious support for this hypothesis, at least for the United States. In fact, quantifying the pass-through from shocks to the price of oil imports to domestic inflation is difficult, as discussed by Hamilton (2012). Indeed, this question is ill posed unless the import price increase is truly exogenous (which is unlikely in practice) and unless we are able to condition on the monetary policy response, which in turn affects the degree of the pass-through. Empirical evidence based on estimated monetary policy rules suggests that unpredictable changes in oil prices have historically been followed by a one-time adjustment in the price level, resulting in a blip in the inflation rate rather than in sustained inflation (e.g., Kilian & Lewis 2011). Kilian & Lewis (2011) also show that the cumulative effects of oil price shocks—after accounting for the estimated monetary policy response—fail to explain the overall evolution of US consumer price inflation, even after accounting for a possible structural break in the mid-1980s.

3.3. The Indirect Effects of an Exogenous Oil Price Shock on Real GDP

The difficulty of explaining large recessions on the basis of exogenous oil price shocks is a common problem in empirical work. Backus & Crucini (2000) illustrated this point within the context of a DSGE model. This type of result was very much at odds with prevailing views among many macroeconomists in the 1980s that exogenous oil price shocks were the primary cause of the stagflation of the 1970s, and prompted several researchers to explore additional amplification mechanisms. A central idea in this literature has been that the response of the economy is asymmetric in positive and negative oil price shocks such that positive oil price shocks generate large recessions, whereas negative oil price shocks have little, if any, effect on the economy. Indeed, such asymmetric models seemed to be the only models potentially capable of explaining how oil price shocks might cause large recessions.

3.3.1. Indirect effects: reallocation effect. The rationale for asymmetric responses of real output to oil price shocks hinges on the existence of additional indirect effects of unexpected changes in the real price of oil. There are mainly three such effects. First, it has been stressed that oil price shocks are relative price shocks that can be viewed as allocative disturbances that cause sectoral shifts throughout the economy (e.g., Hamilton 1988). For example, reduced expenditures on energy-intensive durables, such as automobiles, in response to unexpectedly high real oil prices tend to cause a reallocation of capital and labor away from the automobile sector. As the dollar value of such purchases may be large relative to the value of the energy they use, even relatively small changes in the relative price of oil can have potentially large effects on demand. A similar reallocation may occur within the automobile sector as consumers switch toward more energy-efficient vehicles (Bresnahan & Ramey 1993). If capital and labor are sector specific or product specific and cannot be moved easily to new uses, these intersectoral and intrasectoral reallocations will cause labor and capital to be unemployed, resulting in cutbacks in real output and employment that go beyond the direct effects of positive oil price shocks. For example, workers in the auto sector may be ill equipped to take different jobs short of extensive job retraining. The same effect may arise if unemployed workers simply choose to wait for conditions in their sector to improve.

The reallocation effect arises every time the relative price of oil changes unexpectedly, regardless of the direction of the oil price change. In the case of an unexpected real oil price increase, the reallocation effect reinforces the direct recessionary effect on domestic real GDP, allowing the model to generate a much larger recession than in standard linear models. In the case of an unexpected real oil price decline, the reallocation effect at least partially offsets the direct expansionary effect, causing a smaller economic expansion than implied by the standard model. Thus, in the presence of a reallocation effect, the responses of real output to oil are necessarily asymmetric in unanticipated oil price increases and unanticipated oil price decreases.

The quantitative importance of this channel depends on the extent of expenditure switching in response to real oil price shocks and on how pervasive frictions in capital and labor markets are. There is general agreement that the domestic automobile sector is most susceptible to the reallocation effect. For example, Edelstein & Kilian (2009) suggest that the magnitude of the reallocation effect depends primarily on the size of the domestic automobile industry (as measured by shares in employment and real output) as well as on the extent to which households substitute more energy-efficient imported cars for domestic cars. Edelstein & Kilian provide strong empirical evidence of shifts in demand from energy-intensive vehicles to energy-efficient vehicles in the 1970s and early 1980s. A more extensive study of the US automobile sector is Ramey & Vine (2010). Related evidence for the 2007–2008 recession can be found in Hamilton (2009).

3.3.2. Indirect effects: uncertainty effect. A second example of an amplification mechanism that involves asymmetric responses to positive and negative oil price shocks is Bernanke's (1983) model of uncertainty about the price of oil (see also Pindyck 1991). Bernanke's point is that—to the extent that the cash flow from an irreversible investment project depends on the price of oil—real options theory implies that, all else equal, increased uncertainty about the price of oil prompts firms to delay investments, causing investment expenditures to drop. As in models of the reallocation effect, the relevant oil price variable in these models is the real price of oil. Uncertainty in practice is measured by the expected volatility of the real price of oil over the relevant investment horizon.

