

Measuring Oil-Price Shocks Using Market-Based Information*

Tao Wu[†]

Federal Reserve Bank of Dallas

Michele Cavallo[‡]

Federal Reserve Board

This version: October 2009

First draft: September 2006

Abstract

We study the effects of oil-price shocks on the U.S economy combining narrative and quantitative approaches. After examining daily oil-related events since 1984, we classify them into various event types. We then develop measures of exogenous shocks that avoid endogeneity and predictability concerns. Estimation results indicate that oil-price shocks have had substantial and statistically significant effects during the last 25 years. In contrast, traditional VAR approaches imply much weaker and insignificant effects for the same period. This discrepancy stems from the inability of VARs to separate exogenous oil-supply shocks from endogenous oil-price fluctuations driven by changes in oil demand.

JEL CLASSIFICATION CODES: C32, C82, E31, E32, Q43

*For helpful suggestions and comments, we thank Michelle Alexopoulos, Mark Bergen, John Fernald, Luca Guerrieri, James Hamilton, Oscar Jordà, Lutz Kilian, Daniel Levy, Andrea Pescatori, Monika Piazzesi, Trevor Reeve, Christina Romer, David Romer, Eric Swanson, Rob Vigfusson, and seminar and conference participants at the Bank of England, Bank of Italy, Federal Reserve Banks of Chicago, Dallas, and San Francisco, Federal Reserve Board, IMF, Oxford University, the 2007 SED Annual Meeting, the Monetary Economics workshop at the 2007 NBER Summer Institute, the 2008 SCE Conference, the 2009 North American Summer Meeting and the 2009 Far East and South Asia Meeting of the Econometric Society, and the Impulse and Propagation Mechanism workshop at the 2009 NBER Summer Institute. Thien Nguyen, Nina Ozdemir, Michael Simmons, and Tiffany Smith provided excellent research assistance. The views expressed in this paper are solely our responsibility and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas, the Federal Reserve Board, or the Federal Reserve System in general.

[†]Corresponding author. Economic Research Department, 2200 North Pearl Street, Dallas, TX 75201; Tel.:+1 (214) 922-5230; Fax:+1 (214) 922-5194; Email: Tao.Wu@dal.frb.org

[‡]Division of International Finance, 20th and C Streets NW, Washington, DC 20551; Tel.:+1 (202) 452-2607; Fax:+1 (202) 452-6424; Email: Michele.Cavallo@frb.gov

1 Introduction

The relationship between oil-price shocks and the macroeconomy has attracted extensive scrutiny by economists over the past three decades. The literature, however, has not reached a consensus on how these shocks affect the economy, or by how much. A large body of studies relies on various vector autoregression (VAR) approaches to identify exogenous oil-price shocks and estimate their effects. Nevertheless, estimation results are generally inconsistent with the conventional wisdom, that following a positive oil-price shock, real GDP decreases and the overall price level increases. Moreover, the estimated relationship is often unstable over time. This is why, after a careful examination of various approaches, Bernanke, Gertler, and Watson (1997) conclude that “finding a measure of oil price shocks that ‘works’ in a VAR context is not straightforward. It is also true that the estimated impacts of these measures on output and prices can be quite unstable over different samples.”

Traditional VAR-based measures of oil-price shocks exhibit two recurrent weaknesses: endogeneity and predictability. With regard to the first one, VAR approaches often cannot separate the effects of an exogenous oil-price shock from those of an endogenous response of oil prices to other kinds of structural shocks. For instance, the oil price increases that have occurred since 2002 are viewed by many as the result of “an expanding world economy driven by gains in productivity” (*The Wall Street Journal*, August 11, 2006). Such endogenous fluctuations will undoubtedly lead to biased estimates of the effects of oil shocks.

On the other hand, part of the observed oil price changes might have been anticipated by private agents well in advance; therefore, they are hardly “shocks.” Most measures of oil-price shocks in the literature are constructed using only spot oil prices. However, when the market senses any substantial supply-demand imbalances in the future, changes in the spot prices may not fully reflect such imbalances. A number of authors (e.g., Wu and McCallum, 2005; Chinn, LeBlanc, and Coibon, 2005) have found that oil futures prices are indeed quite powerful in predicting spot oil price movements, indicating that at least a portion of such movements may have been anticipated at least a few months in advance. Both these concerns underscore the need to pursue a different approach to obtain more reliable measures of exogenous oil-price shocks.

