DISCUSSION PAPER SERIES

No. 5994

NOT ALL OIL PRICE SHOCKS ARE ALIKE: DISENTANGLING DEMAND AND SUPPLY SHOCKS IN THE CRUDE OIL MARKET

Lutz Kilian

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP5994.asp**

ISSN 0265-8003

NOT ALL OIL PRICE SHOCKS ARE ALIKE: DISENTANGLING DEMAND AND SUPPLY SHOCKS IN THE CRUDE OIL MARKET

Lutz Kilian, University of Michigan and CEPR

Discussion Paper No. 5994
December 2006

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 5994

December 2006

ABSTRACT

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

Using a newly developed measure of global real economic activity, a structural decomposition of the real price of crude oil in four components is proposed: oil supply shocks driven by political events in OPEC countries; other oil supply shocks; aggregate shocks to the 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 quantifies the magnitude and timing of these shocks, their dynamic effects on the real price of oil and their relative importance in determining the real price of oil during 1975-2005. The analysis also sheds light on the origins of the major oil price shocks since 1979. 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. It is shown 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 mid-1970s has been determined by the external economic shocks driving the real price of oil as opposed to domestic economic factors and policies.

JEL Classification: E31, E32 and Q43
Keywords: oil supply, causality, oil demand and oil price shocks

Lutz Kilian
Department of Economics
University of Michigan
238 Lorch Hall
Ann Arbor, MI 48109-1220
USA
Email: lkilian@umich.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=132598

Submitted 09 November 2006

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 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 real price of oil to macroeconomic outcomes. Thus, we have to move beyond studying changes in the real price of oil and address the problem of identifying the structural shocks underlying the real 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 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 in four components is proposed: oil supply shocks driven by political events in OPEC countries; other oil supply shocks; shocks to the aggregate 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

[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.

demand associated with fears about 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 demand for oil (and thus on the price of oil) has never been estimated. The analysis in this paper sheds light on both of these questions.

The paper quantifies the magnitude and timing of each of these four shocks, their dynamic effects on the real price of oil and their relative importance in determining the real price of oil during 1975-2005. My analysis also 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 analysis also allows for the first time to quantify the extent to which U.S. real growth and CPI inflation since the mid-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 re-thought. 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 during the Persian Gulf War, for example, my analysis suggests that it is less the actual physical supply disruptions than the increased precautionary demand for oil triggered by fears about the availability of future oil supply shortfalls (which may or may not be realized) that is driving the price of oil.

The effect of these demand and supply shocks in global crude oil markets on U.S. real growth and CPI inflation differs substantially. For example, at the 10 percent significance level, global aggregate demand expansions tend to raise U.S. real GDP in the short run, while raising the U.S. CPI in the long run. In contrast, oil-market specific increases in demand persistently lower U.S. real GDP growth and persistently raise the U.S. price level at the same time. Political oil supply disruptions cause a significant decline in real GDP only in the long run, whereas other oil supply disruptions cause a significant decline of real GDP only in the short run. Other oil supply disruptions significantly lower the price level in the long run, whereas political oil supply shocks have no significant effect on the CPI level in the short or the long run.

While a substantial component of U.S. real growth and CPI inflation since the 1970s can be traced to these external shocks, an equally important component is associated with domestic 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 2005 was caused by the external shocks driving the oil market. There also is strong evidence that the disinflation of the early 1980s and of the 1990s was helped 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 during the months leading up to 2003 Iraq War.

While 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 price of oil –

---

[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.

far from being mainly controlled by cartels – is ultimately determined by market forces and 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. Exogenous Shocks to Crude Oil Production driven by Political Events in OPEC Countries

Crude oil is a unique commodity in that its production has at times been severely curtailed by political events such as wars or revolutions in the Middle East that are commonly considered exogenous with respect to global macroeconomic conditions. In assessing the effects of demand and supply shocks on the real price of oil it is important to control explicitly for these shocks to crude oil production. Figure 2 plots a recently proposed measure of the shocks to OPEC crude oil production driven by six major political events (indicated by vertical lines in the plot) due to Kilian (2006a). That measure distinguishes between changes in crude oil production that operate directly in response to the exogenous event in question (such as the temporary production increases in Saudi Arabia during the Persian Gulf War or the Iraqi production increase following the Iranian Revolution) on the one hand and endogenous changes in crude oil production in response to higher oil prices on the other hand. The latter are treated as part of the propagation mechanism to be estimated, whereas the former are explicitly incorporated into the counterfactual underlying the construction of this series. The reader is referred to Kilian (2006a, 2006b) for further discussion of the derivation of this time series.

2.3. Global Oil Production

Another important determinant of the real price of oil is global crude oil production as reported by the U.S. Department of Energy. Figure 3 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 the political supply shocks described in Figure 2 as well as the cartel activities discussed in section 2.1., internal OPEC power struggles in the 1970s and early 1980s, and endogenous responses to changes in the real price of oil in non-OPEC countries (see, e.g., Skeet 1988). They also reflect the effects of Hurricanes Rita and Katrina on crude oil production. The latter events have been excluded from the series shown in Figure 2 because they took place outside of OPEC and were not driven by political events.

2.4. Global Real Economic Activity

The price of oil is determined in global markets. 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 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 properly weight 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.

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.[3] 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

[3] 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.

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 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 net increases in freight rates relative to the recent past 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.[4] 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 4 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 4 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 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 4.[5]

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 4. 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.

---

[4] 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.
[5] 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.

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 since 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 4 confirms this fact. 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 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 4 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 5). Second, Figures 2 and 4 show 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 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 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.

An obvious question is why this type of decomposition has not been attempted before. My identification of the structural shocks that drive the price of oil rests on the availability of two time series: One is a recently proposed measure of OPEC oil supply shocks that can be attributed to exogenous political events such as wars or revolutions in the Middle East (see Kilian 2006a, 2006b). The other is a new time series, specifically constructed for this paper, intended to capture fluctuations in global real economic activity, as they relate to the market for industrial commodities. While these series have been carefully constructed, obviously it will be possible to construct alternative measures of global real economic activity or to explore different measures of politically driven exogenous oil supply shocks. 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 near-VAR model based on monthly data for $z_t = (x_t, \Delta prod_t, rea_t, rpo_t)'$, where x_t denotes the series of oil supply shocks driven by exogenous political events in OPEC countries, $\Delta prod_t$ 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 defers to the real price of oil, as defined in section 3.1. The rea_t and rpo_t series are expressed in logs.

In estimating the model, I allow for up to two years worth of lags. Consider the structural representation

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

where the first row of A_i, $i = 1,...,24,$ has been restricted to zero, reflecting the exogeneity of x_t and its lack of serial correlation and ε_t denotes the vector of serially and mutually uncorrelated structural innovations.

3.3. Identifying Assumptions

I postulate that A_0^{-1} has a recursive structure such that the reduced form errors e_t can be decomposed according to $e_t = A_0^{-1}\varepsilon_t$.

$$e_t \equiv \begin{pmatrix} e_t^{x} \\ e_t^{\Delta prod} \\ e_t^{rea} \\ e_t^{rpo} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \begin{pmatrix} \varepsilon_t^{political\ oil\ \text{supply}\ shock} \\ \varepsilon_t^{other\ oil\ \text{supply}\ shock} \\ \varepsilon_t^{aggregate\ demand\ shock} \\ \varepsilon_t^{oil-specific\ demand\ shock} \end{pmatrix}$$

The restrictions on A_0^{-1} may be motivated as follows: Shocks to crude oil production driven by political events in OPEC countries such as wars and revolutions are exogenous by construction and do not respond to any other structural shocks. These shocks constitute only one type of oil supply shock. There also may be shocks to crude oil production related to cartel activity or to oil production in non-OPEC countries, for example. The latter type of oil supply shocks (referred to as *other oil supply shocks* for short) is allowed to respond to the former type contemporaneously, but is 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 oil supply, if at all.

Innovations to global real economic activity that cannot be explained based on oil supply shocks of either type will be referred to as shocks to the 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 will be viewed as shocks that 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.

Implicit in this model are three more assumptions: First, there are no politically motivated exogenous supply shocks in industrial commodities other than oil. This assumption seems self-evident. Second, there are no precautionary demand shocks in industrial commodities other than crude oil. This assumption is consistent with the common view that crude oil is a unique commodity from the point of view of the oil-importing countries. Third, I assume that idiosyncratic shocks to the demand or supply of dry cargoes average out in the construction of the index of real economic activity.

3.4. Empirical Results

The model is consistently estimated by applying the least-squares method equation-by-equation imposing the relevant restrictions on the lag structure. The resulting estimates are used to construct the restricted structural VAR representation of the model. Inference is based on a fixed-design wild bootstrap with 20,000 replications (see Gonçalves and Kilian 2004).

Initial estimates of the model suggested an initially positive response of real economic activity to adverse political oil supply shocks, followed by a negative response. Inspection of the time series plots reveals that most episodes of major negative political oil supply shocks occur while real economic activity is growing, and many of the subsequent positive political oil supply shocks occur during times of falling real economic activity. We have no reason to believe that this pattern is more than a coincidence since wars and civil unrest are presumed exogenous. Since we know the observed negative conditional correlation to be a spurious small-sample phenomenon, in the estimation results shown below I impose the additional plausible restriction that political oil supply shocks can only affect real economic activity through their effect on the real price of oil. A similar restriction is imposed on the feedback from other oil supply shocks to real economic activity to preserve symmetry. Imposing these additional exclusion restrictions on $A_i, i = 1,...,24,$ has virtually no effect on the remaining impulse responses, while eliminating the counterintuitive patterns of the response of real economic activity to oil supply shocks.

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 6 shows the responses of global oil production, real economic activity and the real price of oil to one-standard deviation structural innovations. The political oil supply shock has been normalized to represent a negative shock, whereas the other shocks are positive.

Political oil supply disruptions cause a temporary decline in global oil production upon impact, followed by an increase in production in the first year that is significant for a few months. They do not significantly raise the real price of oil at any horizon. There is some indication that they lower global real economic activity in the third year after the shock (through their effect on the real price of oil), but that decline is small and not significant at the 10 percent level.

Other oil supply expansions tend to raise the level of global oil production permanently and significantly, but the magnitude of the increase drops to about one half of the initial shock after six months. An exogenous increase in other oil supply also temporarily reduces the real price of oil. The biggest reduction occurs after 5 months. After one year, the effect is essentially zero. The reduction is significant at the 10 percent level for horizons 2 through 7, and at the 5 percent level at horizons 4 and 5. Other oil supply expansions also cause a small and marginally significant increase in global real economic activity in the third year after the shock (through their effect on the price of oil),

The effect of an aggregate demand expansion on real economic activity is very persistent and lasts almost four years, before leveling off. It remains significant at the 5 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. After 18 months the expansion ends. There is some indication that the initial increase is offset by small but persistent decreases at longer horizons, although the latter are not statistically significant. Aggregate demand expansions also cause a large and persistent increase in the real price of oil, albeit with a delay of six months. 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 a 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 temporary decline after two years. The former is marginally significant at the 5 percent level for two months, whereas the latter is clearly 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 decline in oil supply starting in the second year, although that decline is not significant at the 5 percent level.

Perhaps the most striking result in Figure 6 is the fact that political oil supply disruptions have no significant effect on the real price of oil. Part of the explanation, according to Figure 6, is that political oil supply disruptions in OPEC countries tend to trigger endogenous expansions of crude oil production elsewhere in the world that help offset exogenous OPEC production shortfalls. The lack of a significant response of the real price of oil is also consistent given related evidence that such 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 the political oil supply shocks – accounts for the apparent large increases in the real price of oil following the events underlying those political oil supply disruptions. Figure 6 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 6). This explanation will be explored further in the next subsections using historical decompositions of the real price of oil.

3.4.2. What Is the Magnitude and Timing of the Major Structural Shocks?

Figure 7 shows the time series of the structural residuals of model (1). The first panel replicates the series shown in Figure 2. Here we are primarily interested in the other three structural shocks. A natural candidate for explaining sudden shifts in the real price of oil is the oil-market specific demand shock, which by construction captures shifts in precautionary demand driven by increased concerns about the availability of future oil supplies. The last panel of Figure 7 shows that by far the largest spike in this series occurred immediately after the invasion of Kuwait in

August of 1990. There also are much smaller, but persistent positive oil-market specific demand shocks following the Iranian Revolution.

For a more systematic evaluation of these shocks Table 1 shows averages of $\hat{\varepsilon}_{jt},\ j=1,2,3,4,$ for selected subperiods. The period leading up the Iranian Revolution (1975.1-1978.9) on average was characterized by positive shocks to other oil supplies, moderately negative aggregate demand shocks, but moderately positive oil-market specific demand shocks, presumably reflecting uncertainty about OPEC oil supplies. In contrast, the years during and immediately after the Iranian Revolution were characterized by moderately negative shocks to other oil supplies, large positive aggregate demand shocks and even larger positive oil-market specific demand shocks, suggesting that the latter shocks must have played a central role in the 1979/80 real oil price increase. On the other hand, the period following the outbreak of the Iran-Iraq War (1980.9-1982.2) on average was characterized by strongly negative shocks to other oil supplies, negative aggregate demand shocks, and almost no oil-market specific demand shocks.

The period of OPEC supply management (1982.3-1985.10) saw continued moderately large negative aggregate demand shocks. While oil market-specific demand shocks were moderately positive, reflecting increased uncertainty about future oil supplies not unlike in the mid-1970s, other oil supply shocks on balance were essentially absent, reflecting the inability of OPEC to reduce global crude oil production by much. The interim period between the collapse of OPEC and the Persian Gulf War (1985.11-1990.6) was characterized by strongly positive shocks to other oil supplies, reflecting the resumption of Saudi oil production, slightly negative aggregate demand shocks and relatively large negative oil-market specific demand shocks, consistent with a reduction in precautionary demand owing to the collapse of OPEC.

During the Persian Gulf War and its immediate aftermath (1990.7-1991.5) aggregate demand shocks were only slightly negative, but other oil supply shocks turned very strongly negative and oil-market specific demand shocks strongly positive. In the interim period between the Persian Gulf War and the Asian Crisis (1991.6-1997.6) there are no other oil supply shocks on average, aggregate demand shocks are negative and oil-market specific demand shocks slightly negative. During the Asian crisis (1997.7-1998.12) average shocks to other oil supplies turn moderately negative, as do average aggregate demand shocks. Most interestingly, oil-market specific demand shocks are highly negative, consistent with all precautionary motives vanishing, as the real price of oil reaches an all-time low. Since 1991.1 average aggregate demand shocks have been strongly positive, precautionary demand shocks have been moderately positive, and average oil supply shocks have been moderately negative.

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

More important than the timing and magnitude of the shocks in Figure 7 and Table 1 is their cumulative effect on the real price of oil. Figure 8 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 political oil supply shocks overall made little contribution to the real price of oil. The contribution of other oil supply shocks is of a similar magnitude. By far the biggest contribution is 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 an important role for precautionary demand shocks reflecting sharp 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, political and other 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 Chomeini'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.[6] While the oil supply disruption measured by the political oil supply shock also had some effect on the real price of oil, that effect was delayed until the end of the war.

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. 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. 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 8 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 political or other crude oil supply shocks.

The evidence in Figure 8 clearly suggests that not all oil price shocks are alike. There are important differences in the relative contribution of the four 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.

3.4.4. The Role of Precautionary Demand

One of the most important findings of the preceding subsection has been the disproportionate

[6] Kilian (2005) arrives at the same conclusion using a different methodology.