Exactly the same reasoning applies to purchases of energy-intensive consumer durables such as cars. Because any unexpected change in the real price of oil may be associated with higher expected volatility, whether the real price of oil goes up or down, this uncertainty effect may amplify the effects of unexpected oil price increases and may offset the effects of unexpected oil price declines, much like the reallocation effect, resulting in asymmetric responses of real output. The quantitative importance of this channel depends on how important the real price of oil is for investment and durables purchase decisions and on the share of such expenditures in aggregate spending. For example, it seems intuitive that uncertainty about the price of oil would be important for decisions about oil drilling in Texas, but less obvious that it would be as important for other sectors of the economy such as textile production and information technology.

A closely related argument appears in Edelstein & Kilian (2009), who observe that increased uncertainty about the prospects of staying employed in the wake of unexpected changes in the real price of oil may cause an increase in precautionary savings (or, equivalently, a reduction in consumer expenditures). In this interpretation, uncertainty may affect not merely energy-intensive consumer durables such as cars, but other consumer expenditures as well. This argument is logically distinct from the observation that households will smooth their consumption to the extent that unexpectedly higher real oil prices are associated with an increased likelihood of becoming unemployed. The key difference is that consumption-smoothing motives are symmetric in unexpected oil price increases and decreases, whereas precautionary motives are not.

3.3.3. Indirect effects: systematic monetary policy responses. The third example of an amplification mechanism that involves asymmetric responses to positive and negative oil price shocks is due to Bernanke et al. (1997). Bernanke et al. hold the response of the Federal Reserve to oil price shocks responsible for the depth of the recessions following positive oil price shocks. Their premise is that the Federal Reserve chose to respond to incipient or actual inflationary pressures associated with unexpected real oil price increases by raising the interest rate, thereby amplifying the economic contraction. The asymmetry arose because the Federal Reserve, according to Bernanke et al., did not respond as vigorously to unexpected declines in the real price of oil. There is no theoretical model underlying this asymmetry in the Federal Reserve's responses, and the nature of the relevant counterfactual for monetary policy has remained controversial (e.g., Carlstrom & Fuerst 2006, Kilian & Lewis 2011). In fact, in recent years, the notion that policy makers should respond mechanically to oil price shocks has been shown to be at odds with economic theory (e.g., Kilian 2009b, Nakov & Pescatori 2010, Bodenstein et al. 2012).

4. CONSEQUENCES OF EXOGENOUS OIL PRICE SHOCKS: EMPIRICAL EVIDENCE

4.1. Linear Models of the Transmission of Oil Price Shocks

With the endogeneity of oil prices well established, much of the empirical literature focuses on the dynamic responses to unanticipated changes in the price of oil. The standard assumption in empirical work is that oil prices are predetermined with respect to domestic real output, which means that there is no contemporaneous feedback from domestic real macroeconomic aggregates to the price of oil. Kilian & Vega (2011) demonstrate that this identifying assumption is consistent with the data in that daily US macroeconomic news do not predict changes in the price of oil within a month, whereas they do predict changes in bond yields, stock returns, and other asset returns. As a result, linear VAR models can be used to quantify the response of US real GDP associated with an innovation (or surprise change) in the real price of oil. The results of this thought experiment can be represented as an impulse response function that traces out the response of US real GDP over time to a one-time oil price shock.

The assumption of predetermined oil prices does not necessarily merit a casual interpretation of these responses, however. As first discussed in Kilian (2008a, 2009b), the reason is that oil price innovations reflect demand and supply shocks in the global market for crude oil that move not only the real price of oil, but also other macroeconomic aggregates and commodity prices, violating the *ceteris paribus* assumption required for causal analysis. Nevertheless, the expected response to an unanticipated increase in the price of oil provides a first benchmark for thinking about the effects associated with oil price shocks that is commonly used in empirical work.

An additional caveat is that impulse response analysis alone is not sufficient for understanding the dynamic effects of oil price shocks. The aim of much of the literature since the 1980s has been to establish that the reduction in real output associated with an oil price shock is large. The presumption in that literature is that a large response of US real GDP to oil price shocks also means that oil price shocks explain large recessions. This is not necessarily the case, because impulse responses are hypothetical thought experiments that involve a one-time oil price innovation, whereas recessions are driven by sequences of oil price innovations of different magnitudes and signs. The net effect of a large positive oil price innovation followed by a large negative oil price innovation, for example, is to offset much of the dynamic response associated with the initial positive oil price shock. Only a historical decomposition of the US real output data can reveal the cumulative contribution of oil price shocks to US recessions (e.g., Kilian & Lewis 2011).

The empirical literature based on linear VAR models concludes that oil price shocks statistically explain only a modest fraction of the variation in US growth rates over time. In other words, the cumulative effects associated with oil price shocks are not negligible but are far from being the primary determinant of the business cycle. Although alternative theoretical explanations of the coincidence of oil price shocks and stagflation in the 1970s and early 1980s have been proposed in Barsky & Kilian (2002), many macroeconomists have found this type of result hard to accept in light of their prior views that oil price shocks played an important role in generating the major recessions of the 1970s.