In this paper, we combine narrative and quantitative approaches to develop new measures of exogenous oil-price shocks that avoid the endogeneity and predictability concerns. We begin by identifying the events that have driven oil-price fluctuations on a daily basis from 1984 to

2007. To achieve this, we first collect information from daily oil-market commentaries published in a number of oil-industry trade journals, such as *Oil Daily*, *Oil & Gas Journal*, and *Monthly Energy Chronology*. This leads to the construction of a database that identifies major oil-related events that have occurred each day since January 1984. We then classify these daily events into a number of different event types based on their specific features, such as weather changes in the U.S., military actions in the Middle East, OPEC proposals on oil production, U.S. oil inventory announcements, etc. (see Table 1). Next, for each event type we construct a measure of oil-price shocks by running oil-price forecasting equations on a daily basis. Finally, shock series from exogenous oil events are selected and aggregated into a single measure of exogenous oil-price shocks. By construction, these shock measures should be free of endogeneity and predictability problems. For robustness, we provide several alternative definitions of exogenous oil-price shocks and construct corresponding shock measures for each of them.

We employ our new, market-information based measures to study the responses of U.S. output, CPI, and monetary policy to exogenous oil-price shocks. We also compare the estimated responses with those obtained following two traditional VAR-based identification strategies that are very popular in the literature. Estimation results reveal substantial and statistically significant output and price responses to exogenous oil-price shocks identified by our market-based methodology. In contrast, responses implied by the VAR-based approaches are much weaker, statistically insignificant, and unstable over time. Moreover, we find that following a demand-driven oil-price shock, real GDP increases and the price level declines. This finding is consistent with scenarios in which oil-price fluctuations are endogenous responses to changes in the level of economic activity rather than reflecting exogenous oil shocks. We argue that traditional VAR-based approaches cannot separate the effects of these two kinds of shocks and consequently lead to biased estimates of the dynamic responses.

Our approach is similar in spirit to the narrative approach pursued in a number of existing studies. Romer and Romer (2004, 2009) adopt it in their analyses of monetary policy and tax shocks, Alexopoulos (2008) and Alexopoulos and Cohen (2009) in the context of technology shocks, and Ramey (2009) in her analysis of government spending shocks. With regard to oil-price shocks, several earlier studies have tried to isolate some geopolitical events associated with abrupt oil-price increases and examine their effects on the U.S. economy. Hamilton (1983, 1985) identifies a number of “oil-price episodes” before 1981, mainly Middle East tensions, and

concludes that such oil shocks had effectively contributed to postwar recessions in the U.S. Hoover and Perez (1994) revise Hamilton’s (1983) quarterly dummies into a monthly dummy series and find that oil shocks led to declines in U.S. industrial production. Bernanke, Gertler, and Watson (1997) construct a quantitative measure, weighting Hoover and Perez’s dummy variable by the log change in the producer price index for crude oil, yet they are not able to find statistically significant macroeconomic responses to oil shocks in a VAR setting. Hamilton (2003) identifies five military conflicts during the postwar period and reexamines the effects of the associated oil shocks on U.S. GDP growth. Finally, Kilian (2008) also analyzes six geopolitical events since 1973, five in the Middle East and one in Venezuela, and examines their effects on the U.S. economy. Our study contributes to the literature by constructing a database of all oil-related events on a daily basis. This allows us to identify all kinds of oil shocks and conduct a more comprehensive analysis than earlier studies. Extracting the “unpredictable” component of oil-price fluctuations using an oil-futures-price-based forecasting model represents another novelty of our work.

More recently, Kilian (2009) has also used information from the oil market to disentangle different kinds of oil-price shocks. In particular, he constructs an index of global real economic activity and includes it in a three-variate VAR, along with data on world oil production and real oil prices. Using a recursive ordering of these variables, he recovers oil-supply shocks, global aggregate demand-driven shocks, and oil-market-specific demand shocks. Although his approach is completely different from ours, the effects on the U.S. economy of all three kinds of structural shocks estimated in his work are quite close to our empirical estimates. This, in turn, corroborates the validity of our approach. We present detailed evidence in later sections.

Our study is also related to the ongoing debate about how the real effects of oil-price shocks have changed over time. For instance, VAR studies, such as those of Hooker (1996) and Blanchard and Galí (2009), have usually found a much weaker and statistically insignificant relationship between their identified oil-price shocks and real GDP growth for the U.S. and other developed economies during the last two or three decades. These results are often cited as evidence that the U.S. economy has become less volatile and more insulated from external shocks, the result of better economic policy, a lack of large adverse shocks, or a smaller degree of energy dependence (e.g., a more efficient use of energy resources and a larger share of the services sector in the economy), all contributing to a “Great Moderation” starting in the first half of the 1980s.