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 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 be 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, as the example below illustrates, making it difficult to provide more formal evidence in favor of the interpretation of the oil-market specific demand shock as a precautionary demand shock. I focus on the so-called Tanker War in the Persian Gulf between 1984 and 1988. In this war, Iraqi military forces attempted to sink the tankers servicing Iranian ports, and Iranian military forces attacked neutral oil tankers as well as port facilities in neutral Arab Gulf states in an effort to affect the outcome of the Iran-Iraq War. At one point more than 30 crude oil tankers were attacked in a given month. I compiled a monthly time series of the number of attacks on ships, as shown in the first panel of Figure 9. The sources are described in the Data Appendix. I also compiled a database of incidents of maritime military conflict in the oil shipping lanes of the Persian Gulf. Examples include the attack on the USS Stark, the downing of Iranian fighter jets by the Saudi Air Force over the Gulf, missile attacks on oil facilities on and off the shore of neutral Gulf states such as Kuwait and engagements between the Iranian navy and the U.S. Navy, which after May 1987 had begun to secure the sea lanes in the Gulf. That time series is shown in the second panel of Figure 9. The construction of this series is discussed in the Data Appendix.

For the third panel of Figure 9, I add these incidents to the attacks on tankers, as both provide a measure of the overall risks to oil shipping in the Persian Gulf. A natural presumption is that fears about future oil supplies should co-move with this time series. As it turns out, the link from this time series to oil-market specific demand shocks – if it exists – must be highly nonlinear. Figure 10 shows that the increased risk to oil shipping since the early 1980s had very little effect on demand for oil until the end of 1985. The most prominent feature of the data is a sharp drop in oil-market specific demand in early 1986 that arguably reflects increased confidence about future oil supplies following the collapse of the OPEC cartel regime in late 1985. This effect was so strong that it more than offset any effects from the resurgence of attacks on oil tankers at about the same time. There is some indication that precautionary demand may have rebounded in response to these attacks later in 1986 and in early 1987, but by mid 1987 that concern seems to have vanished, as the world and the oil industry got used to continued low-level attacks on oil shipping. Only in late 1987, at the peak of the Tanker War, did oil-market specific demand briefly spike again. There is no indication either that the efforts of the United States after July of 1987 and after April of 1988, respectively, to protect its own and later all countries' oil tankers had any effect on oil-market specific demand shocks. The absence of any apparent comovement between the oil-market specific demand shock and indicators of military conflict in the Persian Gulf is, of course, consistent with the views of oil economists that "a major political crisis will not cause a price shock when capacity cushions exist in other countries", as was the case in the period in question (see Mabro 1998, p.10).

This example illustrates that the task of further decomposing oil-market specific demand shocks into an expectations component and a non-expectations component is likely to be infeasible. Even when the determinants of concerns about future oil supply disruptions are well understood and can be measured, as in the example here, their relationship with oil prices is highly nonlinear. 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.

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

4.1. Methodology

4.1.1 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. The first problem in answering this question is that the latter aggregate is not available at monthly frequency.[7] 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:

[7] 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.

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

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.

The second problem is that the structural shocks, $j = 1, 2, 3, 4$, are by construction strictly exogenous with respect to the information set in model (1), but, with the exception of the first shock, are not necessarily exogenous with respect to U.S. macroeconomic aggregates. Strict exogeneity involves two conditions: (1) Pre-determinedness of the structural shock series and (2) no feedback from lagged macroeconomic aggregates to the shock series. While the second condition is testable, the first one is not (see, e.g., Cooley and LeRoy 1985). Regressions of the form

(2) $$\hat{\zeta}_{jt} = \beta_0 + \sum_{i=1}^{p}\beta_i \Delta y_{t-i} + u_t, \quad j = 2,...,4$$

and

(3) $$\hat{\zeta}_{jt} = \beta_0 + \sum_{i=1}^{p}\beta_i \pi_{t-i} + u_t, \quad j = 2,...,4$$

allow tests of the no-feedback hypothesis $H_0 : \beta_i = 0 \; \forall i$. Given the absence of serial correlation in the dependent variable, there is no need to include lags of the dependent variable. A rejection of this null would also be a rejection of the strict exogeneity hypothesis. Tests are conducted based on a fixed-design wild bootstrap. Although for some choices of p the null of no feedback cannot be rejected at the 5 percent level, for other choices the null is rejected for at least some j, and often the test results are highly sensitive to minor changes in p. Thus, it does not seem advisable to impose the assumption of strict exogeneity in the regression analysis below.

If we are willing to make the weaker assumption that within a given quarter there is no feedback from U.S. real GDP growth (Δy_t) and CPI inflation (π_t) to $\hat{\zeta}_{jt}$, $j = 2,...,4,$ these shocks can be treated as predetermined and we can examine their effects on U.S. macroeconomic aggregates based on the regressions:

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

and

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

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.[8]

Although the assumption of predeterminedness is not testable, there is some evidence that this assumption is not at odds with the data. Table 2 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. Of particular interest are the entries for $j = 2,3,4.$ Table 2 shows that in most cases the correlation of these residuals with $\hat{\zeta}_{jt}$ is quite low, suggesting that there is little feedback one way or the other within the quarter. Specifically, U.S. CPI inflation residuals have only low positive correlation with other oil supply shocks (9 percent). There also is a low negative correlation with aggregate demand shocks (-1.2 percent). The largest correlation by far is with precautionary demand shocks (22 percent). Real growth residuals have a low correlation with other oil supply shocks (8 percent), a negligible correlation with global aggregate demand shocks (1 percent) and a small negative correlation with precautionary demand shocks (-5 percent). This evidence suggests that in general the common component in these shock series is small and the series are nearly uncorrelated. Thus, with the possible exception of the oil-market specific demand shock and U.S. CPI inflation, the regression estimates based on (4) and (5) should provide a good approximation, even if the assumption of predeterminedness were incorrect.

For $\hat{\zeta}_{1t}$, unlike for the other three shocks, exogeneity holds by construction (see Kilian 2006a). Thus, the dynamic effects of $\hat{\zeta}_{1t}$ may be estimated from a model that differs from the model used for $j = 2,3,4$ in that it imposes the no-feedback assumption:

[8] 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.

$$(4') \qquad \Delta y_t = \alpha + \sum_{i=1}^{4} \beta_i \Delta y_{t-i} + \sum_{j=0}^{8} \gamma_j \hat{\zeta}_{1t-j} + u_t$$

and

$$(5') \qquad \pi_t = \delta + \sum_{i=1}^{4} \lambda_i \pi_{t-i} + \sum_{j=0}^{8} \eta_j \hat{\zeta}_{1t-j} + v_t ,$$

where the error terms u_t and v_t are serially uncorrelated. For the derivation of this model see Kilian (2006b). Inference is conducted based on the fixed-design wild bootstrap (see Gonçalves and Kilian 2004). The differences in the impulse responses obtained from models (4) and $(4')$ are negligible for $j = 1$; those between (5) and $(5')$ are not. In the latter case, formal statistical model selection procedures systematically favor model $(5')$.

4.1.2. Alternative Approaches

The baseline model in section 4.1.1. regresses the U.S. macroeconomic aggregate of interest on lags of $\hat{\zeta}_{jt}$, $j = 1,...,4.$ An alternative approach would have been to project the macroeconomic aggregate of interest on lags of the component of the real price of oil that can be predicted based on current and lagged values of $\hat{\zeta}_{jt}$, $j = 1,...,4,$ respectively. I do not include results for that alternative approach for two reasons. First, that approach may be omitting some of the effects of the demand and supply shocks that drive the real price of oil. For example, a global aggregate demand shock may have direct effects on the U.S. economy working through the United States' external balances beyond those related to the implied changes in the real price of oil. Second, substituting the predicted component of the real price of oil amounts to instrumenting for the real price of oil. Results in the working paper version of Kilian (2006a) suggest that there is reason to believe that at least some of the shocks $\hat{\zeta}_{jt}$, $j = 1,...,4,$ are weak instruments for the real price of oil, which would undermine the reliability of the implied impulse response estimates and invalidate standard statistical inference.

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

Figure 11 summarizes the responses of the growth rate and level of U.S. real GDP to each of the four shocks defined earlier. Political oil supply disruptions lower real GDP growth. While there is some evidence of a decline already in the first and second quarter after the shock, the bulk of the reduction in real growth occurs in the second year after the shock. Much of that reduction is significant at the 10 percent level.[9] The corresponding decline in the level of real GDP is initially slow, but accelerates after four quarters. It is statistically significant at the 10 percent level after five quarters. What is apparent from this analysis is that political oil supply shocks have a significant effect on real GDP, although they do not raise the real price of oil significantly, as shown in Figure 6. One possible explanation for such an effect is firms' and households' expectations. Especially if there is a public consensus that 'oil shocks' are bad for the economy, based on the experience of 1973/74, agents may immediately adjust their behavior without waiting for oil prices to adjust, thus generating a self-fulfilling prophecy. A second channel would be a preemptive interest rate hike by the central bank. Kilian (2006b) provides evidence that the real GDP growth responses of G7 countries to exogenous oil supply disruptions vary with the interest rate responses to these shocks.

An expansion of other oil supply significantly raises real GDP growth on impact. There also is some weak indication of a temporary increase in real growth in the second year. With the exception of a marginally significant decline after ten quarters, the response is not significantly different from zero at the 5 percent level after the initial impact. The response of real GDP is positive at all horizons, but marginally significant at the 10 percent level only for the first two quarters. There is no indication that an expansion of other oil supply significantly raises real GDP in the long run.

An aggregate demand expansion causes an initial increase in real GDP growth, followed by a decline 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 significant increase in the level of real GDP for the first three quarters, but the subsequent fall in real GDP 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 quarter.

[9] Unlike the results in Kilian (2005, 2006) these estimates show a significant reduction in real growth as early as after two quarters, while, overall, the effect of the political oil supply shock on real growth is less significant than the earlier estimates suggested. That difference in results is driven entirely by the shorter sample period.

Figure 12 shows the corresponding results for U.S. CPI inflation and the level of the CPI. There is some evidence that political oil supply disruptions raise inflation temporarily. The peak response occurs in the third quarter, but is barely significant at the 10 percent level. Ignoring a marginally significant increase upon impact, there is no significant increase in the level of consumer prices at any horizon at the 10 percent level.[10] An expansion of other oil supply initially has little effect on CPI inflation, but causes a significant increase in CPI inflation at the 10 percent level between quarters five and ten. The corresponding effect on the price level is essentially zero for the first year, followed by an increase that is significant at the 10 percent level after nine quarters, but never significant at the 5 percent level. An aggregate demand expansion causes a persistent increase in U.S. CPI inflation, although that increase is mostly insignificant at the 10 percent level. In contrast, the effect on the price level is significant at the 10 percent level after eight quarters. Again the most clear-cut results are obtained for an oil-market specific demand increase. There is an immediate 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 11 and 12 illustrate important differences in how the demand and 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. Political oil supply disruptions cause a significant decline in real GDP in the long run at the 10 percent level, whereas other oil supply disruptions cause a significant decline of real GDP only in the short run. Aggregate demand expansions significantly raise real GDP in the short run, 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 other oil supply have a delayed effect on the CPI that is not significant at the 5 percent level and political oil supply shocks have no significant effect on the CPI at conventional significance levels.

---

[10] These results are qualitatively similar to the estimates presented in Kilian (2005, 2006) for an extended sample period.

4.2.2. Which Shocks Cause Stagflation?

A common concern is that higher oil prices may be 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 12, 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) and 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 12 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 13 shows that it makes a difference which shock is driving the macroeconomic aggregates. The point estimate of the response to political supply shocks shows virtually no evidence of stagflation, while the point estimates of the response to other oil supply shocks suggests that if there is a stagflationary effect at all, it only occurs about ten quarters after the shock. 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 upon 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 no evidence that the response to political oil supply shocks is stagflationary at any horizon at the usual significance levels. The response to other oil supply shocks is marginally significant at the 10 percent level in the tenth quarter, but otherwise not significantly different from zero. The response to a global aggregate demand shock is marginally significant after four, six, and eight quarters at the 10 percent level. Only the response to oil-market specific demand shocks is highly significant at the 5 percent level. The most significant impact occurs after four quarters with some additional less significant effects at other horizons. I conclude that the risk of stagflationary responses depends very much on the origin of the oil price increase and is much more pronounced for oil demand shocks than for oil supply shocks.

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 what the cumulative effect of a given type of shock on U.S. economic performance has been. The first panel of Figure 14 shows that political oil supply disruptions lowered U.S. real growth following each of the major political events noted in Figure 1. Real growth declines following the Iranian Revolution, following the Iran-Iraq War, following Persian Gulf War and to a much lesser extent following 2003 Iraq War, although typically those declines occur with a considerable delay. Equally important is the observation that exogenous oil supply shocks raised U.S. real GDP growth throughout most of the 1980s and between 1993 and 2001.

The second panel shows a positive effect of other oil supply shocks on real GDP growth after the collapse of OPEC in late 1985, but little systematic impact of OPEC's attempt at supply management in the early 1980s. Moreover, there is no evidence that other oil supply shocks have systematically lowered U.S. real growth since 2003.

The third panel suggests that the U.S. economic expansion in 1978/79 was fueled 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 and briefly in 2003/04, 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 the mid-1980s. They also caused a sustained reduction in real growth in 2002/03. 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 and after 2003.

Figure 15 shows the corresponding results for U.S. CPI inflation. Political oil supply shocks in the first panel are seen to have had little impact on CPI inflation in 1979-1982, while contributing to higher inflation in 1990-92. There also is some evidence of lower inflation owing to political oil supply shocks in 1983-1989 as well as 1993-2002. The second panel shows other oil supply shocks alternately raised inflation and lowered inflation. For example, they raised inflation until 1982, but lowered inflation until 1986. Most recently, having caused lower inflation between 2000 and 2005, their effect has been close to zero.

The third panel shows important 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 driven by these shocks 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 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 14 and 15 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 factors and policies. Figure 16 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 16 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 16 suggests that much of the decline in real growth in the early 1980s and its recovery in the early 1980s was externally caused. Only the brief recovery of real growth in 1981 was driven by domestic factors. There also is evidence that external factors contributed substantially to the 1990/91 contraction and the subsequent recovery. However, some of the decline in real GDP growth in 1991 and the comparatively high real growth rates of 1992 were due to domestic factors and policies.

The second panel shows that a substantial part of the high CPI inflation rates between late 1978 and 1ate 980 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 two thirds of the observed inflation rate. The domestic component of inflation was essentially wiped out later in 1980, at least temporarily. 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. Over time the gap between the two series narrows. 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. Whereas the low inflation rates of the late 1990s were helped by deflationary external shocks, between late 2000 and early 2005 the Fed managed to keep inflation below average even during times of external inflationary pressures. Figure 16 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 or of the exogenous shocks to the production of crude oil driven by political events in OPEC countries. 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. 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 in understanding changes in oil prices is likely to survive.

A 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 oil causes an immediate and persistent increase in the real price of oil, an increase in aggregate demand for all industrial commodities causes a delayed, but sustained increase in the real price of oil; politically driven crude oil production disruptions do not systematically increase the real price of oil (in part because they tend to be offset by production increases elsewhere), whereas other oil production disruptions cause a 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 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 mainly 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 supplies. In practice, it therefore is important to understand the extent to which increases in the price of oil are driven by one shock or another, before formulating appropriate policy responses. It also is important to understand better the relationship of U.S. macroeconomic aggregates and global macroeconomic aggregates.

