

Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market

Lutz Kilian
University of Michigan and CEPR

February 23, 2007

Abstract: Using a newly developed measure of global real economic activity, a structural decomposition of the real price of crude oil into three components is proposed: crude oil supply shocks; shocks to the aggregate global demand for industrial commodities; and demand shocks that are specific to the crude oil market. The latter shock is designed to capture shifts in the price of oil driven by higher precautionary demand associated with concerns about the availability of future oil supplies. The paper estimates the dynamic effects of these shocks on the real price of oil. A historical decomposition sheds light on the causes of the major oil price shocks since 1975. The effects of higher oil prices on U.S. real GDP and CPI inflation are shown to depend on the cause of the oil price increase, suggesting that policies aimed at dealing with higher oil prices must take careful account of the origins of higher oil prices. The paper also quantifies the extent to which the macroeconomic performance of the U.S. since the late 1970s has been determined by the external economic shocks driving the real price of oil as opposed to domestic shocks.

Key words: Oil price; oil demand shocks; oil supply shocks; dynamic effects.

JEL: E31, E32, Q43

Acknowledgements: I thank Ana María Herrera and Eric Sims for helpful discussions.

1. Introduction

A common approach in both empirical and theoretical work on oil price shocks is to evaluate the response of macroeconomic aggregates to changes in the price of oil.¹ Implicit in this approach is a thought experiment in which one varies the price of oil, while holding all other variables constant. This thought experiment is not well defined. For example, Bernanke (2004) notes that, as a professor and textbook author, he “was accustomed to discussing the effects of ... rising oil prices with all other factors held equal. However, as policymakers know, everything else is never held equal. The increases in oil prices this year did not take place in isolation.”

The problem is not just that other factors such as economic expansions, existing inflation, fluctuations in the dollar or changes in interest rates, may cushion or amplify the effects of higher oil prices, as stressed in Bernanke’s (2004) speech, but more importantly that higher oil prices in turn may be driven by global macroeconomic aggregates. This means that cause and effect are no longer well defined when relating changes in the price of oil to macroeconomic outcomes. Thus, we have to move beyond studying changes in the price of oil and address the problem of identifying the structural shocks underlying the price of oil. The first objective of this paper is to propose a model that allows the identification of these shocks and helps us understand their relative importance in determining the real price of oil.

The identification of these shocks is important not just for explaining fluctuations in the real price of oil, but also for the design of macroeconomic policies in response to higher oil prices. Implicit in the literature on the effects of oil price changes and in statements of many policy-makers is the view that an increase in the price of oil has the same effect regardless of the underlying cause of that increase. This interpretation allows one to discuss the effects of higher oil prices as though it did not matter what drove up oil prices in the first place. The second objective of this paper is to demonstrate that this interpretation is incorrect.

Using a newly developed measure of global real economic activity, a structural decomposition of the real price of crude oil into three components is proposed: crude oil supply shocks; shocks to the aggregate global demand for industrial commodities; and global demand shocks that are specific to the crude oil market. The latter shock is designed to capture shifts in the price of oil driven by higher precautionary demand associated with fears about the

¹ For the purpose of this paper it makes no difference whether changes in the price of oil are defined as percent changes in the price of oil, oil price increases, or net oil price increases. All these approaches effectively treat the change in the price of oil as exogenous.

availability of future oil supplies. While it is widely accepted that shifting concerns about future oil supplies that are orthogonal to observable changes in crude oil production are an important source of oil price fluctuations, the magnitude of these shocks and their effect on the real price of oil has never been estimated. The analysis in this paper directly addresses this question.

The paper provides estimates of the dynamic effects of these shocks on the real price of oil and estimates of the cumulative effect of each of these shocks on the real price of oil during 1975-2005. My analysis sheds light on the origin of the observed fluctuations in oil prices, in particular during oil price shocks. For example, it helps gauge the relative importance of these shocks in the rapid build-up of the real price of crude oil since the late 1990s. Distinguishing between the sources of higher oil prices is shown to be crucial for assessing the effect of higher oil prices on U.S. real GDP and CPI inflation, suggesting that policies aimed at dealing with higher oil prices must take careful account of the origins of higher oil prices. The proposed framework also allows me to quantify the extent to which U.S. real growth and CPI inflation since the late 1970s have been driven by external shocks (as embodied in the real price of oil) as opposed to domestic shocks.

One of my key findings is that no two oil price shocks are alike. Nevertheless, there are some regularities. All the major real oil price increases since the mid-1970s can be traced to increased global aggregate demand for industrial commodities and/or oil-market specific increases in demand. The latter demand shifts are consistent with sharp increases in precautionary demand in the wake of exogenous political events in the Middle East. In contrast, disruptions of crude oil production play a less important role, suggesting that the traditional approach of linking oil price increases to exogenous shortfalls in crude oil production must be rethought. The most recent build-up in the real price of oil, in particular, can be attributed almost entirely to positive global aggregate demand shocks. Moreover, when political events do affect oil prices, as happened after the Iranian Revolution or during the Persian Gulf War, my analysis suggests that it is less the actual physical supply disruptions than the increased precautionary demand for oil triggered by fears about future oil supply shortfalls (which may or may not be realized) that is driving the price of oil.

There are important differences in how these demand and supply shocks in global crude oil markets affect U.S. real GDP growth and CPI inflation. For example, at the 10 percent significance level, oil supply disruptions cause a statistically significant decline in real GDP on

impact, and a barely statistically significant and small increase in CPI inflation after one quarter with no evidence of subsequent inflationary pressures. Global aggregate demand expansions tend to raise U.S. real GDP growth in the first year after the shock, but lower it in the second year. They also cause a persistent increase in CPI inflation. In contrast, oil-market specific increases in demand persistently lower U.S. real GDP growth and persistently raise U.S. CPI inflation at the same time. Moreover, only for oil-market specific demand shocks is there statistically significant evidence of stagflationary responses.

A substantial component of U.S. real growth and CPI inflation since the 1970s can be traced to these external shocks. The relative contribution of domestic and external factors varies over time. For example, about half of the U.S. CPI inflation of 1980 was home-made, whereas the bulk of CPI inflation in late 2005 was caused by the external shocks driving the oil market. There also is strong evidence that the disinflation of the early 1980s was facilitated by favorable developments in the crude oil market.

The remainder of the paper is organized as follows. Section 2 describes the data used in identifying the structural shocks underlying the real price of oil. In this context, I introduce a new measure of global real economic activity based on data for dry cargo bulk freight rates. Section 3 focuses on the identification of the structural shocks that drive the real price of oil. I estimate the dynamic effects of these shocks on the real price of oil and quantify their historical contribution to the real price of oil. Section 4 investigates the impact of the shocks identified in section 3 on U.S. macroeconomic aggregates. Section 5 concludes.

2. Data

2.1. Real Price of Oil

Figure 1 plots the real price of oil, expressed in January 1981 dollars, for 1973.1-2005.9. The series is obtained based on the refiner acquisition cost of imported crude oil, as provided by the U.S. Department of Energy since 1974.1 and extended backward as in Barsky and Kilian (2002), and is deflated using the U.S. CPI. It is apparent that the real price of oil, following an all-time low in 1998, has rebounded to levels only surpassed in 1979-83. This fact has made it all the more urgent to understand the underlying causes of that increase.

There has always been a tendency to identify major movements in the price of oil with events that are presumably exogenous to the U.S. macroeconomy. The vertical lines in Figure 1 indicate major events of relevance to the oil market. For example, there are marked increases in

the real price of oil following the Yom Kippur War and Arab oil embargo of 1973/74, the Iranian Revolution of 1978/79 and the outbreak of the Persian Gulf War in 1990. There are much smaller increases after the outbreak of the Iran-Iraq War in late 1980 and in the months leading up to the 2003 Iraq War.

Whereas these events seem primarily relevant for the supply side of the crude oil market, the sharp drop in the price of crude oil during the Asian crisis is a good example of an exogenous demand shock. For completeness, the plot also shows the date of Hurricanes Rita and Katrina at the very end of the sample. The latter exogenous events are best thought of as negative demand shocks for crude oil rather than crude oil supply shocks. The reduction in U.S. crude oil production in the Gulf of Mexico caused by the hurricanes was comparatively minor measured on a global scale. More important was the loss of U.S. refining capacity. As refineries shut down, U.S. demand for crude oil fell and the world price of crude oil dropped.

In addition, Figure 1 shows shaded areas marking periods of active oil supply management. The first period extends into October of 1973. It refers to the end of the old post-war order when U.S. oil companies essentially controlled the price of oil as well as crude oil supplies in Arab oil producing countries. The companies' objective was to keep the price of oil low, while increasing supplies. In contrast, the point of the supply management by OPEC in 1982.3-1982.12 and, after that attempt failed, again in 1983.3-1985.12 was to reduce crude oil supplies in order to stem the decline of the price of oil.² The nature of the OPEC supply management was simple. Saudi Arabia, as the swing producer, would reduce its crude oil production, conditional on the agreed upon production levels of other OPEC countries, as much as required to stem the decline in crude oil prices. This arrangement was ultimately undone by cheating cartel members as well as offsetting increases in crude oil production elsewhere, causing Saudi Arabia to unilaterally withdraw from the cartel in late 1985.

All three periods show that supply management was partially successful (in that the rate of change of the real price of oil was slowed down); in all three cases participants ultimately found the arrangement unworkable; and all three episodes were followed by all the more rapid price adjustments, when supply management ended. This evidence suggests that the real price of oil – far from being mainly controlled by cartels – is ultimately determined by market forces and

² The dating of these periods is based on the analysis in Skeet (1988). There was no coordination of OPEC supply decisions, no concerted supply restraint and no quota system during 1974-1982.

subject to demand and supply shocks like any other industrial commodity. The main objective of this paper will be to quantify the relative importance of various demand and supply shocks for the real price of oil. For this purpose, I will relate the real oil price series to the additional data described below.