Although we do not challenge this general description of the “Great Moderation,” estimation results presented below reveal a substantial and significant adverse effect of exogenous oil shocks on the U.S. economy, even during the last two and a half decades. Results from VAR studies, in particular the time variation in coefficient estimates, may simply reflect a poor identification strategy.

The rest of our paper is organized as follows. Section 2 describes the methodology we follow to identify the oil-related events and construct our oil-price shock measures. Section 3 illustrates the procedure we use to estimate the macroeconomic effects of oil-price shocks. Section 4 presents our empirical results and compares them with those of earlier studies. Finally, Section 5 offers our concluding remarks.

2 Measures of Exogenous Oil-Price Shocks Based on Market Information

This section describes the derivation of our market-information-based measures of oil-price shocks. The methodology consists of three key steps. First, we conduct a thorough and comprehensive examination of the oil-related events that have driven daily oil-price movements since January 1984 and classify them into a number of event types. Second, for each event type, we construct measures of oil-price shocks by conducting an oil-price forecasting exercise at a daily frequency, so as to capture the unpredictable component of oil-price fluctuations. Finally, we aggregate shock series corresponding to exogenous event types and construct a single measure of exogenous oil-price shocks. For robustness, we also provide several alternative definitions of exogenous oil-price shocks, and for each definition we aggregate the daily shock series of the corresponding event types into a single measure of exogenous shocks.

2.1 A Comprehensive Study of Daily Oil-Related Events

The first step of our methodology is to identify the events behind the observed oil-price fluctuations.¹ For this purpose, we collect information from a number of oil-industry trade journals, such as *Oil Daily* and *Oil & Gas Journal*. We then cross check this information with other sources, including such government publications as *Monthly Energy Chronology*, published by the Energy Information Administration, a statistical agency in the U.S. Department of Energy.

¹To be consistent with the literature and oil-industry terminology, throughout the paper we refer to the spot oil price as the price quoted on one-month futures contracts of West Texas Intermediate light sweet crude oil traded on the New York Mercantile Exchange (NYMEX). This is also the spot price that most of the financial press reports every day (see, e.g., *The Wall Street Journal*).

Our sample runs from January 3, 1984, to October 31, 2007, a total of 5,971 trading days. For each trading day in our sample period, we collect information on major oil-related events that occurred on that day from the market commentaries or reviews published in the above-mentioned trade journals. We consider an event as major if it had significantly affected oil prices and had received extensive coverage in the corresponding daily market analysis. After a thorough reading of these market commentaries and reviews, we classify oil-related events into 22 different types (see Table 1), such as weather changes in the U.S., military actions in the Middle East, OPEC development on oil production and U.S. oil inventory announcements. Based on this analysis, we assign one numerical code to each trading day, or more than one code if more than one type of oil-related event occurred on the same day.

We conduct the event study at a daily frequency because the oil market, like other well-developed financial markets, is highly volatile, responding immediately to economic, political, and industry-specific news. Choosing a lower frequency, such as monthly or even weekly, would likely result in a situation in which several events might have happened within the same period, making it difficult to measure the magnitude of the shock that each event brought to the oil market. The daily frequency is the highest for which we can find relevant market information.

To minimize the possibility that both the interpretation of market-based information and the event classification may be biased by the analyst's subjective predispositions, we have conducted a thorough *content analysis*, a practice widely used in marketing literature (see, e.g., Kassarian 1977, and Levy, Dutta, and Bergen, 2002). Specifically, three independent analysts have been engaged in reading the documents and classifying the events. The results have been compared to make sure that they are consistent with each other in most cases.

Column 3 of Table 1 shows the observed relative frequencies of oil-related events from 1984 to 2007. Excluding the days with no particular reason observed or when the price movement was driven by speculation, the most frequent event is "OPEC development on oil production" (741 trading days, 12 percent of the sample), followed by "U.S. oil inventory announcements" (730 days, 12 percent), and "political development in the Middle East" (476 days, 8 percent).² Oil production or transportation disruptions both in the U.S. and outside the U.S. (types 3 and

²It is well known that OPEC announcements on production quota are not always fully enforced. Therefore, on the day of the news, the reaction of oil prices to these announcements or to related developments (e.g., various rumors) can be viewed as reflecting oil traders' probabilistic assessment of the effects of such announcements or developments on the supply of oil.