Data Appendix 1: Military and Terrorist Attacks on Commercial Shipping in the Persian Gulf in 1975.1-2005.9

The data set includes military and terrorist attacks that involve specific ships and dates. It records the number of military and terrorist attacks on commercial shipping by month. I do not include pirate attacks. The data set heavily draws on the detailed data set compiled by Navias and Hooton (1996), which in turn draws on data provided by Lloyd's. This data set was selectively supplemented by data in Cordesman and Wagner (1990, Table 14.1), when the latter source reported additional attacks for the years 1981 through 1984. The annual totals implied by the combined data set are similar to those reported in O'Rourke (1988, 1989) based on data from Lloyd's Shipping Intelligence Unit as well as the Center for Defense Information.

Whereas the data set covers military action against commercial shipping in the Persian Gulf between 1980 and 1991, it excludes attacks on ships bottled up in Iranian ports when the Iraqi ground offensive reached the Iranian port cities. It also excludes attacks on Iraqi ships that occurred as part of the 1990/1991 Persian Gulf War, the subsequent U.N. oil embargo against Iraq, and the 2003 Iraq War. The data set includes the terrorist attack near Aden on the French tanker Limburg on October 6, 2002. Although this attack occurred outside the Gulf, it was widely viewed as indicative of the risk of terrorist attacks along the tanker routes in the Persian Gulf. On the other hand I exclude the mining attacks in the Red Sea in 1984 that are commonly attributed to Libya, which had remained on hostile terms with Egypt ever since Egyptian-Libyan War of July 1977. Of the 18 ships hit by mines in that area in 1984, according to Navias and Hooton (1996, Table 4.3) only one was a tanker, suggesting that the mines endangered mostly coastal traffic, but not oil tankers.

Data Appendix 2: Oil-related Allied Military Engagements in the Persian Gulf in 1975.1-2005.9

This data base collects information on military engagements in the Persian Gulf involving U.S. and allied naval forces, air and naval forces of neutral countries such as Kuwait, Saudi Arabia and the UAE, as well as attacks on shore-based oil facilities of neutral countries. Most engagements took place during 1980-1988. I exclude the border dispute between Iran and the UAE over islands in the Gulf because Iran's occupation of these islands has not led to interference with tanker traffic to date. I include the terrorist attack on the USS Cole in Yemen, however, which was widely viewed as an indication of increased risks to shipping in the Gulf.

The list of incidents has been compiled from Navias and Hooton (1996) and Cordesman and Wagner (1990). The data set excludes two minor incidents in1987 and 1988, in which U.S. forces traded fire with Iranian navy or air force units.[11] The data set also excludes incidents such as the loss of a U.S. helicopter during military operations in the Gulf in April 1988, for which the cause of the crash is not confirmed. Finally, the data set excludes incidents in which Iranian attacks were averted by a show of force such as the October 3, 1987, incident, in which 48-60 Iranian gunboats prepared to attack Saudi shore-based oil facilities, but were turned away by superior U.S. naval forces and the Saudi Air Force.

The data set includes Iranian artillery and missile attacks on Kuwait not already accounted for in other data bases because these attacks were targeted mainly at shore-based oil facilities such as loading docks that are closely linked to the tanker war. It excludes attacks on shore-based oil facilities in Iraq and Iran because the latter do not affect the security of international oil shipping lanes.

When there are multiple military engagements within the same day, these are counted as separate incidents if and only if they are clearly separated geographically or in time. The data base covers (1) attacks on former oil platforms that have been converted to naval bases, (2) incidents in which aircraft were shot down or damaged, (3) missile or artillery strikes on shore-based oil facilities in Kuwait, (4) attacks on warships.

Attacks on former oil platforms or warships are counted individually like attacks on ships. Attacks on gunboats are treated summarily, since the latter tend to operate in groups. The downing of one or more aircraft in the same engagement is treated as one incident. Missile attacks on sea or shore targets are counted even when the attack failed due to evasive action or countermeasures.

List of Military engagements:

June 5, 1984
Saudi F-15s down an Iranian F-4 Phantom over Saudi waters. (1)

May 17, 1987:
The USS Stark is severely damaged by an Iraqi missile. (1)

September 4, 1987
Iran fires two Silkworm missiles at shore-based oil facilities in Kuwait in separate attacks. (2)

[11] On August 8, 1987, U.S. F-14 fighter planes engaged an Iranian F-4 Phantom approaching a U.S. patrol plane, but two U.S. missiles missed their target. On July 12, 1988, U.S. helicopters exchanged fire with Iranian gunboats in the Gulf.

September 5, 1987
Iran fires a third Silkworm missile at Kuwait. (1)

September 21, 1987:
U.S. naval forces intercept and sink an Iranian navy mine-laying vessel. (1)

October 8, 1987
Iranian gunboats open fire on U.S. helicopters. One boat is sunk and two are damaged. (1)

October 19, 1987
U.S. naval forces attack and destroy two Iranian naval bases on former oil platforms. (2)

October 22, 1987
Iran fires shells into northern Kuwait and launches another missile attack on Kuwaiti off-shore oil facilities. Extensive damage to oil terminal. (2)

December 6, 1987
Iranian missile attack on Kuwaiti off-shore oil facilities. (1)

April 14, 1988
The USS Samuel B. Roberts is struck by an Iranian mine and heavily damaged. (1)

April 18, 1988
U.S. naval forces attack and destroy two Iranian naval bases on former oil platforms and attack a third one. (3)
Nearby an Iranian navy ship fires a missile at U.S. navy ships, and is sunk. (1)
Two Iranian F-4 Phantom jets fire missile at the USS Wainwright; one is damaged as the US ship returns missile fire. (1)
U.S. naval planes destroy three Iranian gunboats. (1)
An Iranian frigate approaching U.S. naval forces is sunk by the USS Joseph Strauss and U.S. naval planes. A second Iranian frigate is heavily damaged. (2)
Missile attack on USS Jack Williams (1)

April 20, 1988
Iran fires a Silkworm missile into Kuwait. (1)

July 3, 1988
USS Vincennes sinks two Iranian gunboats and downs an Iranian airliner mistakenly identified as an Iranian fighter. (1)

October 12, 2000
Terrorist attack on USS Cole in Aden, Yemen. (1)

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.

Cordesman, Anthony H., and Abraham R. Wagner (1990), *The Lessons of Modern War. Volume II: The Iran-Iraq War*, Westview Press: Boulder, CO.

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.

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.

Mabro, R. (1998), "OPEC Behavior 1960-1998: A Review of the Literature," *Journal of Energy Literature*, 4, 3-27.

Navias, Martin S., and E.R. Hooton (1996), *Tanker Wars: The assault on merchant shipping during the Iran-Iraq conflict, 1980-1988*, I.B. Tauris Publishers: London.

Oil and Gas Journal, various issues since 1970.

O'Rourke, Ronald (1988), "The Tanker War," Proceedings of the U.S. Naval Institute, 114, 30-34.

O'Rourke, Ronald (1989), "Gulf Ops," Proceedings of the U.S. Naval Institute 4, 115.

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. Average Structural VAR Residuals for Selected Subperiods

	1975.1-1978.9	1978.10-1980.8	1980.9-1982.2	1982.3-1985.10	1985.11-1990.6	1990.7-1991.5	1991.6-1997.6	1997.7-1998.12	1999.1-2005.9
Political Oil Supply Shock	0.00	-0.20	-0.03	0.04	0.03	-0.90	0.05	0.17	-0.02
Other Oil Supply Shock	0.17	-0.08	-0.34	-0.01	0.13	-0.34	0.01	-0.06	-0.04
Aggregate Demand Shock	-0.02	0.17	-0.09	-0.06	-0.03	-0.02	-0.10	-0.10	0.14
Oil-Market Specific Demand Shock	0.06	0.35	-0.02	0.06	-0.11	0.13	-0.02	-0.56	0.05

NOTES: Averages based on structural VAR residuals of model (1).

Table 2. Correlation in Percent of $\hat{\zeta}_{jt}$ with Autoregressive Residuals for U.S. Real GDP Growth and CPI Inflation

	Political Oil Supply Shock	Other Oil Supply Shock	Aggregate Demand Shock	Oil-Market Specific Demand Shock
U.S. Real GDP Growth	-10.7	8.2	1.1	-4.7
U.S. CPI Inflation	1.1	8.9	-1.2	22.2

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

Figure 1: Real Price of Crude Oil
1973.1-2005.9

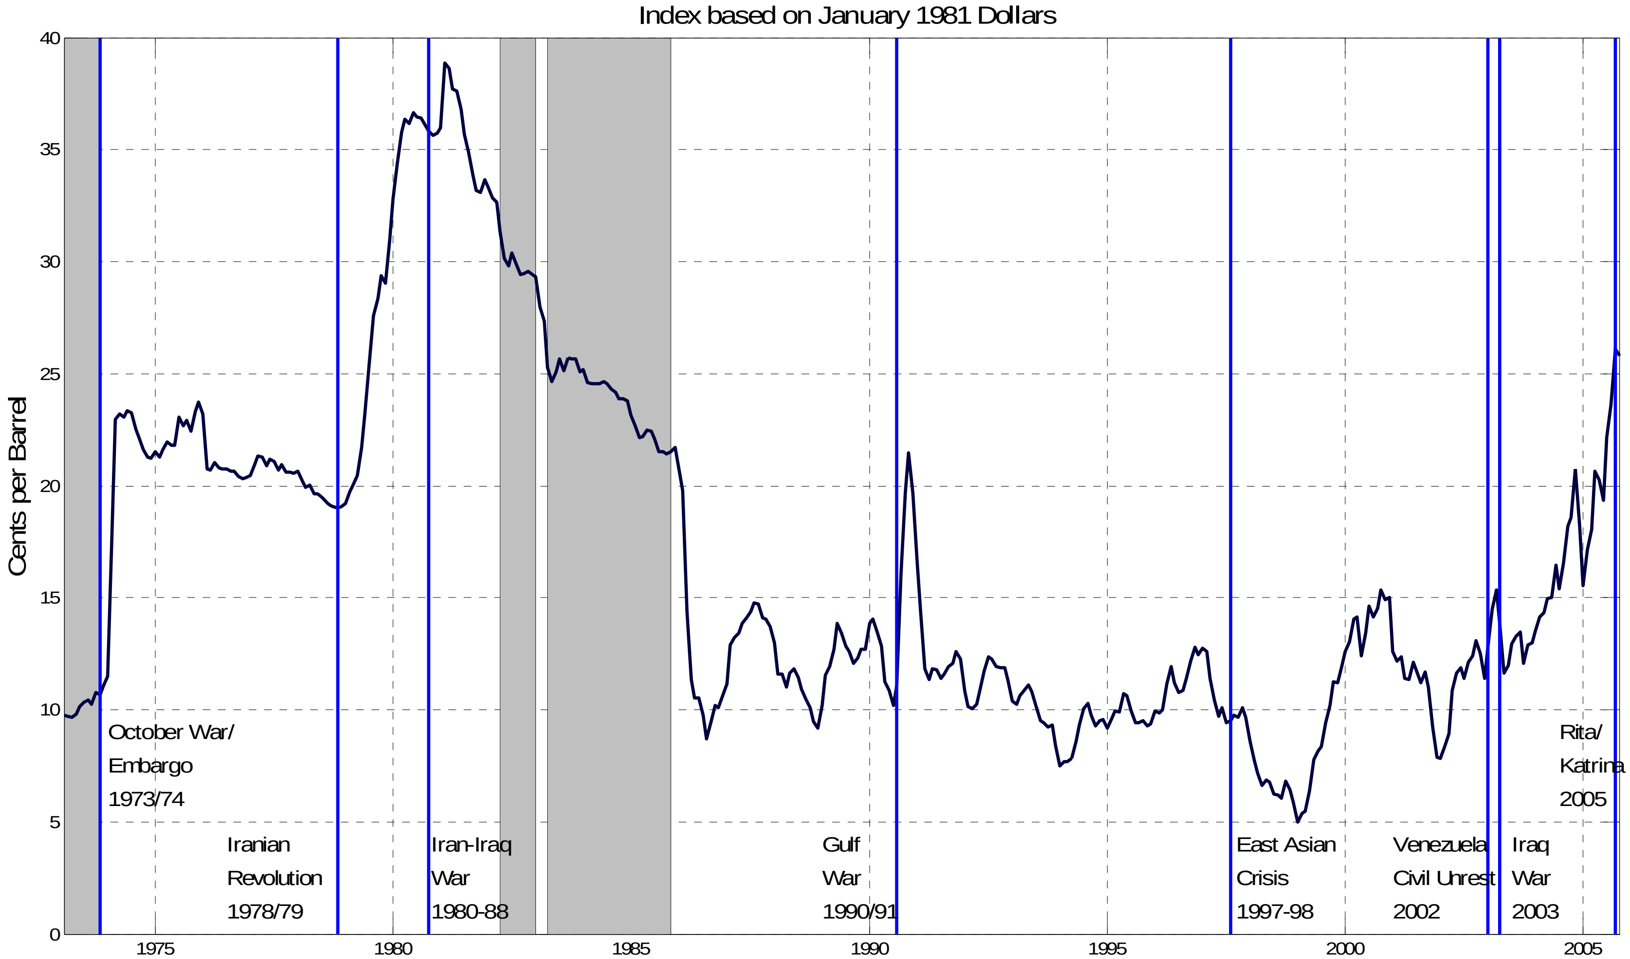

NOTES: The oil price series is the refiner acquisition cost of imported crude oil as reported by the U.S. Department of Energy, extended from 1974.1 back to 1973.1 as in Barsky and Kilian (2002). The oil price has been deflated by the U.S. CPI. Shaded areas mark periods of oil supply management by U.S.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

Figure 2: Exogenous Oil Supply Shocks Driven by Political Events in OPEC Countries 1973.1-2005.9