2.2. Global Oil Production

An important determinant of the real price of oil is global crude oil production as reported by the U.S. Department of Energy. Figure 2 plots this time series (henceforth referred to as *oil supply* for simplicity). The upper panel shows the level of crude oil production in millions of barrels pumped per day (averaged by month). The lower panel shows the same series expressed in annualized percentage changes. The changes in crude oil production shown here reflect exogenous political oil supply shocks, the cartel activities discussed in section 2.1 and internal OPEC power struggles in the 1970s and early 1980s, but they also reflect endogenous responses to changes in the real price of oil in OPEC and non-OPEC countries (see, e.g., Skeet 1988).³

2.4. Global Real Economic Activity

Measures of global oil production are essential in modeling the supply side of the crude oil market. Of equal importance is an explicit measure of global real economic activity because of its effect on the demand for crude oil (see Barsky and Kilian 2002, 2004). Global economic activity is difficult to measure for three reasons. First, the approach to identifying structural shocks to the real price of oil adopted in this paper heavily relies on delay restrictions that are economically plausible only at the monthly frequency, yet for many countries measures of value added are not available at monthly or even at quarterly frequency. This is true not just for emerging economies such as China and India, but also for many of the smaller industrialized economies.

Second, it is not straightforward to weight appropriately each country's contribution to global real economic activity. Commonly used exchange-rate weighted averages are at best crude proxies. To make matters worse, the relative importance of individual countries for global economic activity is shifting over time. For example, the contribution of Asian countries has increased in recent years. Properly accounting for these shifting weights is a daunting task.

³ In an earlier version of this paper, I further decomposed changes in the production of crude oil into crude oil supply shocks driven by exogenous political events in the Middle East, building on recent work in Kilian (2006a,b), and other crude oil supply shocks (see Kilian 2006c). This distinction does not change the main results of this paper.

Third, value added is not the most appropriate measure of real economic activity for understanding industrial commodity markets. For many major industrialized economies an increasing share of value added relates to the service industry, which utilizes industrial commodities far less than manufacturing for example. This structural transformation renders the link from real GDP to the demand for industrial commodities unstable.

An alternative measure of real economic activity would be industrial production. Even for indices of industrial production, however, the problems of weighting and data availability remain and technological changes may over time affect the link from rising production to the demand for industrial commodities. No accepted monthly index of global industrial production exists in the literature.⁴ For these reasons, in this paper, I propose an alternative measure of global real economic activity. This measure is based on a global index of dry cargo single voyage freight rates. The index is measured at monthly frequency and can be constructed as far back as January 1968.

2.4.1. Motivation

My approach to measuring real economic activity is not without precedence. Similar techniques have been used in economic history to measure business cycles. Economists have long observed a positive correlation between ocean freight rates and economic activity (see, e.g., Isserlis 1938, Tinbergen 1959, Stopford 1997, Klovand 2004).

It is widely accepted that world economic activity is by far the most important determinant of the demand for transport services (see, e.g., Klovland 2004). As documented by Stopford (1997), at low levels of freight volumes the supply curve of shipping is relatively flat in the short and intermediate run, as idle ships may be reactivated or active ships may simply cut short layovers and run faster. As the demand schedule for shipping services shifts out due to increased economic activity, the slope of the supply curve becomes increasingly steeper and freight rates increase. At full capacity the supply curve becomes effectively vertical, as all available ships are operational and running at full speed. Only in the long-run will additional ship-building lower freight rates, often at a time when the initial high levels of economic activity have already subsided. Following a global business cycle upswing there is likely to be a rather drawn out trough period in the shipping market, as new ships are still being launched long after

⁴ For example, the index provided in the *International Financial Statistics* of the IMF only dates back to the 1980s, and it is not clear how this index was constructed.

the business cycle peak has passed and excess capacity of shipping prevails. Only gradually scrapping of older ships and rising demand due to the business cycle will offset this depression in the shipping market.

This line of reasoning suggests that increases in freight rates may be used as indicators of strong cumulative global demand pressures. I will use this insight to identify periods of high and low real economic activity. While an index of real economic activity based on global dry cargo freight rates offers clear advantages compared to, for example, measures of global industrial production, it is not free of drawbacks. In particular, the presence of a ship-building and scrapping cycle may weaken the link between real economic activity and the freight rate index. Given the pro-cyclicality of ship-building, one would expect the real freight rate index to lag increases in real economic activity (as spare capacity in shipping cushions the impact of higher demand on freight rates) and to lead decreases in real economic activity (as the arrival of new ships depresses freight rates), thus accentuating upswings in real economic activity. On the other hand, the proposed index is a direct measure of global economic activity that does not require exchange-rate weighting, that automatically aggregates real economic activity in all countries, and that already incorporates shifting country weights, changes in the composition of real output, and changes in the propensity to import industrial commodities for a given unit of real output.

2.4.2. Construction of the Index

The index of global real economic activity derived below is based on representative single voyage freight rates collected by Drewry Shipping Consultants Ltd. for various bulk dry cargoes including grain, oilseeds, coal, iron ore, fertilizer and scrap metal.⁵ Quotes are provided for different commodities, routes and ship sizes. These quotes were entered manually into a spreadsheet, since the data are only available in hardcopy. The upper panel of Figure 3 shows the raw data. Freight rates are typically quoted in U.S. dollars per metric ton. There is no continuous series for the entire sample period. Taking simple averages of the freight rates in Figure 3 would ignore the existence of different fixed effects for different routes, commodities and ship sizes. In constructing an index of dry bulk cargo freight rates I eliminate these fixed effects as follows: I first compute the period-to-period growth rates for each series in the first panel of Figure 4, as far

⁵ Ideally, one would want to restrict the sample to industrial commodities. Grain is included in this index because the earliest data on dry cargo rates are only reported in the form of indices that include grain among other dry cargoes. For later periods, there is no evidence that freight rates for grains behave differently from freight rates for other dry cargoes.

as the data are available. I then take the equal-weighted average of these growth rates, and cumulate the average growth rate, having normalized January of 1968 to unity. The resulting index is shown in the second panel of Figure 3.⁶

The next step is to deflate this series with the U.S. CPI. Finally, the real index must be detrended. As is well known, the cost of shipping dry cargo has fallen dramatically in real terms over the decades. That trend reflects technological advances in ship-building. It may also be related to long-run trends in the demand for sea transport. As my interest in this paper centers on cyclical variation in ocean freight rates rather than on long-term trends, I linearly detrend the real freight rate index. The deviations of the real freight rates from their long-run trend are shown in the last panel of Figure 3. The linear regression analysis in section 3 is based on the assumption that the unobserved level of global real economic activity is proportionate to this index.⁷

2.4.3. The Global Business Cycle

There is little direct evidence on the global business cycle, but some anecdotal evidence. For example, many researchers have noted that the 1972-74 period was characterized by a global boom, as was to a lesser extent the 1978-80 period (see, e.g., Darmstadter and Landsberg 1976). We also know that the mid-1970s and the early 1980s saw worldwide recessions. Finally, we know that there has been a global boom in commodity markets in the early 2000s driven by strong economic growth worldwide.

An important test of the plausibility of the proposed index of real economic activity is whether it is consistent with these stylized facts. The last panel of Figure 3 suggests that the proposed measure passes this test. It also sheds light on the quantitative importance and timing of these fluctuations in real economic activity. The first major peak in real economic activity occurs in October of 1970, followed by a trough in March of 1972. Following a rapid recovery, the next peak occurs in December of 1973, followed by a trough in February of 1976. Real economic activity remains low throughout the mid-1970s. The third major expansion starts in 1977 and

⁶ Ideally, one would like to apply different weights for different growth rates, but such weights are not provided by Drewry's Shipping Consultants. For the same reason, equal weights are routinely used in the construction of commodity price indices.

⁷ It may seem that the plausibility of this index could be assessed by comparing it to a PPP-weighted measure of global real GDP if not at monthly frequency, then at quarterly or annual frequency. Such a comparison is unlikely to be meaningful, however, given the increasing importance of the service sector and the declining dependence of value added on industrial commodity imports. In other words, the purpose of the index of real economic activity proposed in this paper is not to measure global value added, but to measure the component of real economic activity that is reflected in aggregate demand shifts in global industrial commodity markets.

peaks in July of 1979. It is followed in the 1980s by a protracted period of low real activity, punctuated by an initial trough in August of 1982 and followed by an even deeper trough in July of 1986. By March of 1988 real activity has recovered and remains flat until January of 1990, followed by a long period of slightly below average real activity that predates the invasion of Kuwait. That period (with few interruptions) continues until October of 1999. This evidence is consistent with the notion that in the 1990s the business cycle all but vanished. That impression is proven wrong around 1999. The expansion that starts in July of 1998, after the Asian Financial Crisis has run its course, persists until November 2000. It is followed by a decline that starts well in advance of 9/11 and culminates in a shallow trough in November of 2001, before the expansion continues with an apparent peak in December of 2004.

Apart from the timing of the major expansions and contractions, their magnitudes are of immediate interest for understanding fluctuations in commodity prices. The three periods of highest real economic activity are 1968.9-1971.4, 1972.9-1975.2 and 2002.3-2005.9 with additional periods of sustained high real activity in 1979.3-1981.4 and 1999.10-2001.7. Interestingly, the sustained high levels of real economic activity since 2002 are very much reminiscent of those observed in 1972.9-1975.2. The vertical lines in the last panel of Figure 3 correspond to the oil dates shown in Figure 1. They illustrate that many of these events coincided with periods of high real economic activity and hence strong demand for industrial commodities. Thus, one would want to be careful about associating the concurrent increases in the real price of oil with these events. This evidence underscores the importance of disentangling the effects of demand shocks and supply shocks on the real price of oil.

2.4.4. Further Discussion of the Rationale of the Proposed Index of Real Economic Activity

There are a number of further aspects of the proposed measure that seem worth discussing. One obvious concern is that the index may be contaminated by idiosyncratic shocks in the markets for the commodities shipped as bulk dry cargo. Given the fairly large number of dry cargoes and routes, one would expect idiosyncratic supply shocks (as well as idiosyncratic demand shocks) to average out. The main concern is with rare, but large idiosyncratic shocks. A good example is the shock to the demand for grain that occurred in 1972, when Russia experienced a harvest failure and substituted imports for domestic production; although it is not clear to what extent the grain was shipped by commercial vessels as opposed to Soviet state-owned vessels (in which case there would be no effect on freight rates). In any case, that particular shock pre-dates the

sample period used in the econometric analysis of this paper, and there is no evidence of similar shocks for the remainder of the sample.