4.2. Nonlinear Models of the Transmission of Oil Price Shocks

One possible explanation for the comparative lack of explanatory power of oil price shocks in linear models may be time variation in the relationship between oil prices and the macroeconomy. A number of studies modify the basic linear VAR framework to incorporate such nonlinearities. For example, Edelstein & Kilian (2009) study the effects of retail energy price shocks in a model allowing the share of energy in consumer expenditures to evolve over time, while preserving the linearity of the regression model in its parameters. Their impulse response results are qualitatively similar to results from standard linear VAR models of oil price shocks, however. In related work, Ramey & Vine (2011) propose an alternative adjustment to the price of gasoline that reflects the time cost of queuing in gasoline markets during the 1970s. This adjustment lowers the estimated response of industrial production to an oil price innovation in a sample extending from the late 1960s to 1985, making it more similar to the corresponding response estimate for data starting in 1986. Like the adjustment in Edelstein & Kilian, the Ramey & Vine adjustment does not help generate larger responses to energy price shocks.

4.2.1. General asymmetric models. Other researchers attribute the temporal instability of linear VAR models of the transmission of oil price shocks not to measurement problems in the oil price variable, but to an omitted nonlinearity in the regression of real output on the price of oil. Nonlinearities may arise in many different forms. Proponents of the view that positive oil price shocks have been the major cause of recessions in the United States inevitably appeal to the presence of asymmetries in the transmission of oil price shocks. Most popular in applied work are models in which only oil price increases affect real output and models in which real output depends only on the net oil price increase (defined as the maximum of zero and the deviation of the current oil price from the largest value of the oil price in recent years), as discussed in Hamilton (1996, 2003).[2] Indeed, such asymmetries are the only way to rationalize large recessions in the wake of oil price shocks theoretically, as discussed above. Thus, the question of how asymmetric the responses of real output are is central to the larger question of what lessons to draw from the historical evidence of the 1970s and 1980s. It is also important for assessing the effects of major unexpected declines in the price of oil, such as those that occurred in 1986, 1998, and late 2008.

It was customary until recently to suggest that the asymmetry in the transmission of oil price shocks to real output is well established (e.g., Davis & Haltiwanger 2001). Much of the empirical

[2]The latter specification cannot be derived from any of the theoretical models of Section 3 but is often motivated by behavioral arguments.

work cited in support of asymmetric responses to oil price shocks, however, does not directly test the hypothesis of the asymmetric transmission of oil price innovations but rather imposes it in estimation. The question of how to establish asymmetries of the response of real output depending on the sign of oil price innovations has received attention only in recent years. It has been shown that the simple diagnostic tests commonly used in the literature dating back to the 1990s are not informative about the degree of asymmetry of the response functions of real economic activity. An alternative is the impulse response–based test of the null of symmetric response functions proposed by Kilian & Vigfusson (2011a). The latter approach has become increasingly standard for testing and quantifying the asymmetry of responses to energy price shocks (e.g., Herrera et al. 2011, Venditti 2013). This approach is useful in that it allows us to test collectively the empirical relevance of the three indirect effects of oil price shocks discussed in Section 3, each of which generates asymmetries in the response of macroeconomic aggregates to positive and negative oil price innovations.

If the asymmetry of the responses can be established, there is the additional question of how to quantify the degree of asymmetry in the responses of real output triggered by positive and negative oil price innovations. The most common approach to this problem since the mid-1990s has been to include measures of oil price increases or net oil price increases as one of the variables in a structural VAR model under the maintained assumption that the oil price increase or net oil price increase is predetermined with respect to real output. This approach results in inconsistent impulse response estimates, however, as proved in Kilian & Vigfusson (2011a), and cannot be used to quantify the degree of asymmetry in the responses of real output. This problem afflicts some of the leading studies that have shaped the profession's perceptions about the asymmetric transmission of oil price shocks (e.g., Bernanke et al. 1997, Davis & Haltiwanger 2001, Lee & Ni 2002).

Only recently have the findings of these studies been reexamined using more appropriate econometric methods (e.g., Kilian & Vigfusson 2011a,b; Herrera et al. 2011, 2012; Alsalman & Herrera 2013; Herrera & Karaki 2013). The key findings can be summarized as follows: There is no evidence of statistically significant asymmetries in the response of aggregate US industrial production, real GDP, unemployment, employment, or real stock returns to oil price shocks of typical magnitude. For very large shocks of two standard deviations or larger, the evidence is more mixed, but only when the financial crisis period is included. This outcome may simply reflect overfitting because all statistical significance vanishes when one terminates the sample in the last quarter of 2007.

Given that both the reallocation hypothesis and the uncertainty hypothesis imply different responses to oil price shocks across sectors, several studies also investigate US stock prices, industrial production, and job flows at the disaggregate level using sectoral or plant-level data (Herrera et al. 2011, 2012; Alsalman & Herrera 2013). There is some evidence of asymmetries at the disaggregate level, but typically only in response to very large shocks and only for a small number of sectors. An obvious question of interest is whether these sectors are energy intensive in production or in use. The answer is mixed. Although there are asymmetries for some sectors that are energy intensive, there are also asymmetries for sectors that are not energy intensive. Moreover, not all energy-intensive sectors exhibit asymmetries. For example, the evidence for the automobile sector is mixed, which is at odds with what one would have expected on the basis of the theoretical literature. Finally, the degree of asymmetry in the estimated responses need not be large, even when the asymmetry is statistically significant.