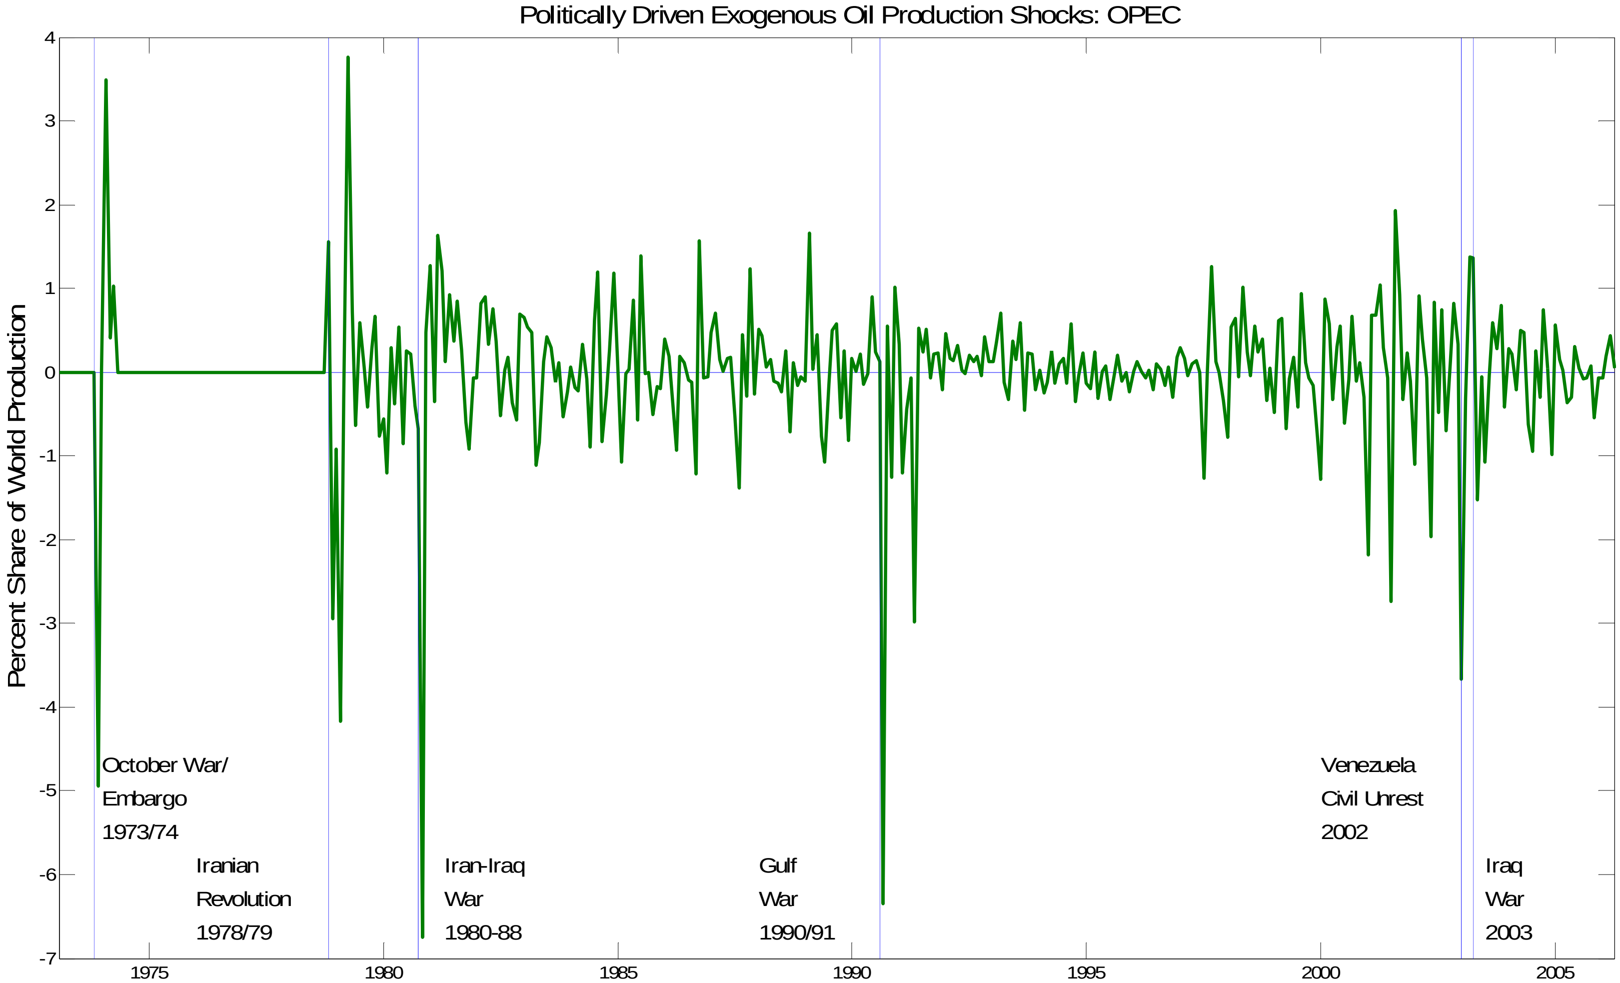

NOTES: The measure of the exogenous crude oil production shocks in OPEC countries is based on Kilian (2005, 2006).

Figure 3: Global Crude Oil Production
1973.1-2005.9

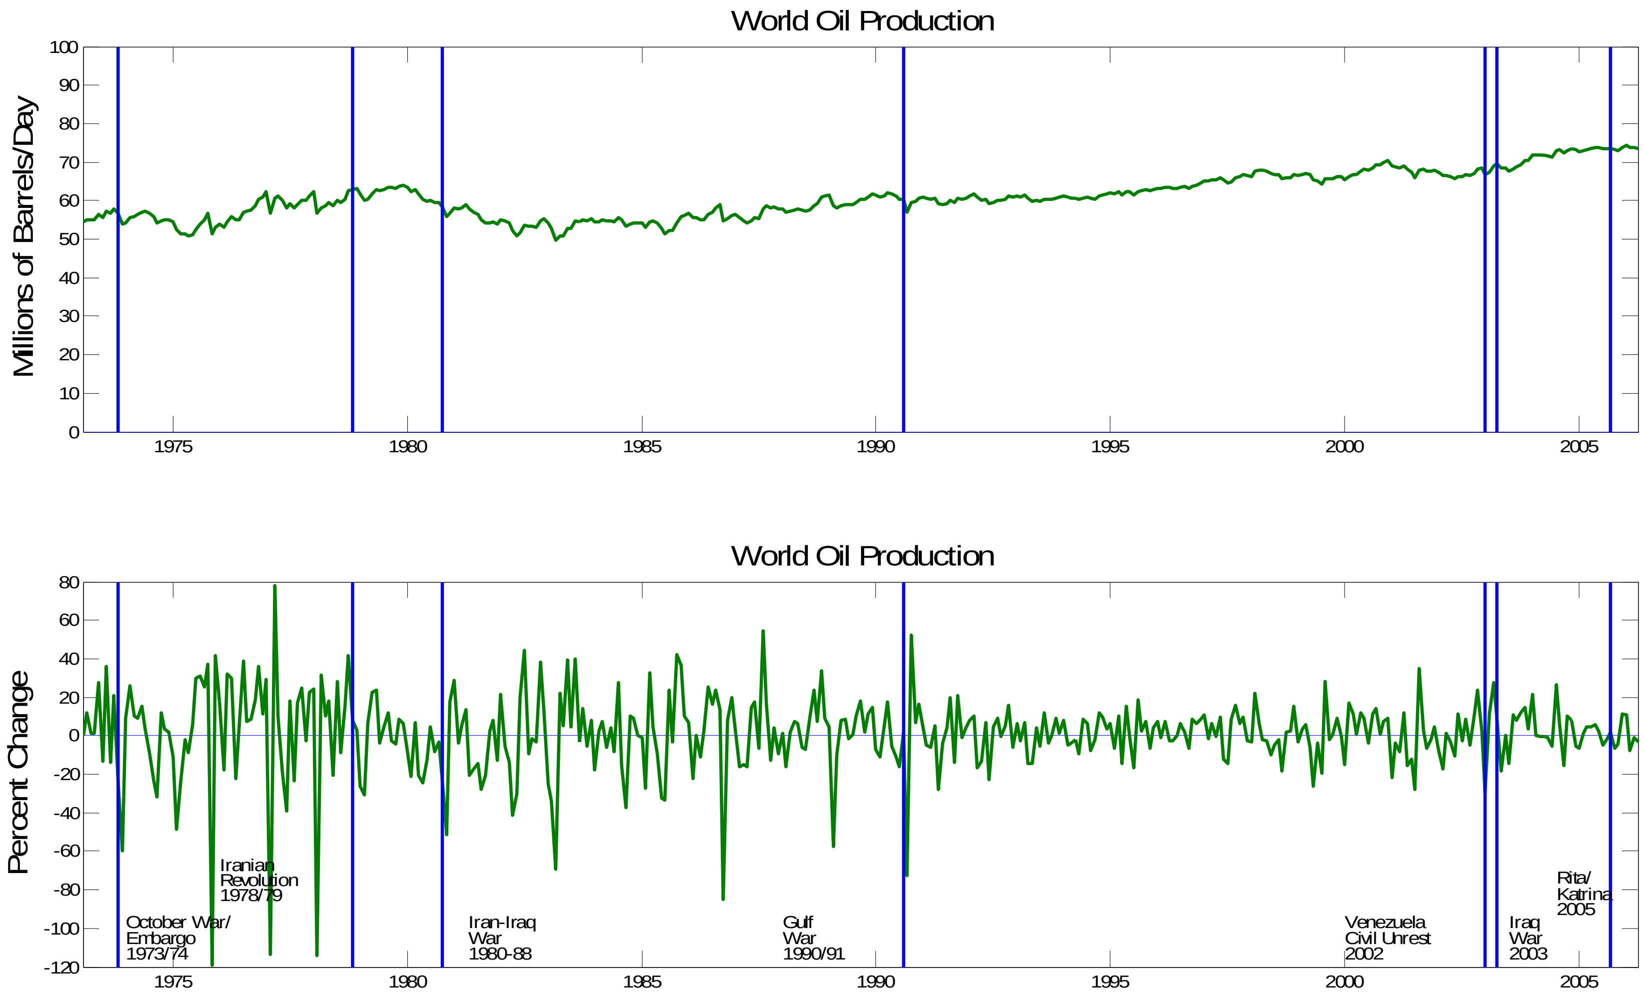

NOTES: World crude oil production as reported by the U.S. Department of Energy. The measure of the exogenous crude oil production shortfall in OPEC countries is based on Kilian (2005, 2006). The percent change has been annualized.

Figure 4: Index of Global Real Economic Activity based on Dry Cargo Bulk Freight Rates 1968.1-2005.9

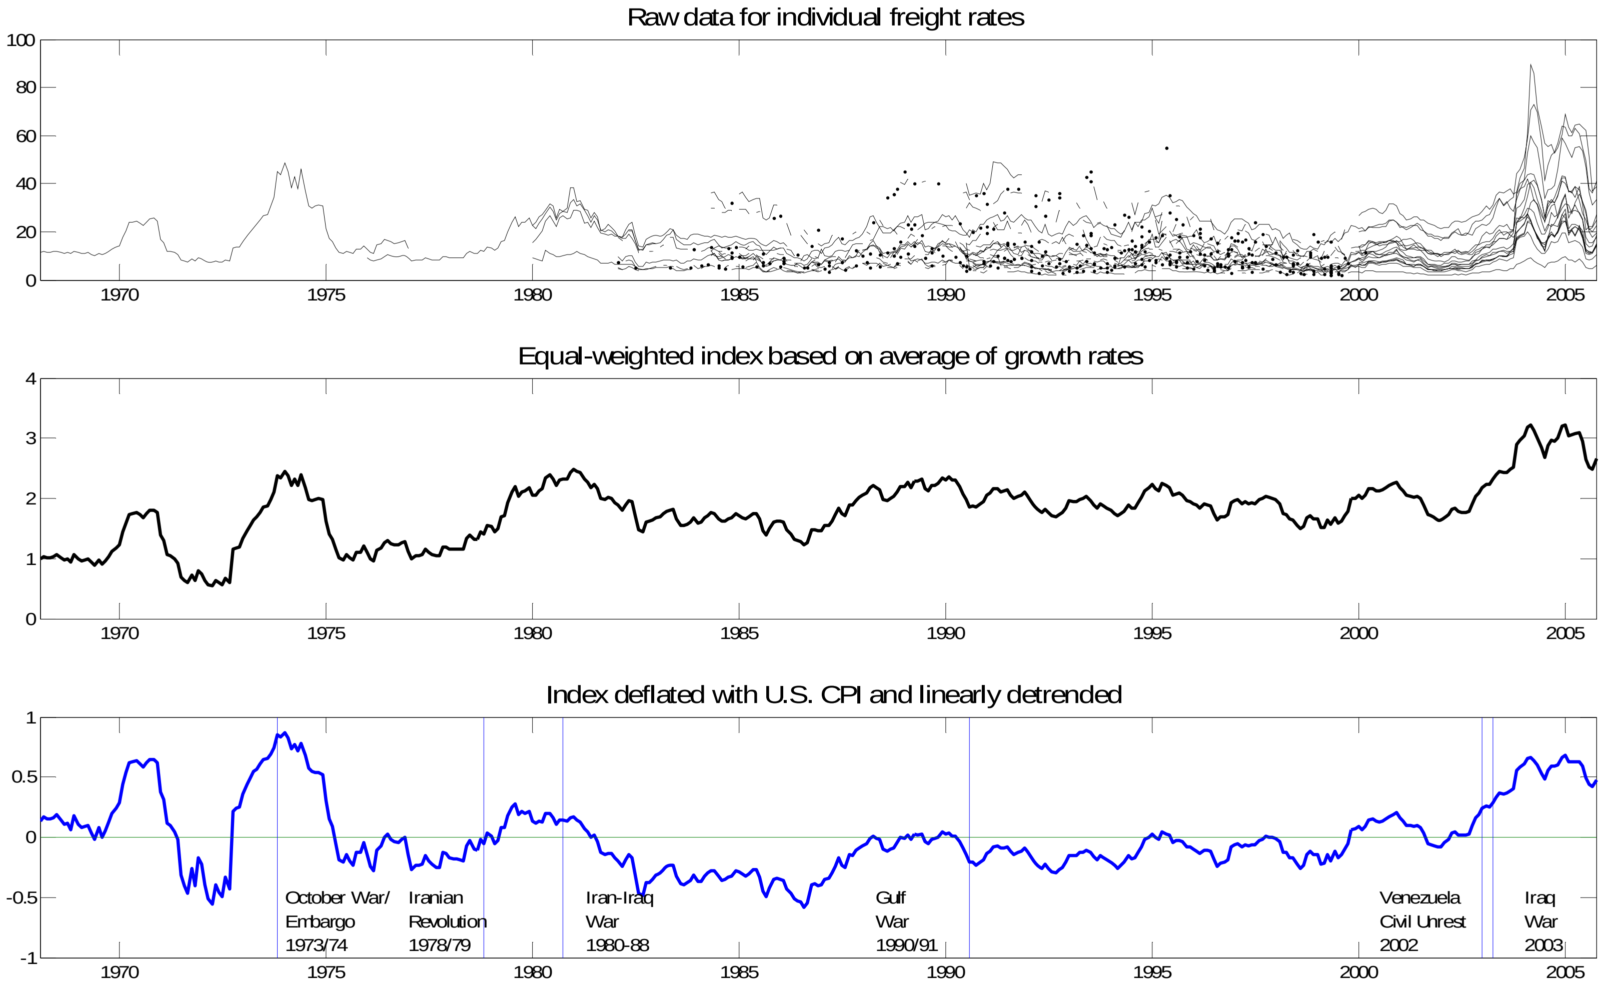

NOTES: The monthly raw data were manually collected from Drewry's Shipping Monthly, various issues since 1970. The two oldest series in the first panel are indices compiled by Drewry's. The remaining series are differentiated by cargo, route and ship size.