A second objection is that dry cargo freight rates may increase during oil price shocks not because both are driven by higher demand for commodities, but because the provision of shipping services uses bunker fuel oil as an input. There are several reasons to think that this link is not quantitatively important. First, records in the *Oil and Gas Journal* indicate that during 1970-1973 the real price of bunker fuel changed very little, yet the index of real economic activity underwent fluctuations of the same magnitude as during later times (see Figure 4). Second, Figure 3 shows that the freight rate index moved very little when the real price of oil dropped sharply in 1985/86. Similarly, during the Persian Gulf War in 1990/91, freight rates first dropped, when oil prices rose sharply, and then rose, as the price of oil dropped again. This evidence is consistent with the view that the cost share of bunker fuel oil in ocean shipping is small.

Finally, one may ask why I did not include the seemingly most relevant information on crude oil tanker rates available from Drewry's Shipping Monthly. The reason is that such rates, while there is typically strong comovement with dry cargo rates, at times may be subject to important oil-market specific supply shocks, which makes them a potentially poor measure of real economic activity. For example, attacks on shipping in the Persian Gulf may raise the insurance premium for tankers (and hence tanker rates). The same applies to transportation surcharges, as tankers are rerouted, although by 1973 most tanker traffic bypassed the Suez Canal, making this argument largely obsolete. While the closure of sea lanes or canals may also force the re-routing of dry-cargo shipping with concomitant increases in average freight rates, in practice that effect is of much less importance for the dry cargo market.

In addition, events such an oil embargo may lower the demand for tankers (and hence tanker rates) simply because there is no oil to be shipped, not because consumers' demand for oil has decreased, making it impossible to gauge the state of demand in the crude oil market. I circumvent this difficulty by using dry cargo rates as a measure of the general state of demand for industrial commodities, and treating shocks to the demand for crude oil as the residual that is not accounted for by either crude oil supply shocks or aggregate demand shocks for industrial commodities, as illustrated in the next section.

3. Decomposing the Real Price of Oil

3.1. Methodology

Numerous empirical and theoretical studies have investigated the response of macroeconomic aggregates to changes in the price of oil. Implicit in this literature is the thought experiment that we can change the price of oil, while holding everything else constant, as would be the case if the price of oil were exogenous. To the extent that the price of oil is actually endogenous with respect to the macroeconomic aggregates of interest, this thought experiment is violated. If there is no well defined cause, it becomes impossible to estimate its effect. This general principle has been recognized dating back to the Cowles Commission. As Cooley and LeRoy (1985, p. 295) summarize, it is inadmissible to inquire about the effect of a change in one endogenous variable on another, when the underlying experiment that led to the assumed variation in the endogenous variable is ambiguous.

This problem has not completely escaped attention. Implicitly or explicitly, many researchers have assumed that at least the major increases in the price of oil can be treated as exogenous. Recent research has demonstrated that this interpretation, which seemed reasonable at the time, does not hold up to scrutiny (see, e.g., Kilian 2006a). This means that, quite simply, without knowing what drove up the price of oil in the first place, it will be impossible to predict the effect of higher oil prices. In this section, I propose a methodology for decomposing unpredictable changes in the real price of oil into mutually orthogonal components with structural economic interpretations. As I will argue below, this decomposition has immediate and important implications for how macroeconomists and policymakers should think about oil price fluctuations.

My identification of the structural shocks that drive the real price of crude oil rests on the ability of the monthly measure of global real economic activity proposed in section 2.4 to capture fluctuations in global real economic activity, as they relate to the market for industrial commodities. While this series has been carefully constructed, obviously it will be possible to construct alternative measures of global real economic activity. While these choices may affect the empirical estimates presented below, the methodological approach proposed in this paper is quite general.

3.2. Structural VAR Model

I specify a structural VAR model based on monthly data for $z_t = (\Delta prod_t, rea_t, rpo_t)'$, where $\Delta prod_t$ is the percent change in global crude oil production, rea_t denotes real economic activity (which is understood to refer to real economic activity that affects industrial commodity markets rather than the usual broader concept of real economic activity underlying real GDP or industrial output), and rpo_t refers to the real price of oil, as defined in section 3.1. The rea_t and rpo_t series are expressed in logs. The sample period is 1973.1-2005.9. In estimating the model, I allow for up to two years worth of lags.

3.3. Identifying Assumptions

Consider the structural VAR representation

$$(1) \quad A_0 z_t = \alpha + \sum_{i=1}^{24} A_i z_{t-i} + \varepsilon_t,$$

where ε_t denotes the vector of serially and mutually uncorrelated structural innovations. I postulate that A_0^{-1} has a recursive structure such that the reduced form errors e_t can be decomposed according to $e_t = A_0^{-1} \varepsilon_t$.

$$e_t \equiv \begin{pmatrix} e_t^{\Delta prod} \\ e_t^{rea} \\ e_t^{rpo} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{pmatrix} \varepsilon_t^{oil \text{ supply shock}} \\ \varepsilon_t^{aggregate \text{ demand shock}} \\ \varepsilon_t^{oil-specific \text{ demand shock}} \end{pmatrix}$$

The restrictions on A_0^{-1} may be motivated as follows: Crude oil supply shocks (referred to as *oil supply shocks* for short) are assumed not to respond to innovations to the demand for oil within the same month. That exclusion restriction is plausible because, in practice, oil-producing countries will be slow to respond to demand shocks, given the costs of adjusting oil production and the uncertainty about the state of the crude oil market. Thus, only persistent increases in demand are likely to prompt an increase in crude oil supply, if at all.

Innovations to global real economic activity that cannot be explained based on crude oil supply shocks will be referred to as shocks to the global demand for industrial commodities (or

aggregate demand shocks for short). This interpretation amounts to imposing the exclusion restriction that increases in the real price of oil driven by shocks that are specific to the oil market will not lower global real economic activity immediately, but with a delay of at least a month. This restriction is consistent with the sluggish behavior of global real economic activity after each of the major oil price increases in the sample.

Finally, innovations to the real price of oil that cannot be explained based on oil supply shocks or aggregate demand shocks by construction will reflect changes in the demand for oil as opposed to changes in the demand for all industrial commodities (referred to as *oil-specific demand shocks* for short). The latter structural shock will reflect in particular fluctuations in precautionary demand for oil driven by fears about the availability of future oil supplies, but it potentially could also reflect other factors such as oil sector-specific changes in inventory policies. As I will discuss below, there is no empirical evidence that exogenous changes to crude oil inventory policies are driving this shock, but the timing of the estimated shocks is broadly consistent with the precautionary demand interpretation.⁸

3.4. Empirical Results

The reduced form model is consistently estimated by the least-squares method. The resulting estimates are used to construct the structural VAR representation of the model. Inference is based on a recursive-design wild bootstrap with 2,000 replications (see Gonçalves and Kilian 2004).

3.4.1. How Do Global Oil Production, Real Economic Activity and the Real Price of Oil Respond to Demand and Supply Shocks in the Crude Oil Market?

Figure 5 shows the responses of global oil production, real economic activity and the real price of oil to one-standard deviation structural innovations. All shocks have been normalized such that an innovation would tend to raise the price of oil.

An unexpected oil supply disruption causes a sharp decline in global oil production upon impact, followed by a partial reversal of that decline within the first year. This pattern is consistent with the view that oil supply contractions in one region tend to trigger production increases elsewhere in the world. At the same time, this shock triggers a small, transitory and

⁸ The implicit assumption that there are no precautionary demand shocks in industrial commodities other than crude oil is consistent with the common view that crude oil is a unique commodity from the point of view of the oil-importing countries. I also abstract from idiosyncratic shocks to the demand or supply of dry cargoes. These shocks are presumed to average out in the construction of the index of real economic activity.

partially statistically significant increase in the real price of oil for about a year. It also causes a small temporary reduction of real economic activity that is partially statistically significant.

The effect of an unanticipated aggregate demand expansion in global commodity markets on global real economic activity is very persistent and lasts almost four years, before leveling off. It remains statistically significant at the 10 percent level for the first two years. Aggregate demand expansions temporarily increase global oil production. There is a delay of half a year before production expands. The production response peaks about one year after the shock and is statistically significant. There is some indication that the initial increase in world crude oil production is subsequently offset by persistent decreases. Aggregate demand expansions also cause a large and persistent increase in the real price of oil. The response of the real price of oil is significant at the 10 percent level for all horizons beyond six months, and at the 5 percent level starting in second year following the shock.

Oil-market specific demand increases have an immediate and persistent positive effect on the real price of oil that is highly significant for the first year. They also are associated with a temporary increase in real economic activity after nine months and a temporary decline after two years. The former is mostly statistically insignificant at the 5 percent level, whereas the latter is significant for half a year. Oil-market specific demand increases do not cause an increase in global oil production. In fact, there is evidence of a subsequent decline in crude oil supply, although that decline is comparatively small.

Perhaps the most striking result in Figure 5 is the fact that oil supply disruptions have only a small positive effect on the real price of oil. Part of the explanation is that oil supply disruptions in one region countries tend to trigger endogenous expansions of crude oil production elsewhere in the world that help offset the initial production shortfall. The small response of the real price of oil is also consistent with related evidence that oil supply shocks have little systematic predictive power for the changes in the real price of oil (see Kilian 2006a); yet it raises the question of what – if not crude oil supply disruptions – accounts for the apparent large increases in the real price of oil following major exogenous political events in the Middle East. Figure 5 suggests that the answer lies in sharp increases in precautionary demand. As shifts in precautionary demand are ultimately driven by expectations about the availability of future oil supplies and such expectations can change almost instantaneously in response to exogenous political events, such shocks tend to trigger an immediate and sharp increase in the real price of

oil (see Figure 5). This explanation will be explored further in the next subsection using historical decompositions of the real price of oil.