In short, linear VAR models appear to provide a good approximation to the responses of the economy to oil price shocks at the aggregate level and for most disaggregates, but in some cases the aggregate results may obscure asymmetric adjustments at the disaggregate level. Thus, on balance, the indirect effects of the oil price shocks discussed above are not an important feature of US macroeconomic data. This result does not mean that there is no reallocation effect or uncertainty effect but means that these effects are small enough to be ignored in practice. At this point, no such work has been conducted for other economies, with the exception of Herrera et al. (2012), who report similar results for real output for a range of countries.

4.2.2. Special cases. The advantage of the econometric approach proposed by Kilian & Vigfusson (2011a) is that it encompasses a range of different theoretical explanations of the origin of asymmetric response functions. An alternative approach in the literature is to focus on one of the three main sources of asymmetry at a time and to embed this explanation within a parametric regression model. Of course, the latter approach would be expected to be informative only if the other sources of asymmetry could be ruled out a priori.

4.2.2.1. Quantifying the uncertainty effect. One set of studies focuses specifically on the uncertainty effect. For example, Elder & Serletis (2010) propose a GARCH-in-mean VAR model for the price of oil and US real output in an effort to assess the empirical content of the uncertainty hypothesis in Section 3.[3] Rather than testing the symmetry of the response functions directly, as proposed in Kilian & Vigfusson (2011a,b), Elder & Serletis test the null of no feedback from the one-quarter-ahead conditional variance of the real price of oil in the conditional mean equation. Elder & Serletis do not provide a formal test of the symmetry of the response functions, but if we take their point estimates at face value, there appears to be evidence of strongly asymmetric responses. Subsequently, Jo (2014) shows that Elder & Serletis's GARCH-in-mean VAR specification is unduly restrictive in that the same model innovation is driving both the conditional variance and the conditional mean. After this restriction is relaxed, the uncertainty effect is reduced to one-third of the estimate in Elder & Serletis, but some asymmetry in the response estimates remains.

These findings raise the question of how economically plausible a large uncertainty effect is. There are two reasons to be skeptical. One is that energy is not an important component of the cash flow of most investment projects in the economy. The other is that the relevant measure of oil price volatility in the theoretical models described by Bernanke (1983) is not the one-month- or one-quarter-ahead volatility of the real price of oil. Rather, the relevant volatility measure is the volatility of the real price of oil at horizons relevant to durables purchase and investment decisions; this volatility is much more stable and is typically measured over years or even decades rather than over months (e.g., Kilian & Vigfusson 2011b, Alquist et al. 2013).

4.2.2.2. Quantifying the role of monetary policy responses. Another set of studies, dating back to Bernanke et al. (1997), focuses on asymmetries arising from the Federal Reserve's response to oil price shocks. The most common approach in this literature has been to impose asymmetry in estimation rather than to test the asymmetry of the responses of the monetary policy maker on the basis of a model that encompasses symmetric as well as asymmetric responses. Bernanke et al.'s own analysis provided some evidence for the quantitative importance of the monetary policy reaction hypothesis, but Bernanke et al. stressed the fragility of their results. Hamilton & Herrera (2004) subsequently illustrated that the evidence reported in Bernanke et al. is sensitive to details of the model specification and vanishes under plausible alternative model specifications. They also showed that the policy counterfactuals discussed in Bernanke at al. are subject to the Lucas

[3]Although the uncertainty effect has played a prominent role in discussions of asymmetric responses of real output for two decades, Elder & Serletis's study is the first study to provide a fully specified model of this transmission mechanism, whereas earlier studies such as Lee et al. (1995) and Ferderer (1996) focus on single-equation models of the effect of oil price uncertainty on real output, in which the price of oil is treated as exogenous. The latter approach is inconsistent with the modern view that the real price of oil contains an important endogenous component.

critique. Subsequently, Kilian & Vigfusson (2011a) demonstrated that the asymmetric regression model proposed by Bernanke et al. and widely used in subsequent research is misspecified, making it impossible to infer the responses of real output to oil price shocks or the reaction of the Federal Reserve to oil price shocks. Kilian & Vigfusson also found that there is no statistically significant evidence of an asymmetry in the real output responses in Bernanke et al.'s model once the regression model is properly specified and evaluated. Their evidence suggests that this particular source of asymmetry cannot be detected in the data.