Figure 5: Index of Global Real Economic Activity and Index of Real Bunker Fuel Prices 1970.1-1973.12

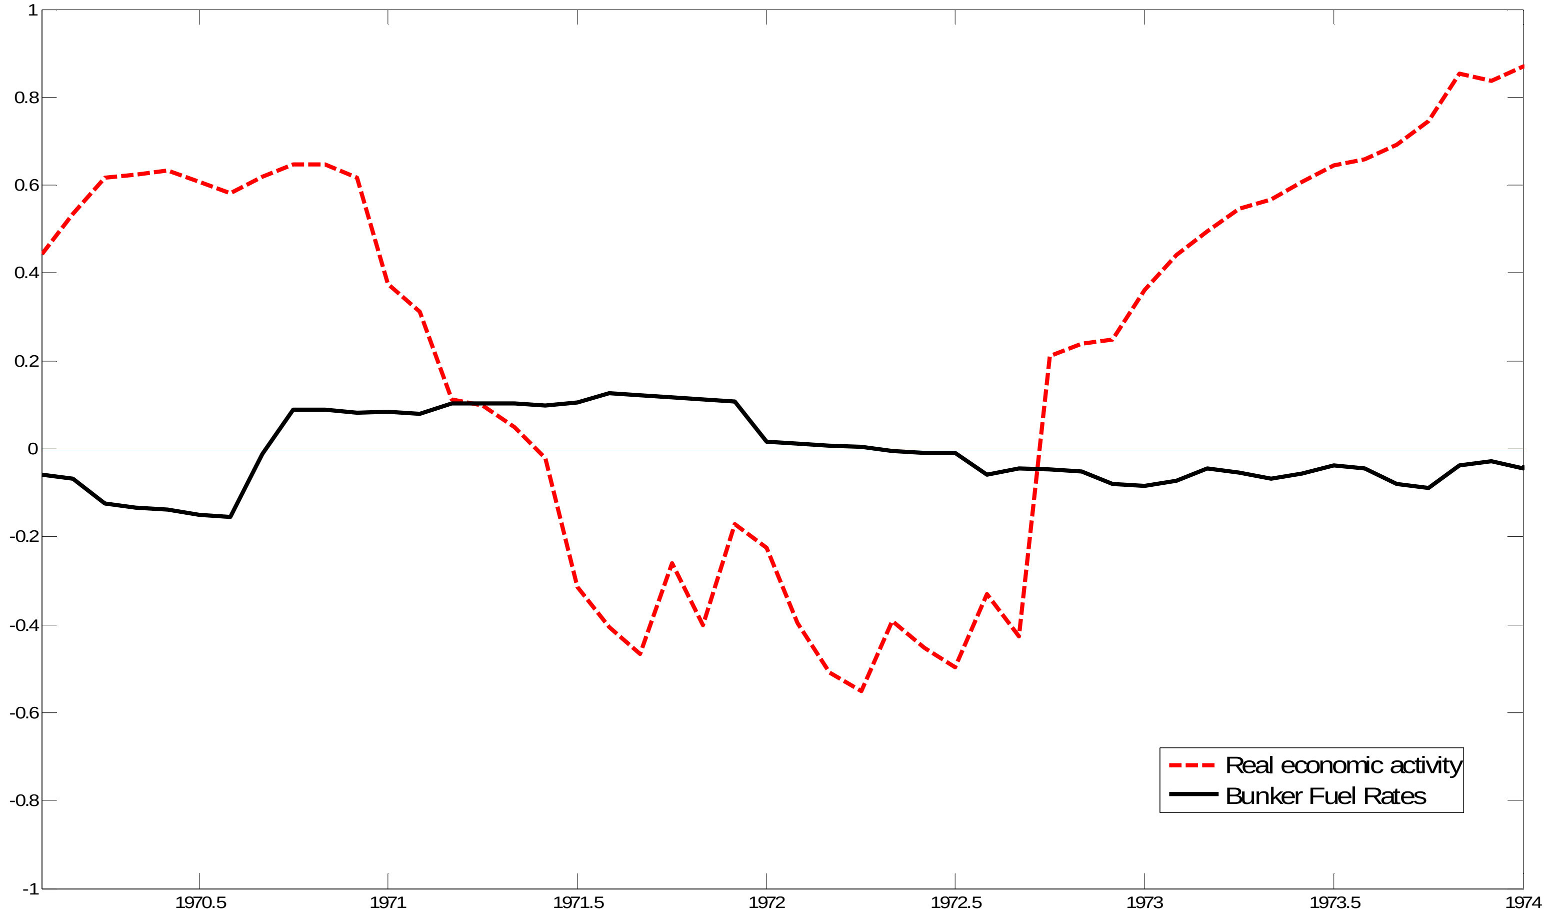

NOTES: The bunker fuel rate data are from the *Oil and Gas Journal*, various issues since 1970. All rates refer to Bunker C fuel, as recorded for the Caribbean, the Gulf Coast and California. The index is based on equal-weighted growth rates, computed using observations for the last week of each month. The real economic activity index is based on Figure 4.

Figure 6: Responses to One-Standard Deviation Structural Shocks
OLS Point Estimates with One and Two-Standard Error Bands

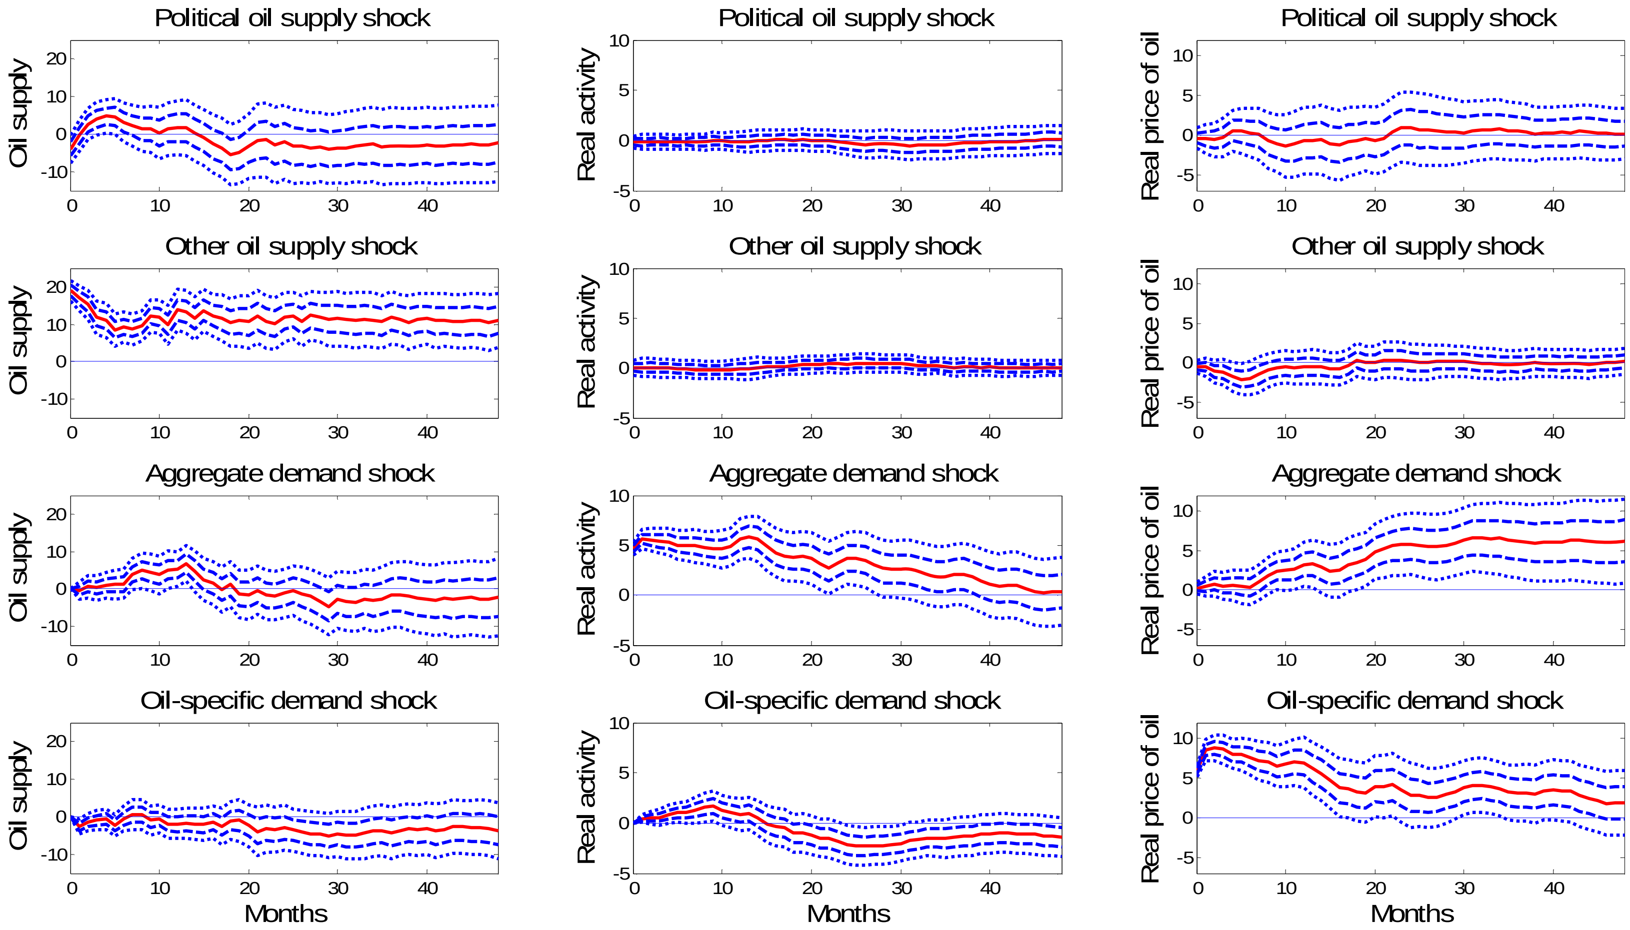

NOTES: Estimates based on restricted VAR(24) system described in text The confidence intervals were constructed using a fixed-design wild bootstrap.

Figure 7: Monthly Time Series of Structural Shocks
1975.1-2005.9

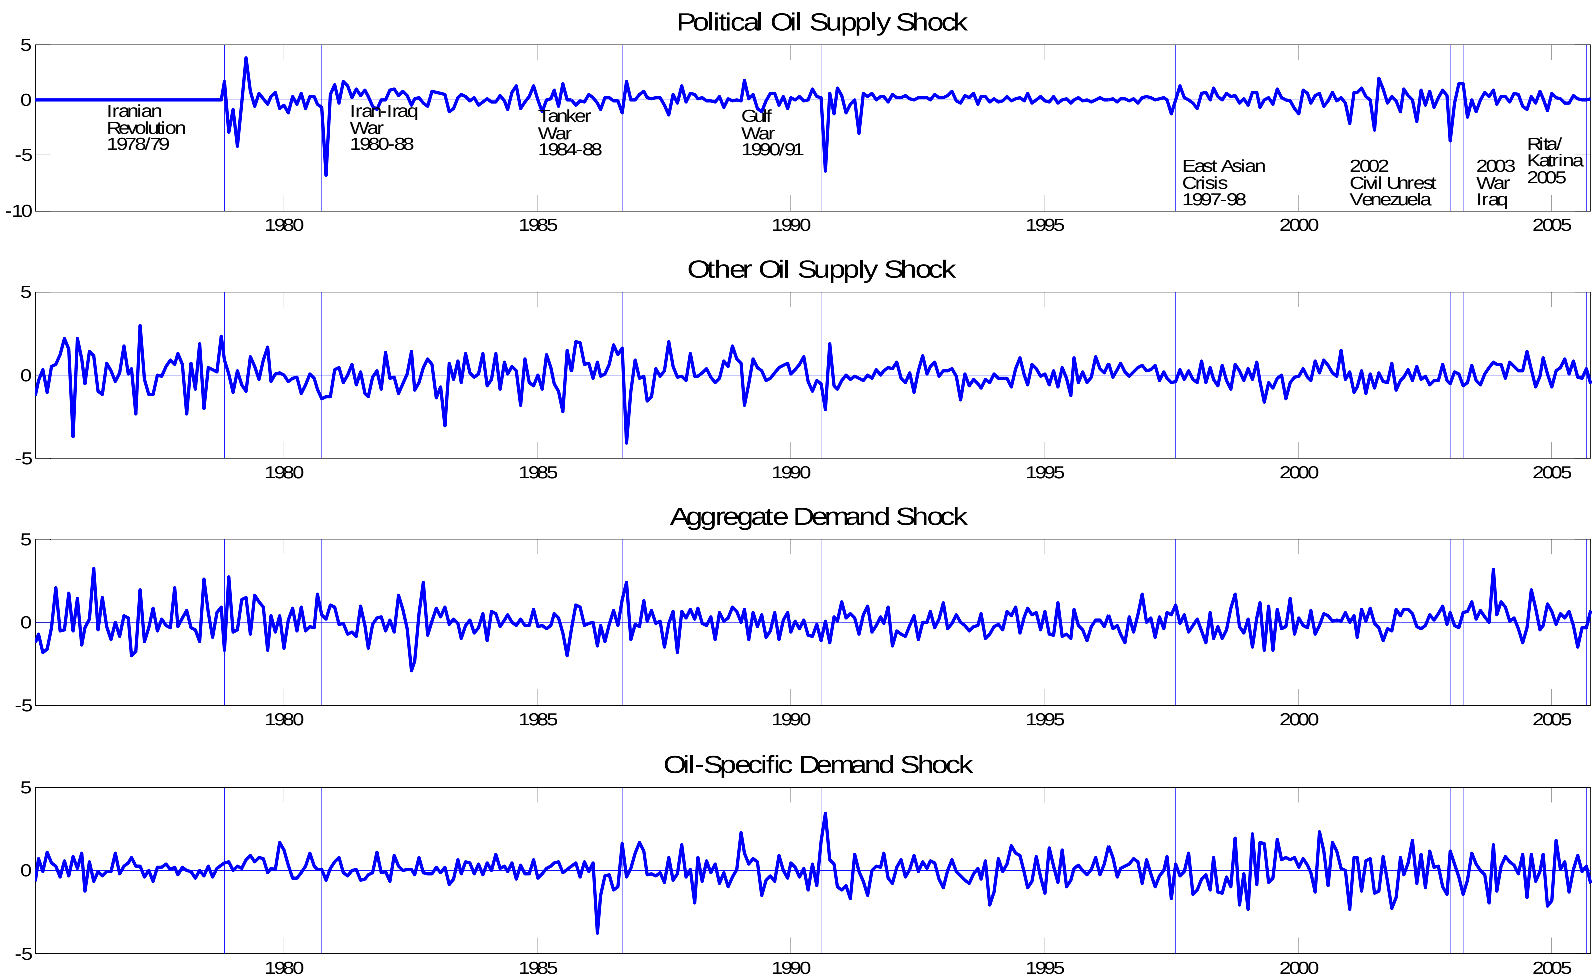

NOTES: See Figure 6.