3.4.2. What Is the Cumulative Effect of Structural Shocks on the Real Price of Oil?

Figure 6 plots the respective cumulative contribution of each structural shock to the real price of oil based on a historical decomposition of the data. The first panel shows that oil supply shocks historically have made comparatively small contributions to the real price of oil. By far the biggest contributions are due to the aggregate demand shock and the oil-market specific demand shock. Whereas the aggregate demand shock caused long swings in the real price of oil, the oil-market specific demand shock is responsible for fairly sharply defined increases and decreases in the price of oil. This fact is consistent with the view that precautionary demand shocks may reflect rapid shifts in the market's assessment of the availability of future oil supplies. Although my econometric methodology does not allow me to identify the expectations component in oil-market specific demand shocks, the timing of the estimated oil-market specific demand shocks lends further credence to their interpretation as precautionary demand shocks, as discussed below.

It is instructive to focus on specific episodes. For example, the rapid rise in the real price of oil in late 1979 and 1980 after the Iranian Revolution appears to be driven mainly by the superimposition of a sharp increase in precautionary demand in 1979 on a slower-moving strong increase in real economic activity that started two years earlier. While the cumulative effect of oil-market specific demand peaked prior to the outbreak of the Iran-Iraq war and slowly subsided in the early 1980s, real economic activity continued to sustain the real price of oil well into the early 1980s. Throughout this period, oil supply shocks only served to amplify some of the short-run dynamics of the real price of oil, sometimes raising the price of oil and lowering it at other times. The increased importance of oil-market specific demand shocks starting in 1979 is consistent with an increase in precautionary demand. 1979 not only was the year of Khomeini's arrival in Iran, but of the Iranian hostage crisis and of the Soviet invasion of Afghanistan, all of which raised persistent fears of a regional war and the destruction of oil fields in Iran and Saudi Arabia. These fears were further fueled by the outbreak of the Iran-Iraq War in 1980.

The sharp fall in the real price of oil following the collapse of the OPEC cartel in late 1985 appears to be due more to a decline in oil-market specific demand than the direct effect of

the increase in Saudi oil production in the second panel or the fall in real economic activity in the third panel. The sudden abundant availability of crude oil and the perception that the breakdown of OPEC was irreversible are likely to have sharply lowered precautionary demand at this point. This sharp drop was partially reversed in 1987, amid attempts by OPEC to reunite. Similarly, the sharp spike in the real price of oil in 1990/91 after the invasion of Kuwait is almost entirely due to an increase in precautionary demand.⁹ Oil supply disruptions also had some effect on the real price of oil in the early 1990s, but their effect was small and persistent.

The disproportionate reduction in oil-market specific demand during the Asian crisis of 1997/98, when the real price of oil fell to an all-time low, suggests that at this point precautionary demand all but vanished.¹⁰ When the price of oil reached an all-time low in real terms, precautionary demand motives all but vanished. This effect was gradually reversed after oil prices recovered starting in 1999. Interestingly, the sharp rise in the real price of oil after 2000 is not driven by global aggregate demand or by the efforts of OPEC to coordinate production, but again by factors specific to the demand for crude oil. The most striking observation in Figure 6 is that the rise in the real price of oil since early 2002 is almost entirely due to a surge in real economic activity that started around 2001. There is no evidence that this price increase is driven either by precautionary demand or by oil supply shocks.

The evidence in Figure 6 clearly suggests that not all oil price shocks are alike. There are important differences in the relative contribution of the three structural shocks to the real price of oil between the Iran-Iraq War and the Iranian Revolution, for example, or between the Persian Gulf War and the period following the Iraq War and the civil unrest in Venezuela.

3.4.3. The Role of Precautionary Demand

One of the most important findings of the preceding subsection has been the disproportionate importance of oil-market specific demand shocks for the real price of oil. My results paint a very different picture of how exogenous political events in the Middle East affect the real price of oil than postulated in the existing literature. The traditional approach has been to quantify exogenous variation in actual crude oil production in OPEC countries and to relate this variation to changes in the price of crude oil (see, e.g., Hamilton 2003; Kilian 2006a). That approach fails

⁹ Kilian (2006a) arrives at the same conclusion using a different methodology.

¹⁰ As discussed below, it can be shown that this drop in oil-market specific demand pre-dates the sharp drop in oil inventories in 1999/2000 and hence is unlikely to be driven by inventory adjustments. Rather inventory policies seem to have changed in response to falling oil prices.

to capture shifts in market expectations that are not reflected in observed changes to crude oil production. Not surprisingly, as has been noted by Barsky and Kilian (2002), production-based accounts of oil price shocks do not match up well with the timing of oil price changes and with historical accounts of the crude oil market during oil crises such as the Iranian Revolution.

The results in this paper, in contrast, suggest that the most important channel by which exogenous events such as wars or revolutions affect the real price of oil is through their effect on precautionary demand for oil. The latter channel can produce immediate and potentially large effects on the real price of oil, even when crude oil production has not changed. It also can amplify the effects of shocks to crude oil production in anticipation of future changes to crude oil production. This point has been recognized for a long time, but it has never been quantified before, the fundamental difficulty being that expectations shifts related to uncertainty about future oil supplies are not observable and not linearly related to observables.

My analysis goes a long way toward capturing these expectations, as they are reflected in oil-market specific demand shocks, but falls short of formally identifying an explicit times series of this expectations-driven component of oil demand. While it is beyond reasonable doubt that oil-market specific demand shocks near certain dates in the sample (such as the outbreak of the Persian Gulf War) reflect shifts in precautionary demand, in general, oil-market specific demand shocks may also reflect other factors such as exogenous changes in crude oil inventory policies. For example, it is sometimes argued that there was a tendency for oil companies to reduce costly crude oil inventories in the mid-1990s, which may have contributed to the decline in oil prices. Data on petroleum inventories in the OECD are available from the U.S. Department of Energy starting in January 1988. Not only is there no sustained shift toward lower inventories in the 1990s in these data, but there is no evidence that the unusually low inventories of 1996, 2000 or 2003 (relative to trend) were systematically followed by positive oil-market specific demand shocks. Rather the evidence is consistent with crude oil inventories falling in response to lower real oil prices, and rising after the real price of oil recovered in 1999/2000. Thus, we can rule out changing inventory policies in the 1990s as one of the chief driving forces for the oil-market specific demand shock, leaving shifts in the uncertainty of oil importers about the availability of future oil supplies as the likely driving force behind fluctuations in the oil-market specific demand shock.

While expectations about the availability of adequate future oil supplies are likely to be

driven by observables, that link is highly nonlinear and elusive, even when the determinants of concerns about future oil supply disruptions are well understood and can be measured, as illustrated in Kilian (2006c) using data on military attacks on ships in the Persian Gulf in the 1980s. When the source of increased anxieties about future oil supplies is diffuse, as in 1979/80, when the Iranian Revolution, the Iranian Hostage Crisis, the Soviet invasion of Afghanistan and the gradual transition of the crude oil market toward spot contracts all contributed to increased uncertainty, the effect on expectations will be even harder to isolate.¹¹

4. Understanding the Effects of Oil Price Disturbances on the U.S. Economy

4.1. Baseline Model

A question of considerable interest is how the structural innovations in model (1) relate to U.S. macroeconomic aggregates such as CPI inflation or real GDP growth. The main problem in answering this question is that the latter aggregate is not available at monthly frequency.¹² While one could construct an analogous structural VAR model (1) on data aggregated to quarterly frequency, at that frequency the identifying assumptions would no longer be credible. Instead I construct measures of the quarterly shocks by averaging the monthly structural innovations for each quarter:

$$\hat{\zeta}_{jt} = \frac{1}{3} \sum_{i=1}^3 \hat{\varepsilon}_{j,t,i}, \quad j = 1, \dots, 4,$$

where $\hat{\varepsilon}_{j,t,i}$ refers to the estimated residual for the j th structural shock in the i th month of the t th quarter of the sample.

Under the identifying assumption that within a given quarter there is no feedback from Δy_t and π_t to $\hat{\zeta}_{jt}$, $j = 1, \dots, 3$, these shocks can be treated as predetermined and we can examine their effects on U.S. macroeconomic aggregates based on the regressions:

$$(4) \quad \Delta y_t = \alpha + \sum_{i=0}^{12} \phi_i \hat{\zeta}_{jt-i} + u_t, \quad j = 2, \dots, 4$$

¹¹ An interesting question is the extent to which oil futures data which have existed since 1983 help capture shifts in expectations driven by precautionary demand. This is the subject of ongoing research.

¹² I do not use interpolated real GDP growth data. One reason is that interpolation is known to cause spurious dynamics in general. In this case, in particular, there also is a second reason not to use interpolation. Standard interpolation methods use monthly data on industrial production to infer movements in real GDP within the quarter. Since industrial production is a measure of gross output that may behave very differently from real GDP (a measure of value added) in response to oil shocks, as discussed in Barsky and Kilian (2002), the concern is that interpolated “monthly” real GDP data will behave like industrial production, thus invalidating the analysis.

and

$$(5) \quad \pi_t = \delta + \sum_{i=0}^{12} \psi_i \hat{\zeta}_{jt-i} + v_t, \quad j = 2, \dots, 4$$

where u_t and v_t are potentially serially correlated errors. In this regression model the impulse response coefficients at horizon h correspond to ϕ_h and ψ_h , respectively. Thus, the number of lags is determined by the maximum horizon of the impulse response function, which is set to 12 quarters. In conducting inference on the response estimates implied by (4) and (5), the presence of serial correlation in the error term is dealt with by using block bootstrap methods.¹³

By allowing for feedback from lagged U.S. macroeconomic aggregates to the shock series, this approach recognizes that the shocks $\hat{\zeta}_{jt}$, $j = 1, 2, 3$, are by construction strictly exogenous with respect to the information set in model (1), but are not necessarily strictly exogenous with respect to U.S. macroeconomic aggregates. All we impose is the assumption that all $\hat{\zeta}_{jt}$ are predetermined in that they do not respond within the quarter to innovations in U.S. real GDP growth or CPI inflation. Although the assumption of predeterminedness is not testable, there is some evidence that this assumption is not at odds with the data. Table 1 shows estimates of the correlation of $\hat{\zeta}_{jt}$ with reduced form autoregressive residuals for U.S. real GDP growth and CPI inflation. The results shown allow for eight autoregressive lags. Qualitatively similar results are obtained with four lags. Table 1 shows that in most cases these residuals are nearly uncorrelated within the quarter with $\hat{\zeta}_{jt}$. Of particular interest are the correlations for $j = 2$, which help dismiss the potential concern that domestic innovations to U.S. real GDP growth (or CPI inflation) may be positively correlated with aggregate demand shocks in global industrial commodity markets, which would invalidate the assumption that $\hat{\zeta}_{2t}$ is predetermined with respect to U.S. macroeconomic aggregates.