This evidence leaves open the question of whether the monetary policy authority reacts to oil price shocks symmetrically. Kilian & Lewis (2011)—using a linear (and hence symmetric) version of the Bernanke et al. (1997) VAR model—conclude that, even in the 1970s and early 1980s, endogenous monetary policy responses to oil price shocks had negligible cumulative effects on US real output and inflation. They also provide indirect evidence that the US policy response has differed depending on whether oil price shocks reflected oil supply or oil demand shocks, confirming a conjecture in Kilian (2009b) and casting doubt on the existence of a mechanical policy response to oil price shocks. The theoretical analysis in Nakov & Pescatori (2010) reinforces this point. Finally, Bodenstein et al. (2012), on the basis of an estimated global DSGE model with fully endogenous oil prices, confirm that the estimated weight on oil price shocks in the monetary policy rule is negligible, as is the weight in the welfare-optimized monetary policy reaction function. Hence, the potential role of monetary policy in amplifying the US economy's responses to oil price shocks may be safely ignored in practice.

5. INSIGHTS FROM STRUCTURAL OIL MARKET MODELS

An alternative explanation of the seeming instability of the response of macroeconomic aggregates to oil price shocks has been proposed by Kilian (2008a, 2009b). These studies emphasize that, in traditional VAR models with predetermined oil prices, any positive oil price innovation triggers a reduction in real output. Although this result is correct as a statistical average for the sample period starting in 1973, such response estimates are misleading whenever the oil price innovation reflects primarily positive shocks to the flow demand for oil, as was the case during 2003–2008, for example. In that case, the oil price shock is merely a symptom of unexpected global economic strength, and the resulting stimulus for domestic real output may neutralize (or even more than offset), in the short run, the growth-retarding effects of higher prices for imported oil and other imported commodities. Only as this stimulus dies out over time does the response of real output become negative. In contrast, the same oil price innovation may be even more negative than implied by the average response over the sample if, for example, the innovation reflects shocks to precautionary demand or speculative demand. In related work, Kilian (2009a) illustrates that conventional VAR models would have predicted a recession in 2005–2006 that never materialized because the oil price surge of 2003–2008 was driven by unexpectedly strong demand for oil from emerging Asia; this demand offset the negative effects of higher import prices for oil and other commodities. In contrast, response estimates from conventional VAR models would have provided a somewhat more accurate prediction during 1990, when oil prices moved for reasons other than shifts in flow demand.

Another implication of this point is that, as the importance of flow demand shocks relative to the importance of other shocks evolves over time, so does the estimated response of US real output to an oil price shock. This time variation in the estimates may occur even in the absence of structural change in the economy if the estimation sample is short, as is typically the case. This fact explains why researchers estimating the same conventional VAR model of oil price innovations may obtain very different and seemingly conflicting estimates of the effects of oil price shocks, depending on the choice of the estimation period. Only as the sample period grows large does the response estimate converge to the expected or average response. This point has been elaborated on in several contexts. For example, Kilian & Park (2009) empirically demonstrate that seemingly contradictory results in the literature about the dynamic relationship between oil prices and stock prices can be reconciled by observing how the composition of oil demand and oil supply shocks has changed over time. Indeed, there are situations in which higher oil prices are associated with higher stock prices, contrary to conventional wisdom.

This discussion illustrates that the common question about the causal effects of oil price shocks is ill posed. This question ignores the fact that oil price shocks are merely a symptom of the underlying structural oil demand and oil supply shocks. Because these structural shocks set in motion a variety of changes in the global economy, all of which affect the oil-importing economy in question, one cannot attribute the resulting dynamic response to the oil price shock alone as though it had occurred holding everything else constant. This point was first made by Kilian (2008a, 2009b) in the context of structural VAR models, but it can also be made using fully specified DSGE models (e.g., Bodenstein et al. 2012). To understand, for example, the consequences of the surge in oil prices in recent years, one needs to model the sources of growth in the global economy rather than thinking about exogenous oil price shocks.

It is sometimes claimed that whether a given oil price shock was caused by unexpected demand from China or by an exogenous supply disruption in the Middle East should not make a difference, as long as both shocks occur abroad (e.g., Blanchard & Galí 2010). This claim is incorrect, as Bodenstein et al. (2012) demonstrate within the framework of a global DSGE model with fully endogenous oil prices. Even when the magnitude of the oil price shock on impact is controlled for, the responses of the economy depend on the type of structural shock in the model and on where in the world this structural shock occurs.

This fact also points to a potential limitation of structural VAR models in the tradition of Kilian (2009b) and Kilian & Murphy (2014). These models cannot differentiate between flow demand shocks originating in different parts of the world or recognize the difference between a flow demand shock driven by foreign productivity shocks and one driven by foreign monetary policy shocks, for example. At best, they are able to capture average responses. Although DSGE models can make such fine distinctions, the ability of DSGE models to provide more detailed answers about the transmission of oil price shocks comes at the cost of having to fully specify the microeconomic structure of the model, which often involves ad hoc assumptions, so neither approach is without limitations. Nevertheless, the substantive implications of these two classes of models with regard to the surge in the price of oil from 2003 to mid-2008 are similar (e.g., Bodenstein & Guerrieri 2011).