Figure 8: Historical Decomposition of Real Price of Oil
1975.1-2005.9

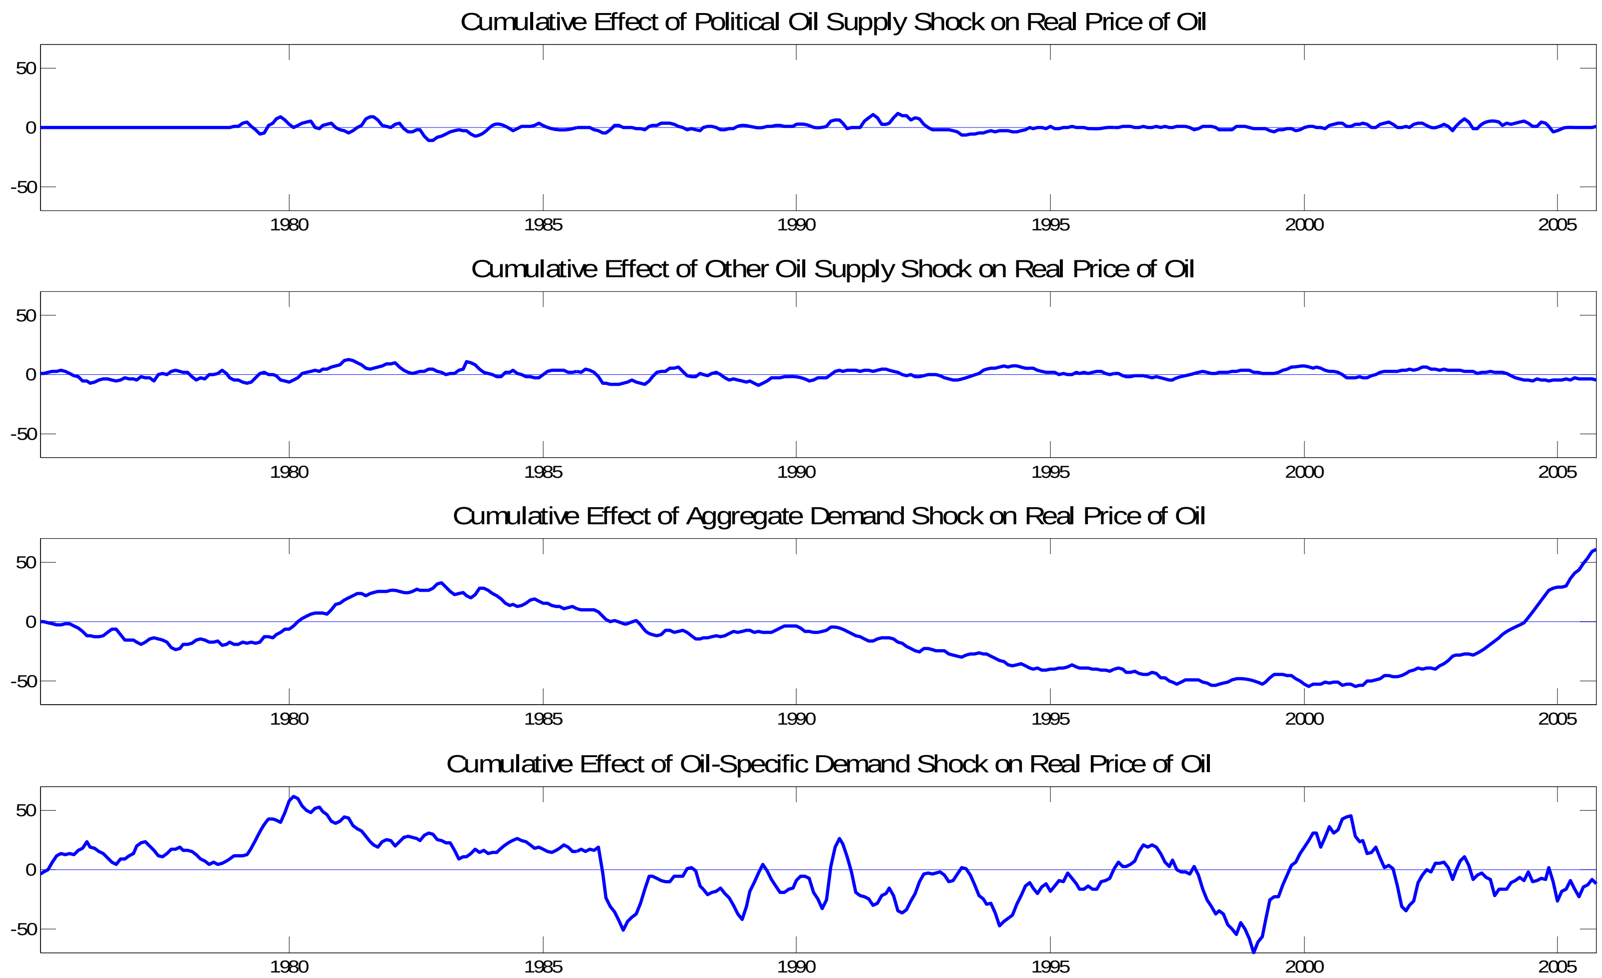

NOTES: See Figure 6.

Figure 9: Number of Military and Terrorist Attacks on Tankers in the Middle East 1973.1-2005.9

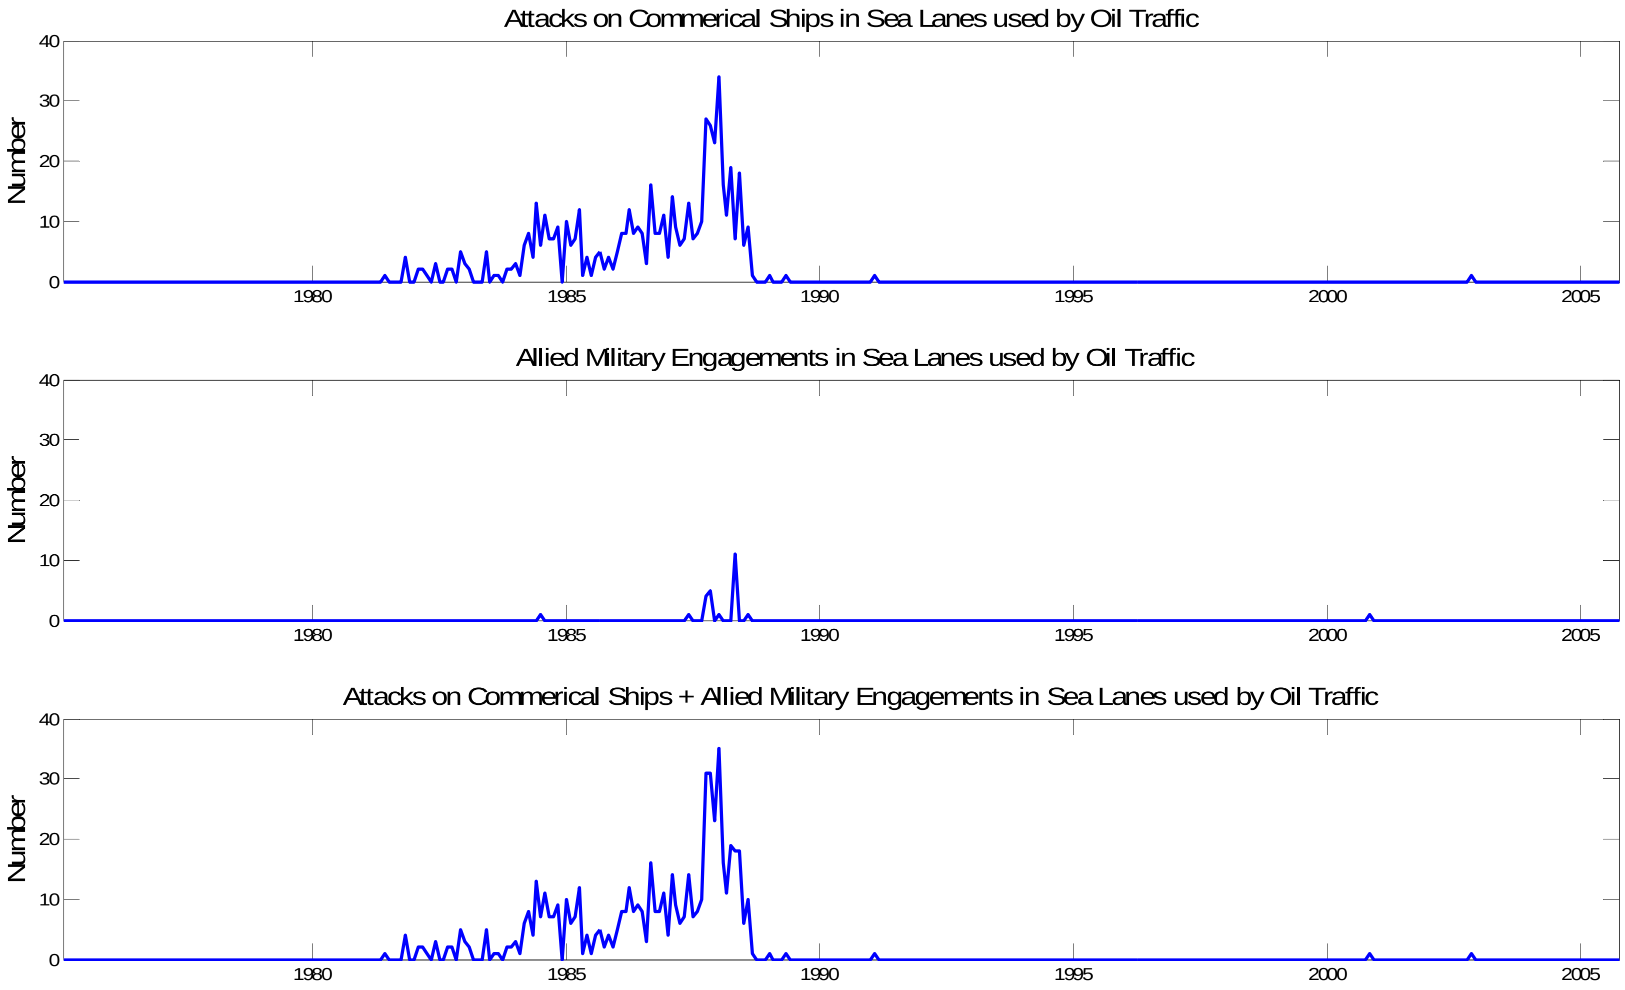

NOTES: The data sources are described in the Data Appendices.

Figure 10: Military Conflict in Oil Sea Lanes and Oil-Specific Demand Shocks During Iran-Iraq War

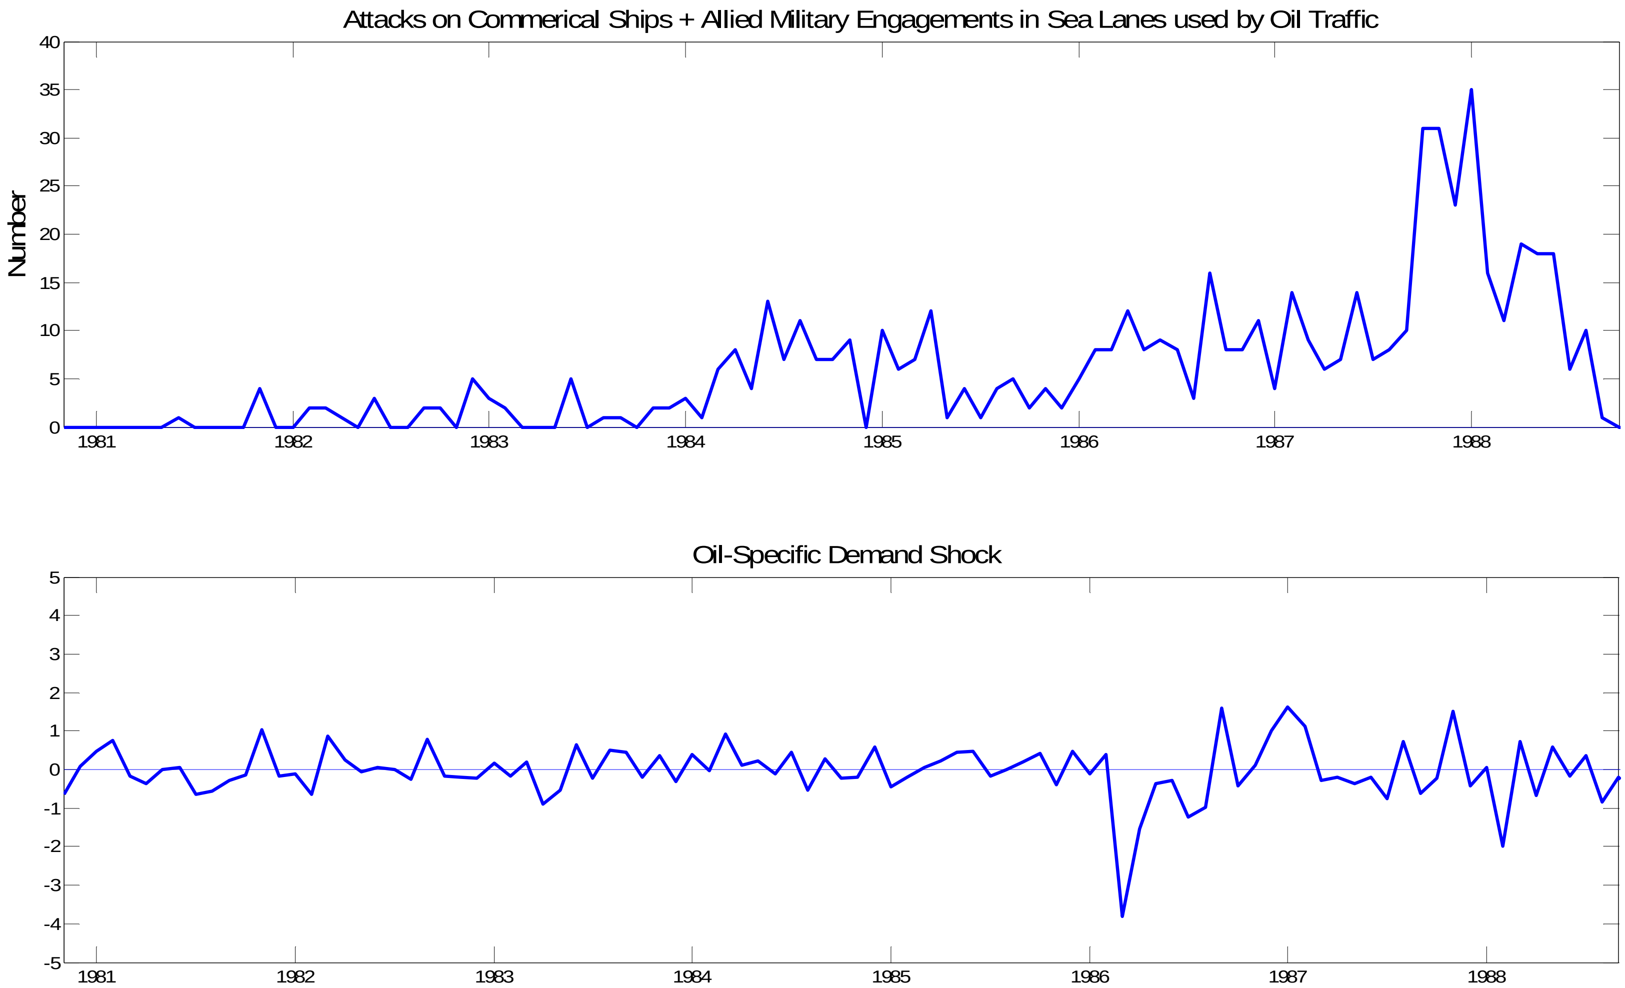

NOTES: See Figures 8 and 9.