4.2. Does it Matter for U.S. Macroeconomic Performance Why the Price of Oil Increased?

4.2.1. The Dynamic Effects of Oil Demand and Oil Supply Shocks on Real GDP Growth and CPI Inflation

¹³ All results shown below are for block size 4 and 20,000 bootstrap replications. Results based on block size 8 are very similar. It is important to note that these confidence intervals do not account for the fact that the residuals used in the regressor matrix are generated regressors. Controlling for this problem is complicated by the fact that the regression procedure involves data at both monthly and quarterly frequency.

Figure 7 summarizes the responses of the growth rate and level of U.S. real GDP to each of the three shocks defined earlier. Unanticipated oil supply disruptions significantly lower real GDP growth on impact. The impact response is statistically significant at the 5 percent level. At longer horizons, there is only weak evidence of a decline. The response of real GDP is negative at all horizons, but significant at the 5 percent level only for the first two years.

An aggregate demand expansion causes an initial increase in real GDP growth in the first year, followed by a decline below the initial level in the second year. In the third year, the response reverts to zero. Both the initial increase and the decline in real growth are partially significant at the 10 percent level. There is a marginally statistically significant increase in the level of real GDP for the first three quarters, but the subsequent fall in real GDP below the initial level is not significant at the 10 percent level.

Oil-market specific demand increases cause a sustained reduction in real GDP growth that is significant at the 10 percent level at most horizons. At some horizons the response is significant even at the 5 percent level. The level of real GDP declines throughout. That decline is significant at the 5 percent level for all horizons but the first three quarters.

Figure 8 shows the corresponding results for U.S. CPI inflation and the level of the CPI. There is some evidence that oil supply disruptions raise inflation temporarily, but not by much. The peak response occurs in the first quarter, but is barely significant at the 10 percent level. There is no significant increase in the level of consumer prices at any horizon at the 10 percent level. In fact, the price level falls in the long run. An aggregate demand expansion causes an only partially significant, but persistent increase in U.S. CPI inflation, starting in the first quarter. In contrast, the effect on the price level is significant at the 10 percent level after six quarters. Again the most clear-cut results are obtained for an oil-market specific demand increase. There is an immediate, large and persistent increase in CPI inflation that is mostly significant at the 5 percent level for the first two years. The effect on the CPI is significant at the 5 percent level for all horizons.

The results in Figures 7 and 8 illustrate important differences in how the oil demand and oil supply shocks underlying the real price of oil affect U.S. macroeconomic aggregates. The chief results can be summarized as follows: The only shock to lower U.S. real GDP growth significantly at the 5 percent level for extended periods is an oil-market specific increase in demand. Oil supply disruptions cause a temporary decline in real GDP in the first two years that

is marginally significant at the 10 percent level. Aggregate demand expansions significantly raise real GDP in the short run at the 10 percent significance level, but have no significant long run effects. Whereas global aggregate demand increases raise the price level in the long run, oil-market specific demand increases raise the price level immediately and persistently. Only the latter response is significant at the 5 percent level. Disruptions in oil supply have no significant effect on the price level in the first year, but there is some evidence of a long run decline.

4.2.2. Which Shocks Cause Stagflation?

A common concern is that higher oil prices are responsible for stagflation. Again the answer may depend on which type of shock is driving the real price of oil. In this subsection, I address this question based on a statistical measure of conditional co-movement developed by Den Haan (2000). In Figure 9, this measure is applied to the responses of CPI inflation and real GDP growth to each of the demand and supply shocks in the crude oil market, allowing us to assess which – if any – of these shocks have stagflationary effects. Following Den Haan and Summer (2004, p. 1340), the plot shows conditional covariances rather than conditional correlations. This normalization facilitates a comparison of the statistic across horizons. The conditional covariance at horizon h is constructed as

$$C(h) = \Delta y_h^{imp} \pi_h^{imp}$$

where z_h^{imp} denotes the response of variable z_t at horizon h to a given structural innovation in equations (4) and (5), respectively (see Den Haan 2000, p. 8). Stagflation in the form of rising prices and falling output means that this measure will be negative. It is natural to conduct a one-sided test of the null of zero conditional covariance against the stagflationary alternative. Figure 9 plots 80 percent and 90 percent bootstrap confidence intervals along with the point estimates. The coverage rates are chosen such that the rejection probability in the upper tail corresponds to 10 percent and 5 percent, respectively.

Figure 9 shows that it makes a difference which shock is driving the macroeconomic aggregates. The point estimate of the response to oil supply shocks shows virtually no evidence of stagflation. The estimated response to a global aggregate demand shock shows some stagflationary effects between four and eight quarters after the shock. Lastly, the estimated response to an oil-market specific demand shock is stagflationary on impact and continues to be stagflationary to varying degrees across all horizons.

The confidence intervals indicate that these responses are very imprecisely estimated as a rule. There is virtually no evidence that the response to oil supply shocks is stagflationary at the usual significance levels. The response to a global aggregate demand shock is marginally significant after eight quarters at the 10 percent level. The response to oil-market specific demand shocks is the only response to be partially significant at the 5 percent level. The most significant impact occurs after four quarters with some additional less significant effects at other horizons. Thus, 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.

4.2.3. How Much of U.S. Real GDP Growth and CPI Inflation Can Be Attributed to Each Shock?

While impulse responses are informative about the dynamic effects of a one-time shock, an equally important question is how much of U.S. economic performance can be attributed to the cumulative effect of each shock. The first panel of Figure 10 shows that oil supply disruptions lowered U.S. real growth in the early 1980s, while raising it in the mid-1980s. There is no evidence that oil supply shocks have lowered U.S. real growth since the 2003 Iraq War.

The second panel suggests that the U.S. economic expansion in 1978/79 was fueled in part by global aggregate demand shocks, as was the subsequent contraction in the early 1980s. Global aggregate demand shocks also contributed to the slump in U.S. real growth in 1989-1991. Interestingly, while global aggregate demand shocks raised U.S. real growth after the Asian Crisis, their cumulative effect has been to lower U.S. real GDP growth in 2004/05.

The last panel shows that oil-market specific demand shocks played an important role in lowering U.S. real GDP growth from 1979 through 1982. They also caused a sustained reduction in real growth in 2001/02. On the other hand, they did not have much effect following the invasion of Kuwait in 1990, and they raised real growth in the late 1980s, in 1994, and after 2003.

Figure 11 shows the corresponding results for U.S. CPI inflation. The first panel shows that oil supply shocks raised CPI inflation in 1979-1982, but lowered it in 1983-1985. They also contributed to higher inflation in 1990-92. Most recently, having caused lower inflation between 2000 and 2004, their effect has been close to zero. The second panel shows large cumulative effects from aggregate demand shocks on U.S. CPI inflation, especially in 1979-1982 and from 2003 onwards. Even stronger effects arise from oil-market specific demand shocks. There is

evidence of a sharp increase in inflation in 1979 that peaks in 1980 and persists into the early 1980s. There also is evidence of a sharp decline in CPI inflation in 1986 and again in 1993/94 and during the Asian Crisis. While oil-market specific demand shocks contributed to higher inflation in 2000-2002, their effect more recently has been to lower CPI inflation.

4.2.4. To What Extent Is U.S. Macroeconomic Performance Driven by the Real Price of Crude Oil as Opposed to Domestic Policies? External Factors versus Domestic Factors

The evidence in Figures 10 and 11 raises the question of how much of U.S. macroeconomic performance since the 1970s must be attributed to the external shocks driving the real price of oil and how much must be attributed to domestic shocks. Figure 12 attempts to answer this question by plotting the sum of the cumulative responses shown earlier (referred as the “cumulative effect” of external factors) against the actual (demeaned) realizations of U.S. real growth and CPI inflation. The difference between the two lines is the component of each series that must be attributed to domestic factors and policies. In essence, Figure 12 answers the question of how U.S. real growth and CPI inflation would have evolved, if real oil prices had not moved. As noted earlier, this analysis is based on the crucial assumption that all shocks that drive the real price of oil are predetermined relative to U.S. real GDP growth and CPI inflation.

The first panel of Figure 12 suggests that much (but not all) of the decline in real growth in the early 1980s and its recovery in the early 1980s was externally caused. The brief recovery of real growth in 1981 and the excess growth of 1983/84 were driven by domestic factors. There also is evidence that external factors contributed substantially to the 1990/91 contraction and the subsequent recovery. However, an equally important part of the decline in real GDP growth in 1991 was due to domestic shocks. Moreover, there is evidence of a decline in real growth driven by domestic factors in 2004.

The second panel shows that a substantial part of the high CPI inflation rates between late 1978 and mid-1980 had nothing to do with external factors, although external factors became an increasingly important contributor over time. Whereas in late 1978, essentially zero percent of U.S. inflation was due to external factors, by early 1980 they accounted for about half of the observed inflation rate. The domestic component of inflation was essentially wiped out later in 1980. This evidence is consistent with the view that this component was monetary in origin, and was erased by a shift in monetary policy under Volcker. From 1981 through 1983 the decline in U.S. inflation was mainly driven by external factors. As the effect of these factors declines, so

does inflation. Starting in late 1982, external factors had deflationary effects, allowing domestic policy to pursue a more gradualist and more inflationary policy, which may explain the excess growth for the same time period in the first panel. A similar tendency to smooth inflation rates over time can be observed after the Persian Gulf War. Between 1992 and 1999 external effects on inflation were deflationary. Based on this evidence, it seems fair to say that the Federal Reserve's ability to lower inflation was helped by external factors in 1982-83 and 1991-93. Figure 12 also suggests that much of the build-up in U.S. CPI inflation in late 2005 can be attributed to shocks driving the oil market, whereas real growth has been remarkably stable and seemingly unaffected by external developments since 2004.