An important question is how to quantify the impact of the deeper structural shocks such as flow supply shocks or flow demand shocks on US macroeconomic aggregates and other variables of interest. The answer is straightforward in the context of an estimated or calibrated DSGE model. When working with structural VAR models, researchers have used two approaches. One is to augment the oil market VAR model with additional variables. For example, Kilian (2010) proposes a VAR model with two blocks. The first block is the oil market block, which is identical to the model in Kilian (2009b), and the second block represents the US gasoline market. The additional identifying assumption is that the global oil market variables are predetermined with respect to US gasoline prices and consumption. Another example is Kilian & Park (2009), who augment the oil market model with US real stock returns, allowing stock returns to respond within the month to oil demand and supply shocks, but not the other way around. A similar approach may also be applied when sign restrictions are used for identification, in which case we can dispense with the assumption of predeterminedness (e.g., Baumeister & Peersman 2013b). An alternative approach has been to use the structural VAR shocks (possibly aggregated across time) as regressors in a second-stage distributed lag regression of the variable of interest on a constant and current and lagged values of the orthogonal shocks. This approach has been used in Kilian (2009a,b) and Kilian et al. (2009), for example. Estimation and inference in this case are complicated by the fact that the regressors are generated from another regression.

Although structural VAR models of oil markets go a long way toward explaining seeming time variation in the responses of real output to oil price shocks, there may be additional time variation even after one controls for the evolution of the structural shocks. A recent development in the literature is the introduction of structural oil market models with time-varying parameters (Baumeister & Peersman 2013a). This class of oil market models may be motivated, for example, by the facts that oil production as well as inventories may be subject to capacity constraints; that the energy intensity of economies has evolved over time; that financial markets for oil products have evolved, providing new opportunities for diversification and risk sharing; and that an oil-importing economy may become an oil exporter. One unresolved question is whether such models are too flexible to provide a good approximation in practice. Another problem is that these structural models tend to be computationally infeasible when monthly data are used.

6. CONCLUSION

The literature on oil markets and on the transmission of oil price shocks has made more progress than would have seemed possible 15 years ago. Some of the key insights are that the real price of oil is endogenous with respect to economic fundamentals and that oil price shocks do not occur *ceteris paribus*, making it necessary to account for the deeper structural shocks underlying oil price shocks when studying their transmission to the domestic economy. There are two paradigms for tackling such questions. One is the use of parsimonious structural VAR models of the global oil market based on minimal identifying assumptions, and the other is the use of fully specified DSGE models with explicit microeconomic structure.

The structural VAR approach to modeling oil markets has probably been advanced as far as currently possible. The problem is not only one of identifying additional shocks, but also one of the reliability of these regressions when one is fitting ever more parameters. At best, we will be able to augment the models by one or two additional macroeconomic variables as more data become available or perhaps to embed the model within a factor structure. For macroeconomic policy analysis, an even richer shock structure is required, which necessitates the development of truly global DSGE models in which the oil market is one of many markets with endogenously determined oil prices. The DSGE framework allows users to differentiate between, for example, fiscal and monetary policy shocks as well as productivity and oil intensity shocks while at the same time differentiating shocks by geographic origin. A prototype for such a model was recently introduced by Bodenstein et al. (2012).

These limitations do not mean that DSGE models will make structural VAR models obsolete. One reason is that at this point we lack the international data required to build reliable global DSGE models disaggregated by geographic region. An additional challenge is modeling the interaction of the financial sector and the real sector, especially at the global level (e.g., Kilian et al. 2009). In fact, the literature has barely begun to introduce expectations-driven shifts in inventory demand into DSGE models (e.g., Arseneau & Leduc 2012). A reasonable conjecture is that, as this literature evolves, international commodity markets will play a much more important role in global DSGE models than in the past.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I thank the reviewer of this article, Christiane Baumeister, and Ana María Herrera for helpful comments.

LITERATURE CITED

Almoguera PA, Douglas C, Herrera AM. 2011. Testing for the cartel in OPEC: noncooperative collusion or just noncooperative? *Oxf. Rev. Econ. Policy* 27:144–68

Alquist R, Kilian L. 2010. What do we learn from the price of crude oil futures? *J. Appl. Econ.* 25:539–73

Alquist R, Kilian L, Vigfusson RJ. 2013. Forecasting the price of oil. In *Handbook of Economic Forecasting*, Vol. 2, ed. G Elliott, A Timmermann, pp. 427–507. Amsterdam: North-Holland

Alsalman Z, Herrera AM. 2013. *Oil price shocks and the US stock market: Do sign and size matter?* Unpub. Pap., Dep. Econ., Univ. Ky.

Arseneau DM, Leduc S. 2012. *Commodity price movements in a general equilibrium model of storage*. Int. Financ. Discuss. Pap. 1,054, Fed. Reserv. Board

Atkeson A, Kehoe PJ. 1999. Models of energy use: putty-putty versus putty-clay. *Am. Econ. Rev.* 89:1028–43

Backus DK, Crucini MJ. 2000. Oil prices and the terms of trade. *J. Int. Econ.* 50:185–213

Barsky RB, Kilian L. 2002. Do we really know that oil caused the Great Stagflation? A monetary alternative. In *NBER Macroeconomics Annual 2001*, Vol. 16, ed. B Bernanke, K Rogoff, pp. 137–83. Cambridge, MA: MIT Press

Barsky RB, Kilian L. 2004. Oil and the macroeconomy since the 1970s. *J. Econ. Perspect.* 18(Fall):115–34

Basak S, Pavlova A. 2013. *A model of financialization of commodities*. Unpub. Pap., Lond. Bus. Sch.