Figure 11: 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

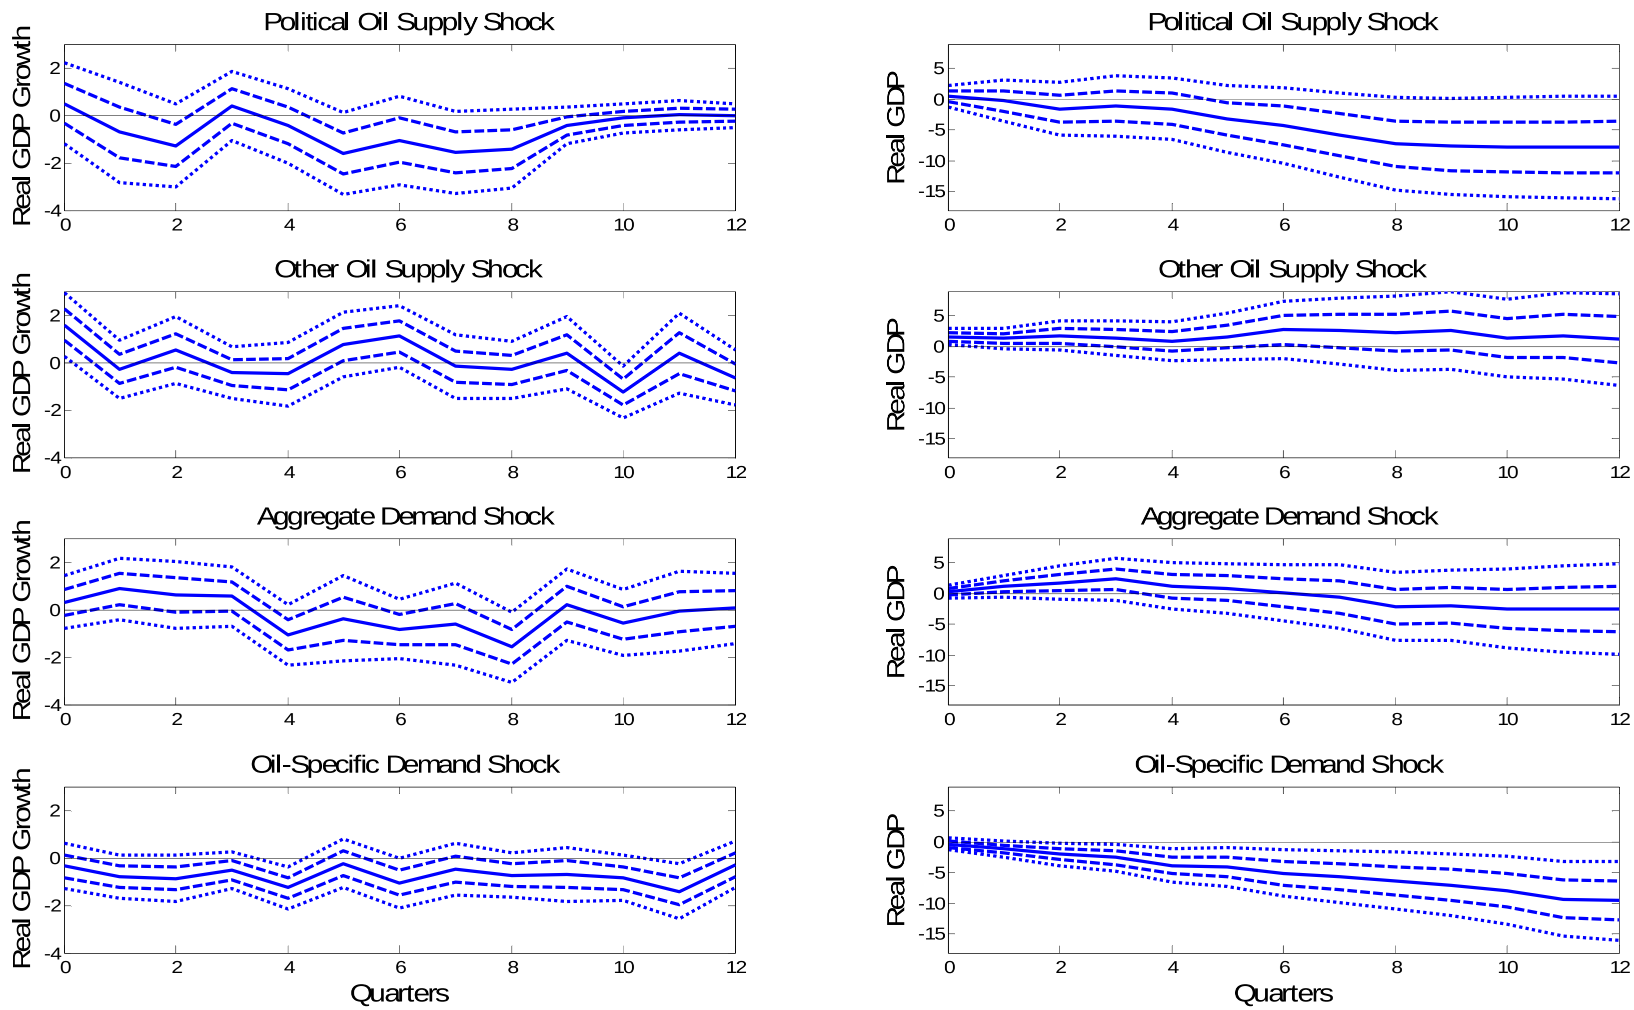

NOTES: All results are based on equations (4) and (4′)

Figure 12: Responses of U.S. CPI Inflation and CPI Price Level to Each Structural Shock
OLS Point Estimates with One and Two-Standard Error Bands

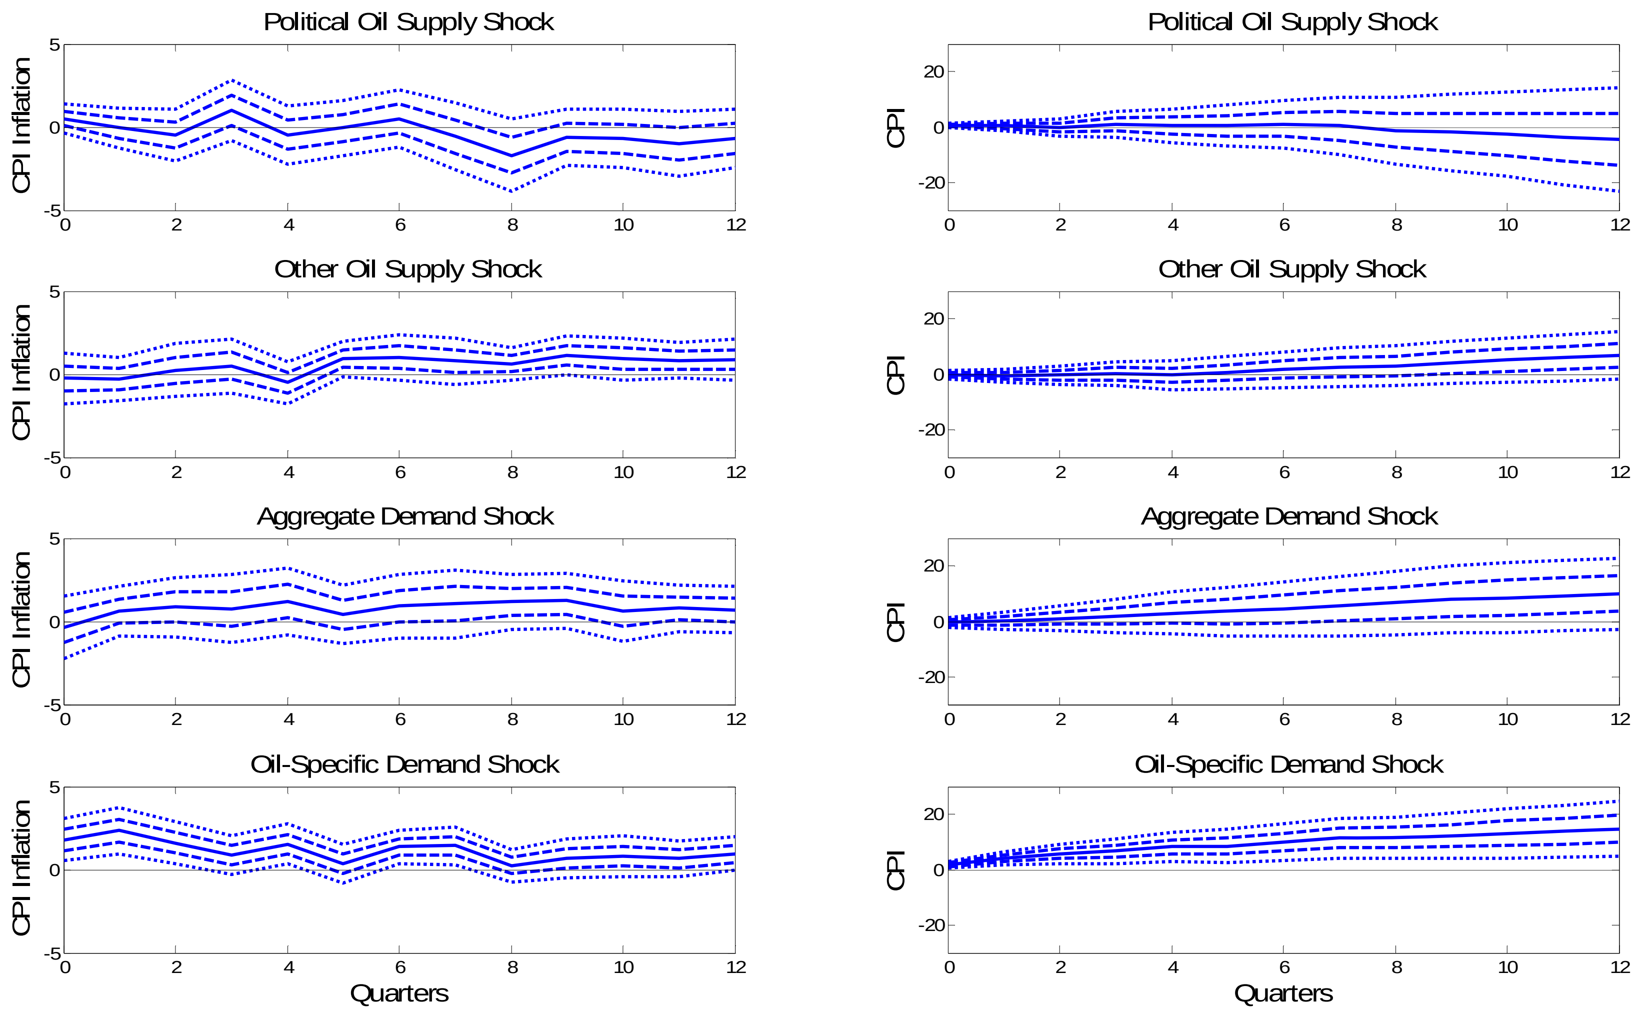

NOTES: All results are based on equations (5) and (5′).

Figure 13: Stagflationary Effects of a Unit Innovation in the Structural Shocks
Conditional Covariance of Inflation and Real GDP Growth with 80% and 90% Confidence Bands

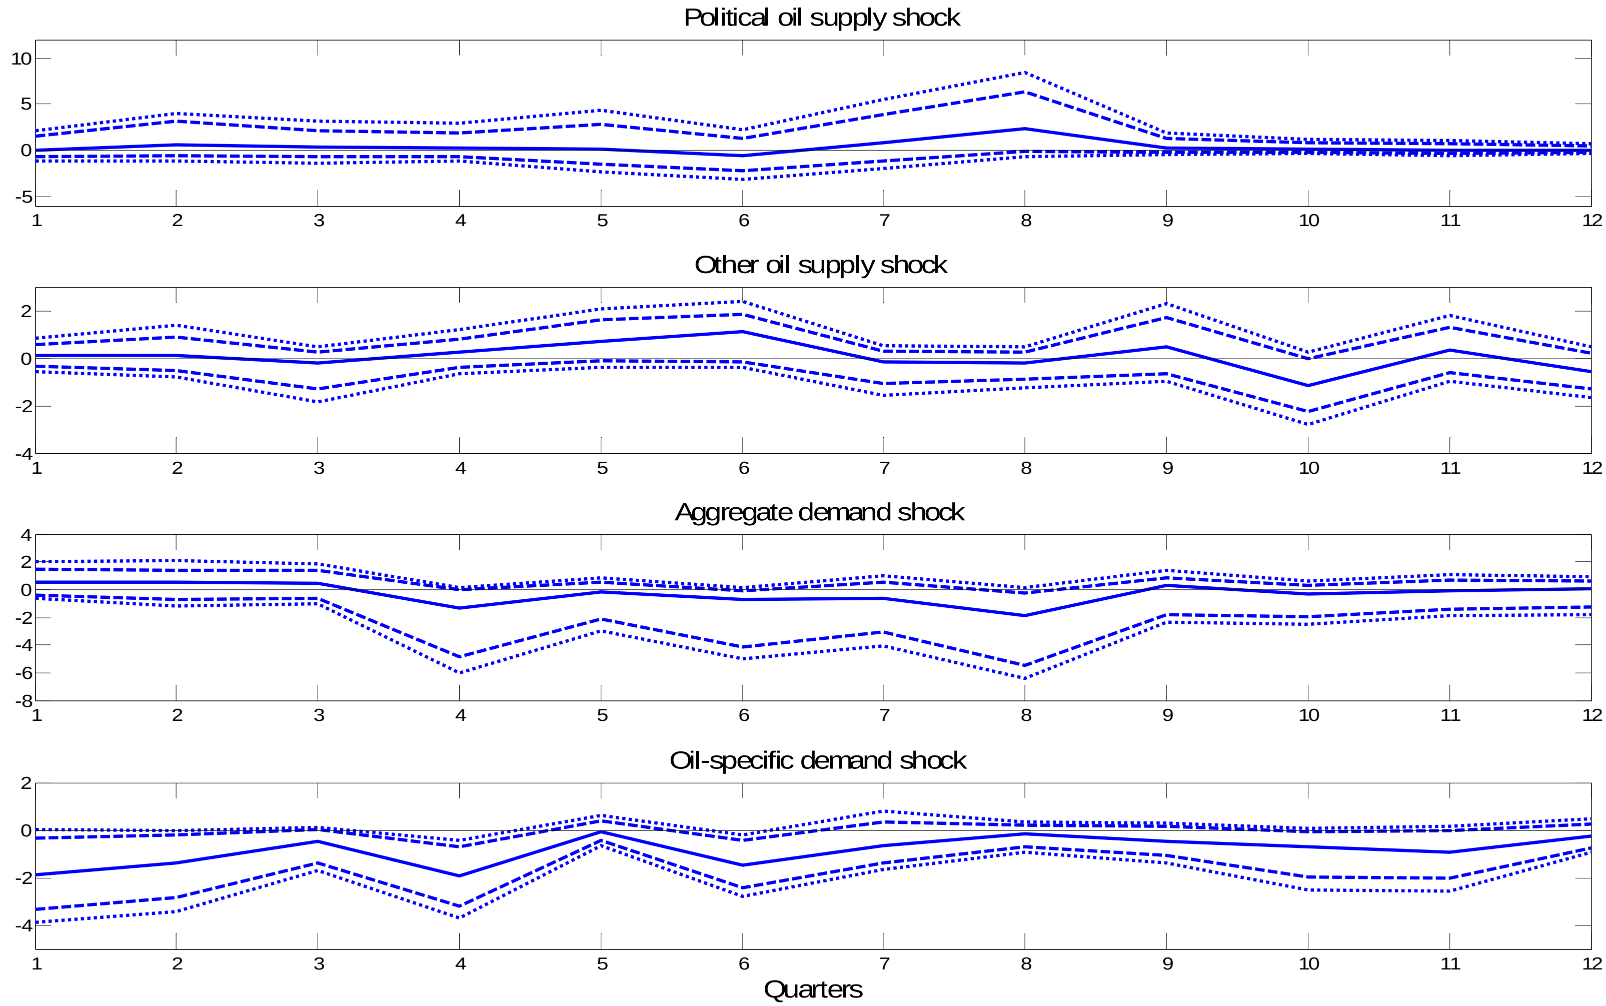

NOTES: The plot shows a statistical measure of the conditional co-movement between real GDP growth and CPI inflation, as defined in Den Haan (2000) and Den Haan and Summer (2004), based on the impulse responses in Figures 11 and 12. Stagflation in the form of rising prices and falling output means that this measure will be negative.