5. Conclusion

While it is uncontroversial that the real price of oil is determined in global markets, global data on many of the key determinants of the real price of oil are hard to come by. This paper has explored some novel approaches to circumventing these difficulties. Some of the analysis therefore is necessarily tentative, and subject to reexamination using alternative measures of global real economic activity. Nevertheless the current analysis demonstrates that the traditional approach to thinking about oil price changes and oil price shocks must be rethought.

My analysis suggests that shocks to the production of crude oil, while not trivial, are far less important in understanding changes in the real price of oil than shocks to aggregate demand and shocks to the precautionary demand for oil that reflect fears about future oil supplies. This result would continue to hold if oil supply shocks were further decomposed into exogenous political oil supply shocks and other oil supply shocks (see Kilian 2006c). While obviously the details of my results may be affected by changes in the data set, the fact that aggregate demand shocks and oil-market specific demand shocks (in particular in the form of shocks to precautionary demand for oil) are central to understanding changes in the real price of oil is likely to survive.

The central message of this paper is that oil price increases may have very different effects on the real price of oil, depending on the underlying cause of the price increase. For example, an increase in precautionary demand for crude oil causes an immediate, persistent and large increase in the real price of crude oil, an increase in aggregate demand for all industrial commodities causes a delayed, but sustained increase in the real price of oil that is also substantial; crude oil production disruptions cause a small and transitory increase in the real price

of oil within the first year.

Historical decompositions of fluctuations in the real price of oil suggest that oil price shocks historically have been driven mainly by a combination of aggregate demand shocks and precautionary demand shocks, rather than oil supply shocks, as is commonly believed. For example, the increase in the real price of oil after 1979 was driven by the superimposition of strong global demand and a sharp increase in precautionary demand in 1979 with only minor contributions from oil supply shocks. Likewise the build-up in oil prices after 2003 was driven entirely by the cumulative effects of positive global demand shocks.

These findings have important implications for thinking about the effects of oil price changes on the U.S. economy. As the paper demonstrated, an oil price change driven by a global aggregate demand shock, for example, will have a very different effect than an oil price change driven by an increase in precautionary demand driven by fears about future oil supply shortfalls. Moreover, only precautionary demand shocks are likely to generate stagflation in the U.S. economy. It therefore is important that we understand the extent to which increases in the real price of oil are driven by one shock or another, before formulating appropriate policy responses.

Finally, my analysis provided a comprehensive assessment of the extent to which shocks in the crude oil market contributed to U.S. real GDP growth and CPI inflation since 1978. A historical decomposition of the data revealed, for example, that a substantial part of the high U.S. inflation rates between late 1978 and late 1980 had nothing to do with external shocks, although external shocks became an increasingly important contributor over time. They also show that the disinflation policy under Paul Volcker was greatly facilitated by falling oil prices in the early 1980s.

References

- Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.
- Barsky, R.B., and L. Kilian (2004), "Oil and the Macroeconomy Since the 1970s," *Journal of Economic Perspectives*, 18(4), 115-134.
- Bernanke, B.S. (2004), "Oil and the Economy", Remarks by Governor Ben S. Bernanke at the Distinguished Lecture Series, Darton College, Albany, Georgia, October 21.
- Cooley, T.F., and S. LeRoy (1985), "Atheoretical Macroeconometrics: A Critique," *Journal of Monetary Economics*, 16, 283-308.
- Darmstadter, J., and H.H. Landsberg (1976), "The Economic Background," in R. Vernon (ed.), *The Oil Crisis*, Norton: New York.
- Den Haan, W.J. (2000), "The Comovement Between Output and Prices," *Journal of Monetary Economics*, 46, 3-30.
- Den Haan, W.J., and S.W. Summer (2004), "The Comovement between Real Activity and Prices in the G7," *European Economic Review*, 48, 1333-1347.
- Drewry Shipping Consultants Ltd, *Shipping Statistics and Economics*, monthly, various issues since 1970.
- Gonçalves, S., and L. Kilian (2004), "Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form," *Journal of Econometrics*, 123, 89-120.
- Hamilton, J.D. (2003), "What is an Oil Shock?" *Journal of Econometrics*, 113, 363-398.
- Isserlis, L. (1938), "Tramp Shipping Cargoes and Freights," *Journal of the Royal Statistical Society*, 101, 53-134.
- Kilian, L. (2006a), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" mimeo, Department of Economics, University of Michigan, <http://www-personal.umich.edu/~lkilian/paperlinks.html>.
- Kilian, L. (2006b), "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries," mimeo, Department of Economics, University of Michigan, <http://www-personal.umich.edu/~lkilian/paperlinks.html>.
- Kilian, L. (2006c), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," mimeo, Department of Economics, University of Michigan, <http://www-personal.umich.edu/~lkilian/paperlinks.html>.

Klovland, J.T. (2004), "Business Cycles, Commodity Prices and Shipping Freight Rates: Some Evidence from the pre-WWI Period," paper presented at the *Workshop on Market Performance and the Welfare Gains of Market Integration in History*, European University Institute, Florence, Italy, July 1-4, 2004.

Oil and Gas Journal, various issues since 1970.

Skeet, I. (1988), *OPEC: Twenty-Five Years of Prices and Politics*, Cambridge University Press: Cambridge, UK.

Stopford, M. (1997), *Maritime Economics*. 2nd ed., Routledge: London, U.K.


Tinbergen, J. (1959), "Tonnage and Freight," in *Jan Tinbergen Selected Papers*, North Holland: Amsterdam, 93-111.

**Table 1. Correlation of $\hat{\xi}_{jt}$ with Autoregressive Residuals for
U.S. Real GDP Growth and CPI Inflation**

	Oil Supply Shock	Aggregate Demand Shock	Oil-Market Specific Demand Shock
U.S. Real GDP Growth	0.146	0.055	-0.035
U.S. CPI Inflation	0.081	-0.128	0.243


NOTES: Autoregressive residuals for U.S. data based on AR(8) models.

**Figure 1: Real Price of Crude Oil
1973.1-2005.9**




NOTES: The oil price series is the refiner acquisition cost of imported crude oil as reported by the U.S. Department of Energy, extended from 1974.1 back to 1973.1 as in Barsky and Kilian (2002). The oil price has been deflated by the U.S. CPI. Shaded areas mark periods of oil supply management by U.S. oil companies (1973.1-1973.10) and by the OPEC cartel (1982.3-1982.12 and 1983.3-1985.10).

**Figure 2: Global Crude Oil Production
1973.1-2005.9**




NOTES: World crude oil production as reported by the U.S. Department of Energy. The measure of the exogenous crude oil production shortfall in OPEC countries is based on Kilian (2005, 2006). The percent change has been annualized.

**Figure 3: Index of Global Real Economic Activity based on Dry Cargo Bulk Freight Rates
1968.1-2005.9**



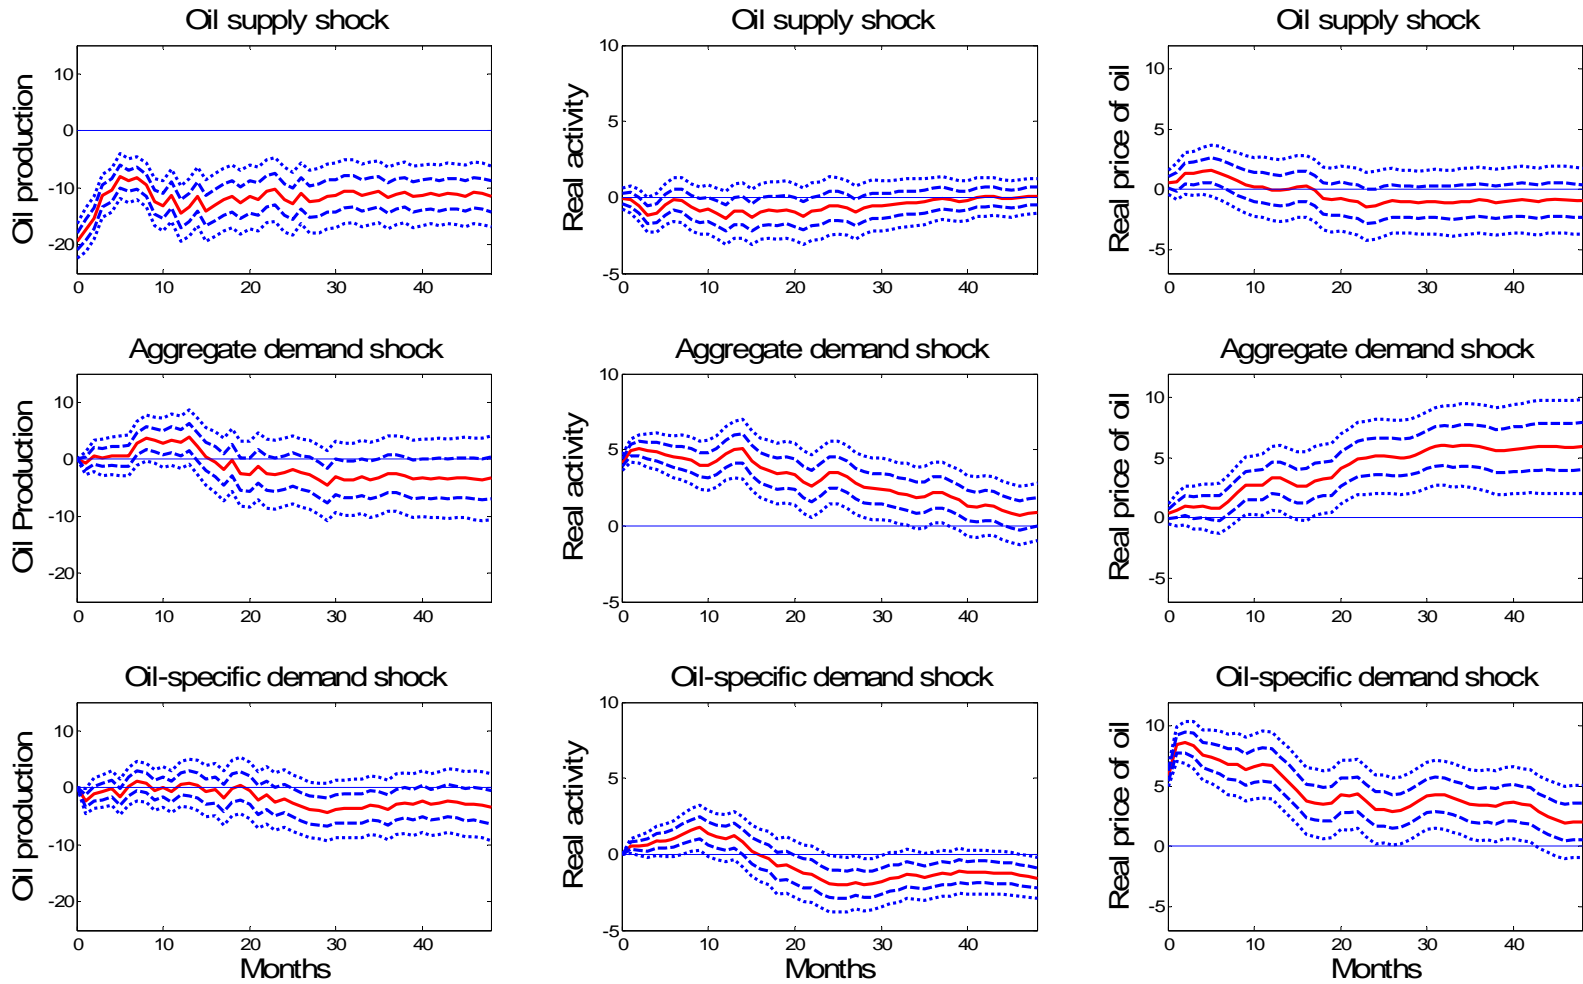
NOTES: The monthly raw data were manually collected from Drewry's Shipping Monthly, various issues since 1970. The two oldest series in the first panel are indices compiled by Drewry's. The remaining series are differentiated by cargo, route and ship size.