Baumeister C, Peersman G. 2013a. The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *J. Appl. Econ.* 28:1087–109

Baumeister C, Peersman G. 2013b. Time-varying effects of oil supply shocks on the US economy. *AEJ Macroecon.* 5:1–29

Bernanke BS. 1983. Irreversibility, uncertainty, and cyclical investment. *Q. J. Econ.* 98:85–106

Bernanke BS, Gertler M, Watson MW. 1997. Systematic monetary policy and the effects of oil price shocks. *Brookings Pap. Econ. Act.* 1:91–142

Blanchard OJ, Galí J. 2010. The macroeconomic effects of oil price shocks: Why are the 2000s so different from the 1970s? In *International Dimensions of Monetary Policy*, ed. J Galí, M Gertler, pp. 373–421. Chicago: Univ. Chicago Press

Bodenstein M, Guerrieri L. 2011. *Oil efficiency, demand and prices: a tale of ups and downs*. Unpub. Pap., Fed. Reserv. Board

Bodenstein M, Guerrieri L, Kilian L. 2012. Monetary policy responses to oil price fluctuations. *IMF Econ. Rev.* 60:470–504

Bresnahan T, Ramey V. 1993. Segment shifts and capacity utilization in the US automobile industry. *Am. Econ. Rev. Pap. Proc.* 83:213–18

Bruno M, Sachs J. 1982. Input price shocks and the slowdown in economic growth: the case of UK manufacturing. *Rev. Econ. Stud.* 49:679–705

Carlstrom CT, Fuerst TS. 2006. Oil prices, monetary policy, and counterfactual experiments. *J. Money Credit Bank.* 38:1945–58

Davis SJ, Haltiwanger J. 2001. Sectoral job creation and destruction responses to oil price changes. *J. Monet. Econ.* 48:465–512

Dhawan R, Jeske K. 2008. Energy price shocks and the macroeconomy: the role of consumer durables. *J. Money Credit Bank.* 40:1357–77

Dvir E, Rogoff K. 2010. *Three epochs of oil*. Unpub. Pap., Harvard Univ.

Edelstein P, Kilian L. 2007. The response of business fixed investment to energy price changes: a test of some hypotheses about the transmission of energy price shocks. *B.E. J. Macroecon.* 7

Edelstein P, Kilian L. 2009. How sensitive are consumer expenditures to retail energy prices? *J. Monet. Econ.* 56:766–79

Elder J, Serletis A. 2010. Oil price uncertainty. *J. Money Credit Bank.* 42:1137–59
Fama EF, French KR. 1987. Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage. *J. Bus.* 60:55–73
Fattouh B, Kilian L, Mahadeva L. 2013. The role of speculation in oil markets: What have we learned so far? *Energy J.* 34:7–33
Fattouh B, Mahadeva L. 2013. *Causes and implications of shifts in financial participation in commodity markets.* Unpub. Pap., Oxford Univ.
Ferderer JP. 1996. Oil price volatility and the macroeconomy: a solution to the asymmetry puzzle. *J. Macroecon.* 18:1–16
Finn MG. 2000. Perfect competition and the effects of energy price increases on economic activity. *J. Money Credit Bank.* 32:400–16
Green EJ, Porter RH. 1984. Noncooperative collusion under imperfect price information. *Econometrica* 52:87–100
Hamilton JD. 1983. Oil and the macroeconomy since World War II. *J. Polit. Econ.* 91:228–48
Hamilton JD. 1988. A neoclassical model of unemployment and the business cycle. *J. Polit. Econ.* 96:593–617
Hamilton JD. 1996. This is what happened to the oil price–macroeconomy relationship. *J. Monet. Econ.* 38:215–20
Hamilton JD. 2003. What is an oil shock? *J. Econom.* 113:363–98
Hamilton JD. 2009. Causes and consequences of the oil shock of 2007–08. *Brookings Pap. Econ. Act.* 1(Spring):215–61
Hamilton JD. 2012. Commentary: import prices and inflation. *Int. J. Cent. Bank.* 8:271–79
Hamilton JD, Herrera AM. 2004. Oil shocks and aggregate economic behavior: the role of monetary policy. *J. Money Credit Bank.* 36:265–86
Hamilton JD, Wu JC. 2014. Risk premia in crude oil futures prices. *J. Int. Money Financ.* 42:9–37
Herrera AM, Karaki MB. 2013. *The effects of oil price shocks on job reallocation.* Unpub. Pap., Univ. Ky.
Herrera AM, Lagalo LG, Wada T. 2011. Oil price shocks and industrial production: Is the relationship linear? *Macroecon. Dynam.* 15:472–97
Herrera AM, Lagalo LG, Wada T. 2012. *Asymmetries in the oil price–industrial production relationship? Evidence from 18 OECD countries.* Unpub. Pap., Univ. Ky.
Hu C, Xiong W. 2013. *Are commodity futures prices barometers of the global economy?* Unpub. Pap., Princeton Univ.
Jo S. 2014. The effect of oil price uncertainty on global real economic activity. *J. Money Credit Bank.* Forthcoming
Kilian L. 2008a. The economic effects of energy price shocks. *J. Econ. Lit.* 46:871–909
Kilian L. 2008b. Exogenous oil supply shocks: How big are they and how much do they matter for the US economy? *Rev. Econ. Stat.* 90:216–40
Kilian L. 2009a. Comment on "Causes and consequences of the oil shock of 2007–08" by James D. Hamilton. *Brookings Pap. Econ. Act.* 1(Spring):267–78
Kilian L. 2009b. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *Am. Econ. Rev.* 99:1053–69
Kilian L. 2010. Explaining fluctuations in us gasoline prices: a joint model of the global crude oil market and the US retail gasoline market. *Energy J.* 31:87–104
Kilian L, Hicks B. 2013. Did unexpectedly strong economic growth cause the oil price shock of 2003-2008? *J. Forecast.* 32:385–94
Kilian L, Lee TK. 2014. Quantifying the speculative component in the real price of oil: the role of global oil inventories. *J. Int. Money Financ.* 42:71–87
Kilian L, Lewis LT. 2011. Does the Fed respond to oil price shocks? *Econ. J.* 121:1047–72
Kilian L, Murphy DP. 2012. Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models. *J. Eur. Econ. Assoc.* 10:1166–88
Kilian L, Murphy DP. 2014. The role of inventories and speculative trading in the global market for crude oil. *J. Appl. Econom.* 29:454–78
Kilian L, Park C. 2009. The impact of oil price shocks on the US stock market. *Int. Econ. Rev.* 50:1267–87