Figure 14: Cumulative Contribution of Each Structural Shock to U.S. Real GDP Growth 1978.IV-2005.III

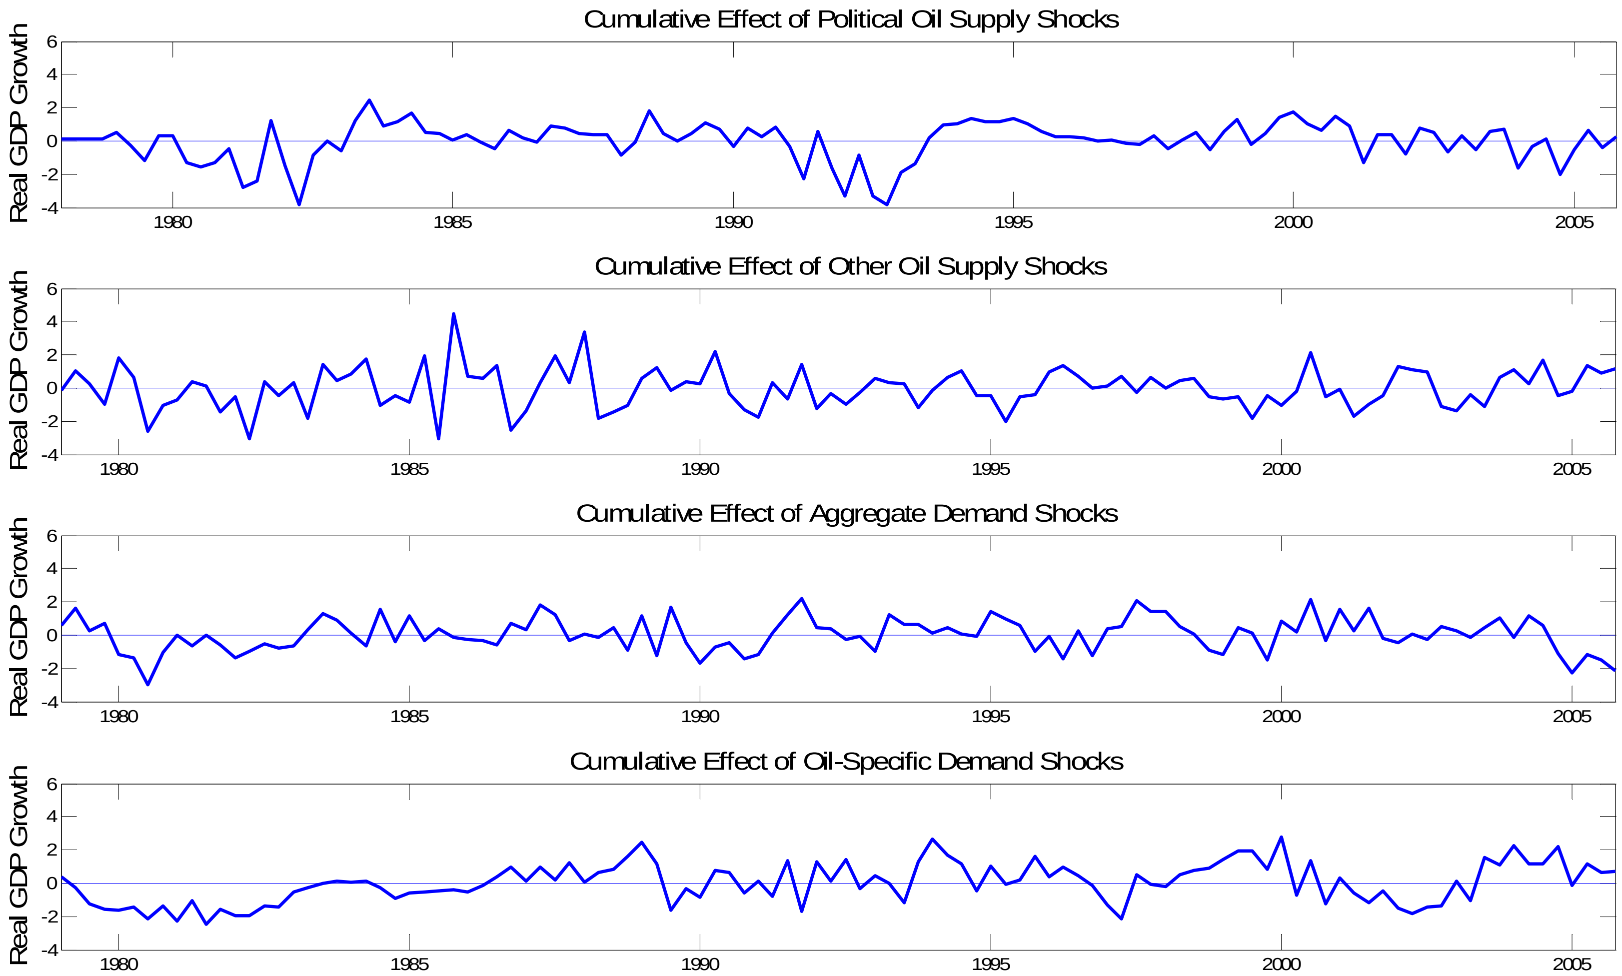

NOTES: All results based on regression models (4) and (4′). The effect is computed by simulating the predicted value of regressions (4) and (4′), respectively, for $x_t = 0 \; \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. All rates are annualized.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

Figure 15: Cumulative Contribution of Each Structural Shock to U.S. CPI Inflation 1978.IV-2005.III

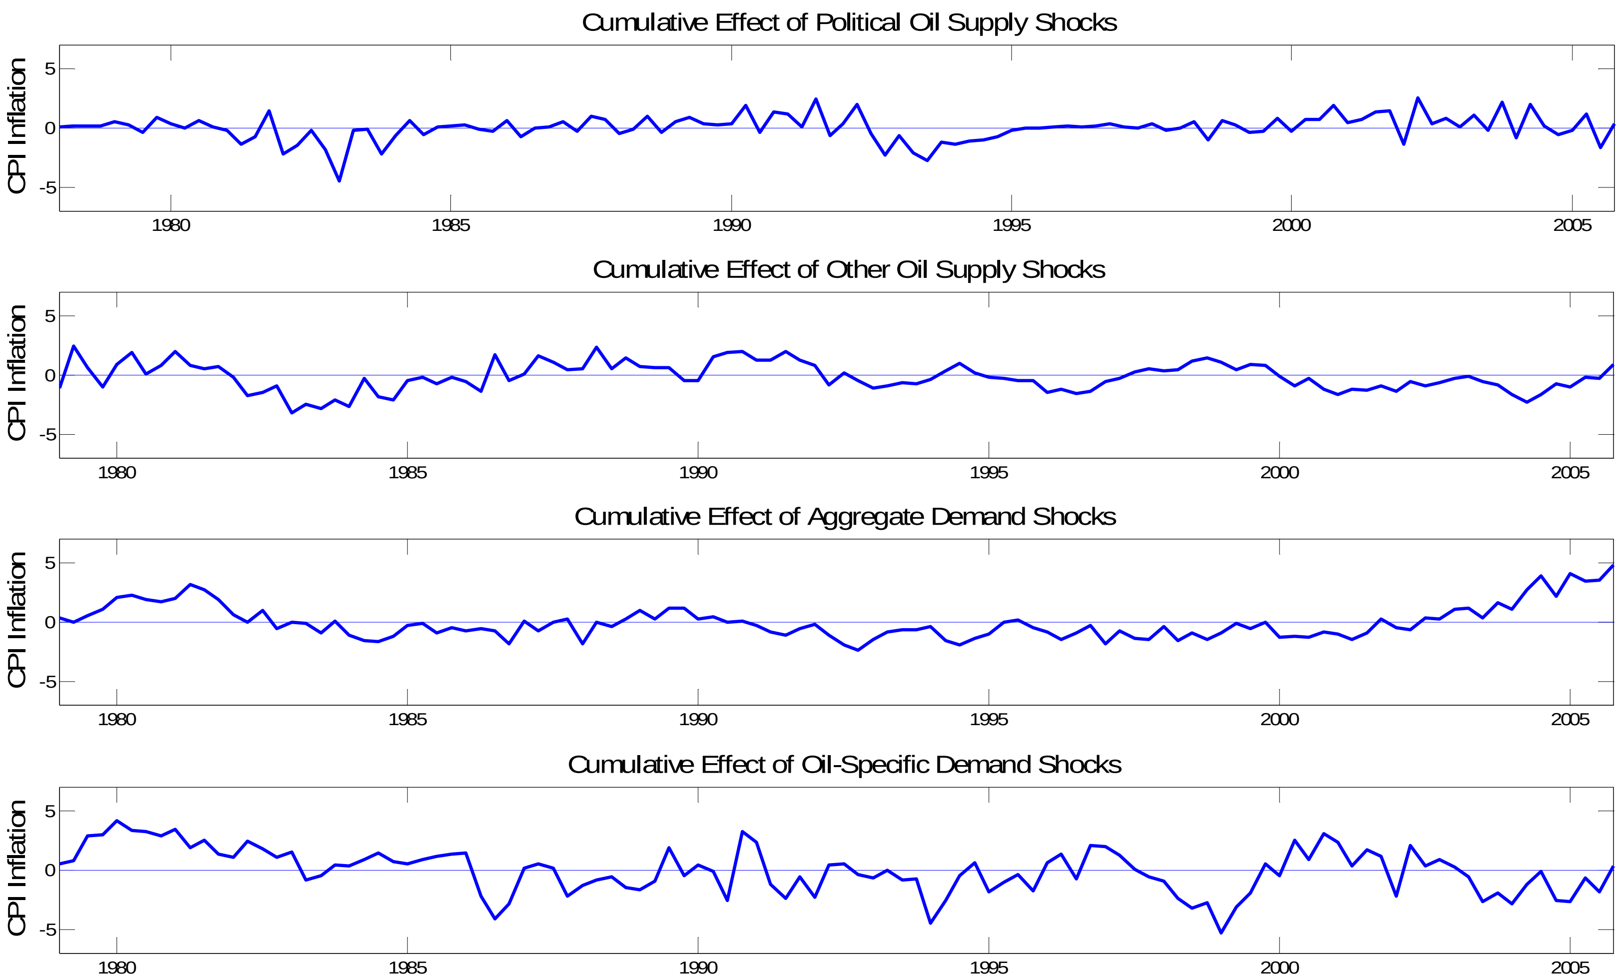

NOTES: All results based on regression models (5) and ($5'$). The effect is computed by simulating the predicted value of regressions (5) and ($5'$), respectively, for $x_t = 0 \; \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. All rates are annualized. The first three observations were discarded as transients.

Figure 16: Cumulative Contribution of All Structural Shocks Combined to U.S. Macroeconomic Real GDP Growth and CPI Inflation 1978.IV-2005.III

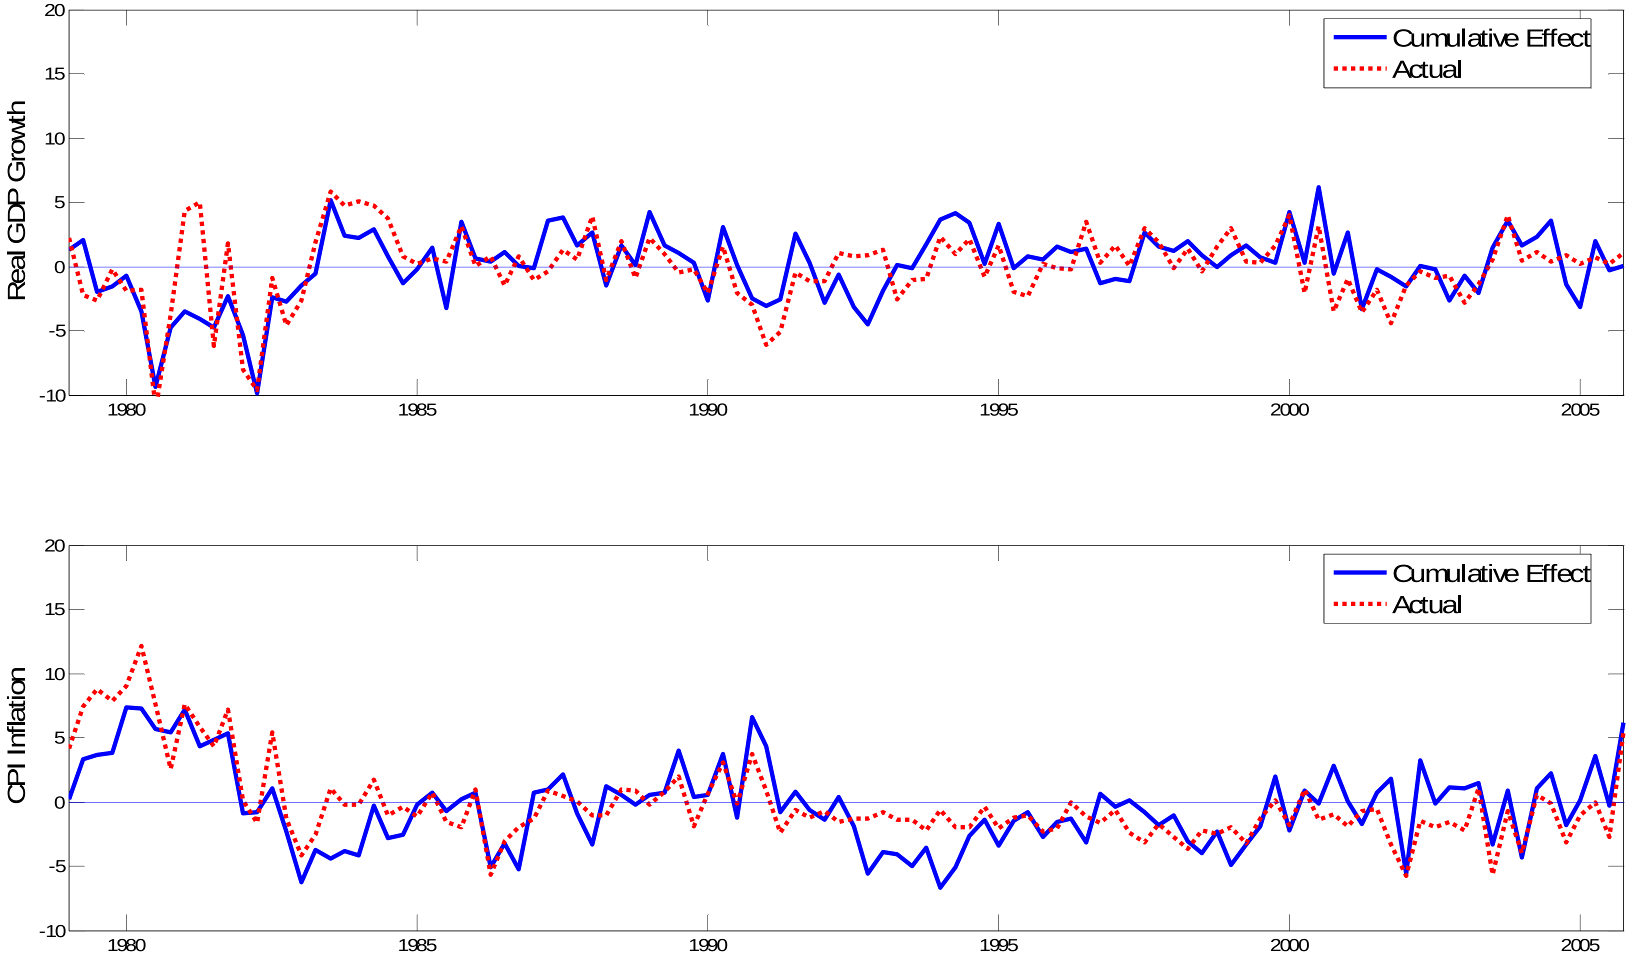

NOTES: See Figures 14 and 15.