**Figure 4: Index of Global Real Economic Activity and Index of Real Bunker Fuel Prices
1970.1-1973.12**




NOTES: The bunker fuel rate data are from the *Oil and Gas Journal*, various issues since 1970. All rates refer to Bunker C fuel, as recorded for the Caribbean, the Gulf Coast and California. The index is based on equal-weighted growth rates, computed using observations for the last week of each month. The real economic activity index is based on Figure 3.

**Figure 5: Responses to One-Standard Deviation Structural Shocks
OLS Point Estimates with One and Two-Standard Error Bands**




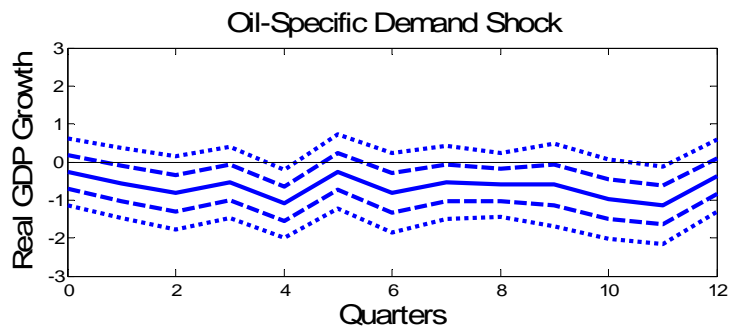




NOTES: Estimates based on restricted VAR(24) system described in text The confidence intervals were constructed using a recursive-design wild bootstrap.

**Figure 6: Historical Decomposition of Real Price of Oil
1975.2-2005.9**



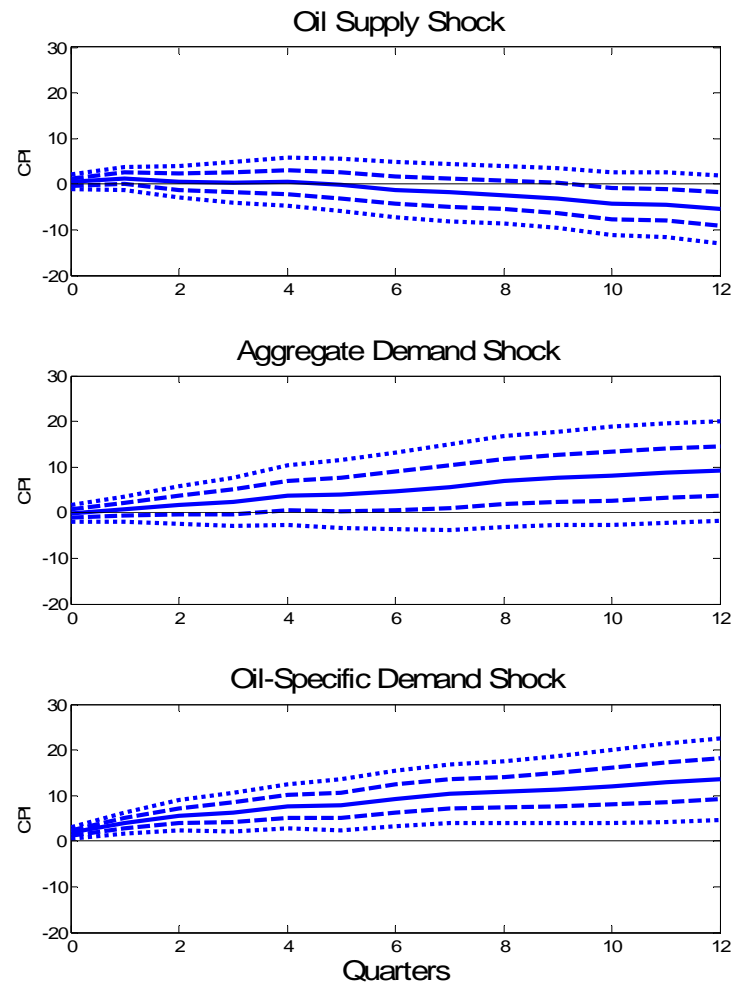
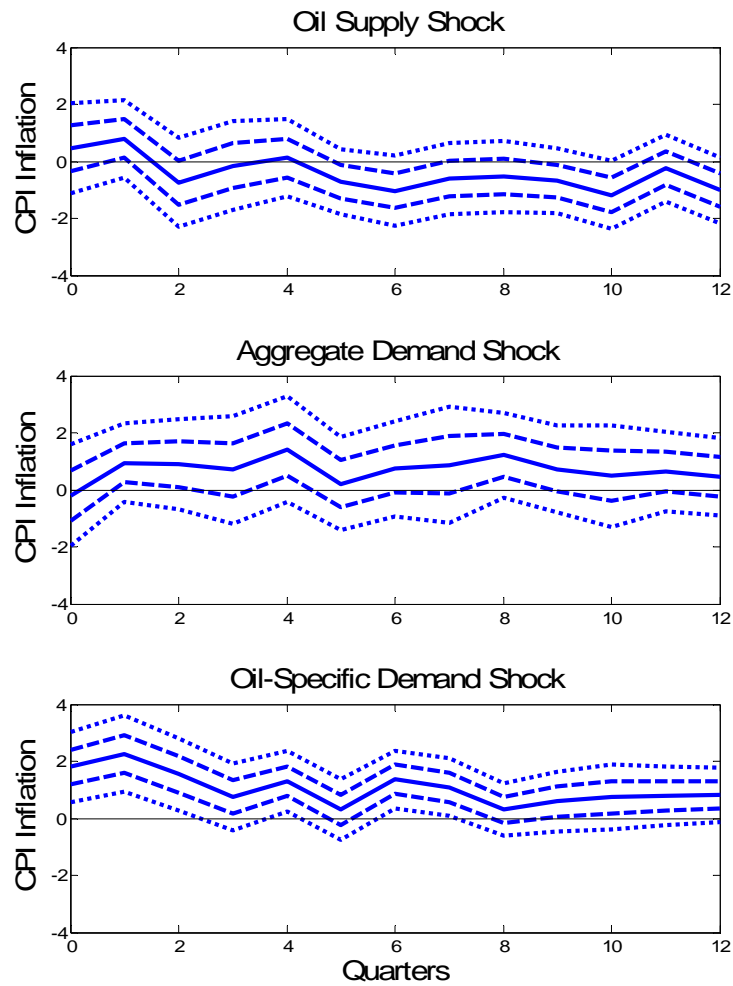
NOTES: See Figure 5.