Kilian L, Rebucci A, Spatafora N. 2009. Oil shocks and external balances. *J. Int. Econ.* 77:181–94
Kilian L, Vega C. 2011. Do energy prices respond to US macroeconomic news? A test of the hypothesis of predetermined energy prices. *Rev. Econ. Stat.* 93:660–71
Kilian L, Vigfusson RJ. 2011a. Are the responses of the US economy asymmetric in energy price increases and decreases? *Quant. Econ.* 2:419–53
Kilian L, Vigfusson RJ. 2011b. Nonlinearities in the oil price–output relationship. *Macroecon. Dynam.* 15(Suppl. 3):337–63
Kilian L, Vigfusson RJ. 2013. Do oil prices help forecast US real GDP? The role of nonlinearities and asymmetries. *J. Bus. Econ. Stat.* 31:78–93
Kim IM, Loungani P. 1992. The role of energy in real business cycle models. *J. Monet. Econ.* 29:173–89
Knittel CR, Pindyck RS. 2013. *The simple economics of commodity price speculation*. Unpub. Pap., Sloan Sch. Manag., Mass. Inst. Technol.
Leduc S, Sill K. 2004. A quantitative analysis of oil price shocks, systematic monetary policy, and economic downturns. *J. Monet. Econ.* 51:781–808
Lee K, Ni S. 2002. On the dynamic effects of oil price shocks: a study using industry-level data. *J. Monet. Econ.* 49:823–52
Lee K, Ni S, Ratti RA. 1995. Oil shocks and the macroeconomy: the role of price variability. *Energy J.* 16:39–56
Nakov A, Nuño G. 2013. Saudi Aramco and the oil market. *Econ. J.* 123:1333–62
Nakov A, Pescatori A. 2010. Monetary policy trade-offs with a dominant oil producer. *J. Money Credit Bank.* 42:1–32
Pindyck RS. 1991. Irreversibility, uncertainty and investment. *J. Econ. Lit.* 29:1110–48
Ramey VA, Vine DJ. 2011. Oil, automobiles, and the US economy: How much have things really changed? *NBER Macroeconomics Annual 2010*, Vol. 25, ed. D Acemoglu, M Woodford, 333–68. Chicago: Univ. Chicago Press
Rotemberg JJ, Woodford M. 1996. Imperfect competition and the effects of energy price increases on economic activity. *J. Money Credit Bank.* 28:549–77
Skeet I. 1988. *OPEC: Twenty-Five Years of Prices and Politics*. Cambridge, UK: Cambridge Univ. Press
Smith JL. 2005. Inscrutable OPEC? Behavioral tests of the cartel hypothesis. *Energy J.* 26:51–82
Sockin M, Xiong W. 2013. *Feedback effects of commodity futures prices*. Unpub. Pap., Princeton Univ.
Venditti F. 2013. From oil to consumer energy prices: how much asymmetry along the way? *Energy Econ.* 40:468–73