4) affected oil-price movements on 486 trading days, about 8 percent of the sample.³

2.2 Two Measures of Oil-Price Shocks

The next step is to quantify the magnitude of the shocks implied by each oil-related event on a daily basis. Two approaches are adopted. The first one is based on a modified version of the oil-price forecasting model in Wu and McCallum (2005). In particular, for each trading day, we regress the realized oil-price changes on the spreads between oil spot and futures prices at different horizons quoted by the end of the previous trading day, with a rolling sample consisting of the previous 200 trading days.⁴ Our estimating equation is:

$$\log P_{t+1}^S - \log P_t^S = \alpha + \sum_{j=2}^6 \beta_j (\log P_{j,t}^F - \log P_t^S) + \varepsilon_{t+1}, \quad (1)$$

where P_t^S and P_{t+1}^S are the spot prices at t and $t + 1$, respectively, $P_{j,t}^F$ denotes the j -month oil futures price at time t , α and β_j 's are the estimation coefficients, and ε is a white-noise error term. We then calculate the unpredicted change in spot price as realized at $t + 1$ and define the “predicting error” as our shock measure for the day.

Equation (1) incorporates term structure information on futures-spot spreads in forecasting future oil-price movement. This equation is in the same spirit as the bond-yield forecasting model in Cochrane and Piazzesi (2002, 2005), who also use information embodied in term spreads of interest rates at all available horizons to forecast future bond-yield movement, without imposing the Expectations Hypothesis. Wu and McCallum (2005) compare the out-of-sample forecasting performance of such a “futures-spot spread” model with that of several other models and conclude that the futures-spot spread model performs the best, particularly when the forecasting horizons are within the next few months. On the other hand, we exclude price quotes on futures contracts beyond six months from the equation, as the futures market becomes substantially less liquid for those horizons, and consequently, the quoted futures prices become a much less accurate measure of oil-price expectations. Wu and McCallum (2005) have also found that the out-of-sample performance of the futures-spot spread model becomes much worse when the forecasting horizon goes beyond one year.

Alternatively, we measure the magnitude of the shock as simply the change in the logarithm

³Other types of less frequent events include changes in the market’s expectations of U.S. oil inventories, U.S. weather changes, and changes in oil demand in the U.S. and outside the U.S.

⁴Changing the length of the rolling sample has negligible effects on the forecasting results.

of the spot oil price on the day. This quantitative approach is consistent with the belief that oil futures prices do not have any predictive content for future oil price (Alquist and Kilian, 2009) and that the log oil price follows a random walk. We call the shock measure based on this approach the “log-price change” measure and, in later econometric analysis, use it along with the “predicting error” measure described above.⁵ However, it is important to bear in mind that both measures are formulated on a daily basis and that, through the identification of the oil-related events, they will be constructed around days of exogenous events, implying that both will be legitimate measures of exogenous oil shocks.

2.3 What Does Exogeneity Mean?

Having classified the daily oil-related events and constructed daily shock measures for each event type, the next step is to construct a single series of exogenous oil-price shocks, by combining shock series related to all types of events that are exogenous. For this purpose, we need to first explicitly define which notion of “exogeneity” we are referring to when we talk about exogenous oil-price shocks.

Ideally, genuine exogenous shocks would be defined as exogenous with respect to the U.S. economy in the most rigorous sense. Therefore, any event that could possibly correlate with the U.S. economy cannot be a genuine exogenous shock. For instance, weather changes in New England (event-type 1 in Table 1) do not qualify as exogenous oil shocks, as such changes may affect not only oil demand but also utility output, construction activities, and retail sales. For the same reason, oil production and transportation disruptions in the U.S. (event-type 3) do not qualify as exogenous oil-supply shocks, as the energy industry is a substantial component of the U.S. economy. Furthermore, military conflicts (event-types 10 through 12) do not qualify, as they may affect U.S. defense spending. Even when the U.S. is not directly involved, one can still argue that the military buildup following such conflicts would make them correlate with real GDP in addition to their impacts on oil supply and demand (Ramey and Shapiro 1998, Ramey 2009). For the same reason, political developments (event-types 7 through 9) may not qualify as genuine exogenous shocks. These concerns essentially rule out most of the oil-related events

⁵Hamilton (2009a) conducts a literature survey and notes that, while many empirical studies “found that the spot oil price provides as good or even a better forecast of the future oil price than does the futures price,” they generally also “failed to reject the hypothesis that the oil futures price embodies a rational expectation of the future spot price.” Apparently, his conclusion also supports our choice of running two alternative forecasting models to extract the unpredictable component.

that many researchers consider exogenous. In fact, during the past 25 years, there were only six trading days in which non-U.S. weather changes (event-type 2) had significantly affected the oil market, and there was no new oil field discovered anywhere (event-types 5 and 6) that had a noticeable impact on the oil market.