**Figure 7: Responses of U.S. Real GDP Growth and Real GDP Level to Each Structural Shock
OLS Point Estimates with One and Two-Standard Error Bands**



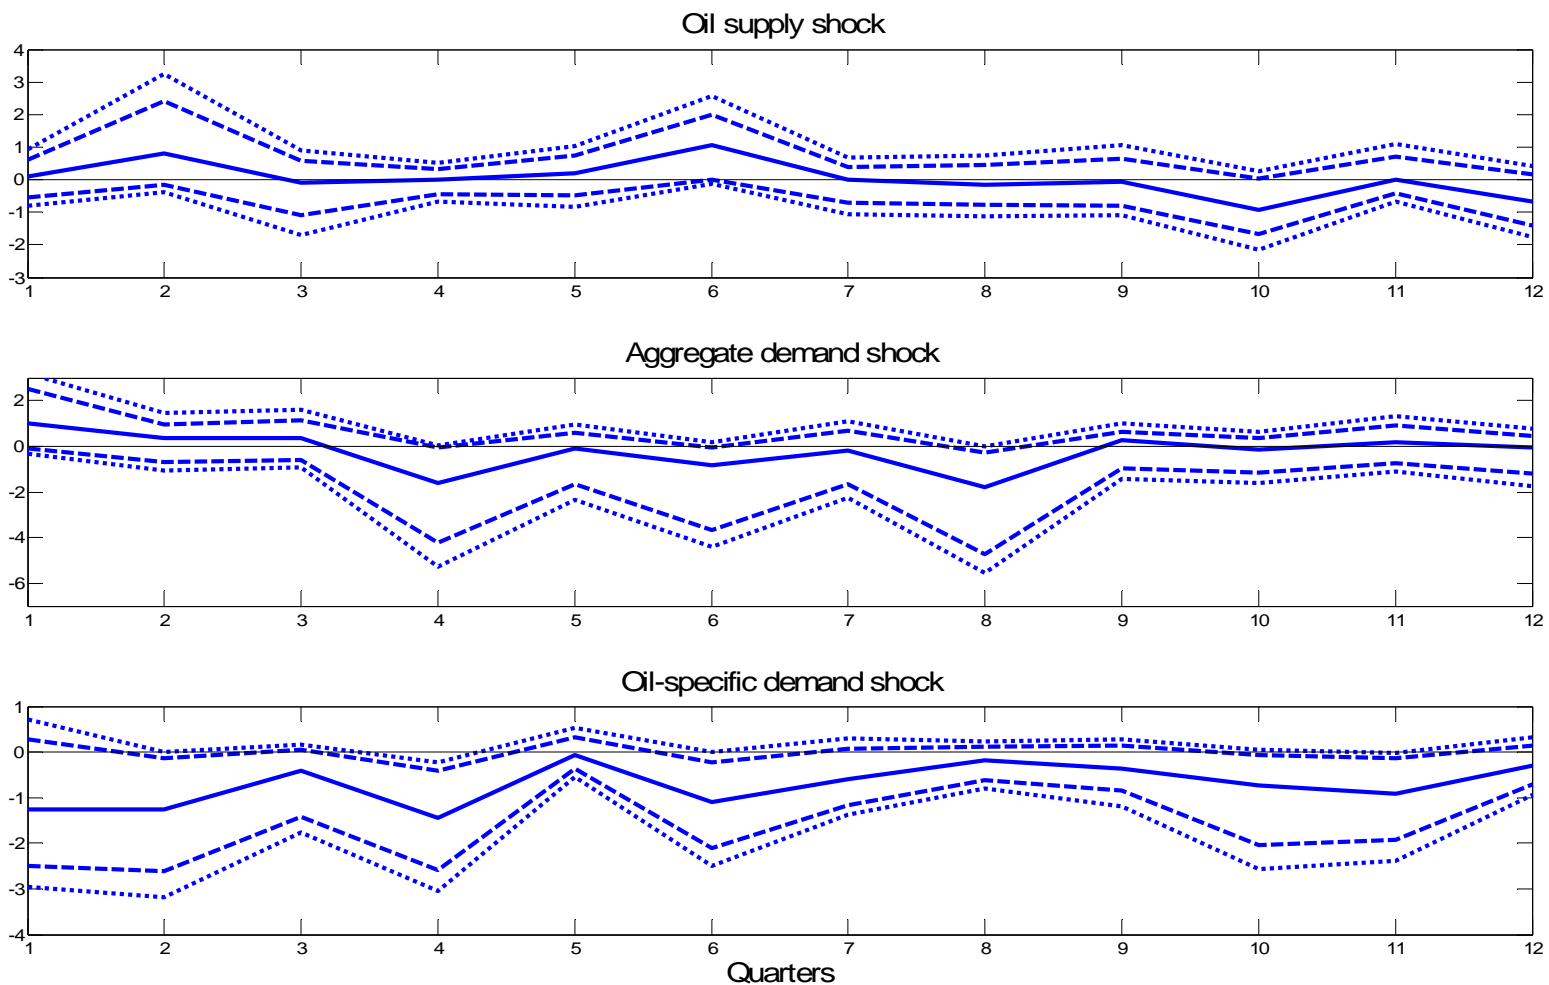
NOTES: All results are based on equations (4).

Figure 8: Responses of U.S. CPI Inflation and CPI Price Level to Each Structural Shock
OLS Point Estimates with One and Two-Standard Error Bands



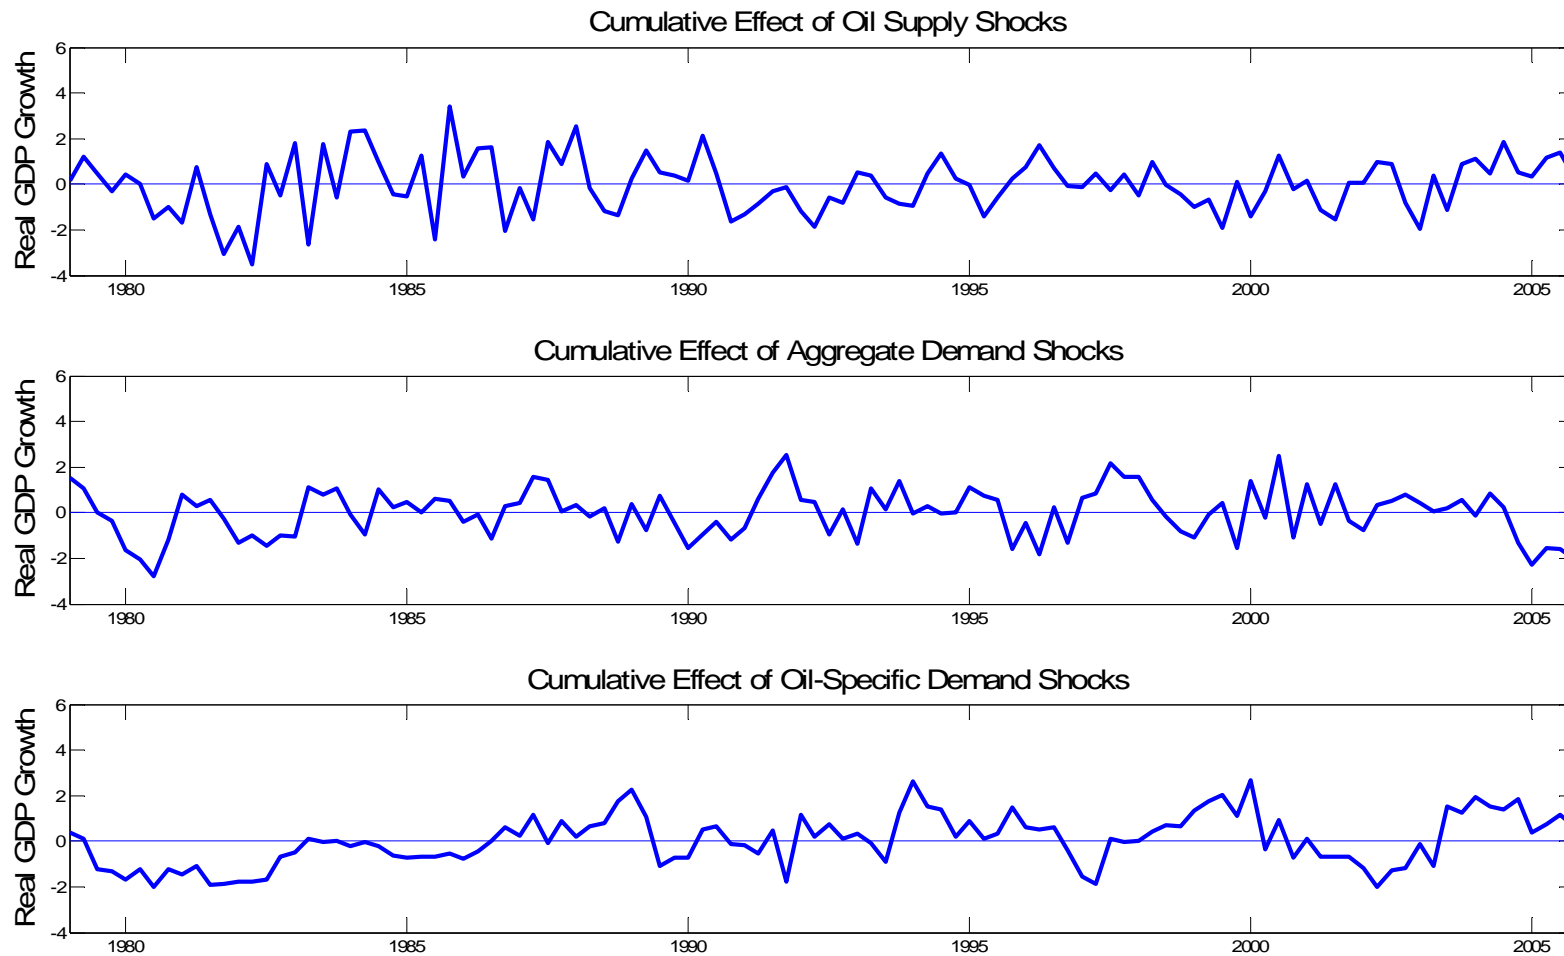
NOTES: All results are based on equation (5).

**Figure 9: Stagflationary Effects of a Unit Innovation in the Structural Shocks
Conditional Covariance of Inflation and Real GDP Growth with 80% and 90% Confidence Bands**




NOTES: The plot shows a statistical measure of the conditional co-movement between real GDP growth and CPI inflation, as defined in Den Haan (2000) and Den Haan and Summer (2004), based on the impulse responses in Figures 7 and 8. Stagflation in the form of rising prices and falling output means that this measure will be negative.

Figure 10: Cumulative Contribution of Each Structural Shock to U.S. Real GDP Growth 1978.IV-2005.III




NOTES: All results based on regression model (4). The effect is computed by simulating the predicted value of regression (4) for $x_t = 0 \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. All rates are annualized. The first three observations were discarded as transients.

**Figure 11: Cumulative Contribution of Each Structural Shock to U.S. CPI Inflation
1978.IV-2005.III**



NOTES: All results based on regression model (5). The effect is computed by simulating the predicted value of regression (5) for $x_t = 0 \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. All rates are annualized. The first three observations were discarded as transients.

Figure 12: Cumulative Contribution of All Structural Shocks Combined to U.S. Macroeconomic Real GDP Growth and CPI Inflation 1978.IV-2005.III



NOTES: See Figures 10 and 11.