In light of these considerations, we take a step back and allow for some degree of ambiguity in our econometric analysis. We reiterate that different assumptions may lead to quite different interpretations of “exogeneity.” Therefore, instead of providing one single series and treating it as the exogenous oil shock measure, we provide three definitions of exogenous oil-related events and construct the corresponding series of exogenous oil-price shocks,⁶ examining their dynamic macroeconomic effects in later sections:

(1) Our “baseline” definition consists of event-types 1 through 12, including the U.S. and non-U.S. weather changes, oil production or transportation disruptions, and political and military actions. These are typically the kinds of events many researchers consider exogenous.⁷

(2) Alternatively, we make a “narrow” definition, consisting of only event-types 2 through 9, that is, non-U.S. weather changes, oil production or transportation disruptions, and political developments. Compared with our “baseline” specification, this definition excludes U.S. weather changes and military actions around the world. We exclude these events because they are more likely to be correlated with real GDP growth in the U.S.

(3) Finally, our “broad” definition of exogenous oil-related events consists of event-types 1 through 12 and 15 through 17. This definition includes not only the event types in our baseline definition, but also events such as oil and gas inventory announcements (for instance, the Energy Information Administration’s weekly inventories reports) or changes in market expectations of oil inventories. These oil-related events are described in Kilian (2009) as “precautionary demand shocks,” as they are likely to be associated with concerns about the availability of future oil supplies.

We choose not to include in our three definitions events such as OPEC or non-OPEC oil exporters’ changes of their production plans or proposals (event-types 13 and 14). In fact, these events are likely to reflect oil producer’s endogenous responses to developments in other sectors

⁶We thank Christina and David Romer for this suggestion.

⁷Interestingly, we also find that political developments and military tensions in non-oil-producing countries (event-types 9 and 12) were not mentioned even once in oil-market analyses during the past 25 years. This indicates that such events had essentially no effect on global oil supply and demand.

of the world economy. For a similar reason, we choose not to include in any of our definitions changes in oil demand, such as global economic development, and improvements in oil usage efficiency (event-types 18 and 19). However, as these events represent a very important portion of the developments that have occurred in the global oil market, in the following sections we examine their macroeconomic effects separately.⁸

2.4 Constructing Monthly Oil Shock Series

The final goal of our work is to explore the effects of various kinds of oil shocks on the U.S. economy. As the highest available frequency for most macroeconomic data is monthly, to facilitate the econometric analysis, we aggregate our daily oil shock series into monthly series. Specifically, for each trading day, we attribute the daily shocks to an event type based on the code assigned to that particular day.⁹ We then aggregate the 22 daily shock series into the same number of monthly series. Finally, for each of our definitions of exogenous oil-related events, as well as for any other combination of oil-related events that is of potential interest, we construct a monthly oil-price shock measure to be used in our econometric analysis later.

Figures 1A and 1B display our market-information-based measures of oil-price shocks, with the shocks defined as the “predicting error” from equation (1) and as the “log-price change,” respectively. To improve the readability of the time plot of shock series, in these figures we display the annual average of monthly series (the original monthly series are shown in Figures 2A and 2B). The three market-information-based measures are quite similar. Consider, for example, the shock measures constructed as the “predicting error”: The correlation between the “baseline” measure and the “broad” measure is 78 percent, between the “baseline” and the “narrow” measures is 72 percent, and between the “narrow” and the “broad” measures is 55 percent. On the other hand, shock series constructed following the two quantitative approaches (i.e., the “predicting error” and the “log-price change”) are also very similar: The correlation between the “baseline” measures constructed in these two different ways is 90 percent, and the corresponding correlations are 87 percent for the “broad” measures and 88 percent for the “narrow” measures.

⁸To facilitate possible future work by other researchers, in constructing our database we have preserved as much primitive information as possible about oil-market developments. Interested readers can select the definitions of exogenous oil-related events of their own choice and construct the corresponding alternative measures.

⁹If multiple codes are assigned to the day, the shocks will be equally divided among corresponding event types.

Figure 1C displays two VAR-based measures that are widely used in the literature. The first one is based on the “net oil price increase” (NOPI) indicator of the oil market proposed by Hamilton (1996), and the second one is based on the log change in the producer price index (PPI) for crude oil, as, for example, in Bernanke, Gertler, and Watson (1997) and Blanchard and Galí (2009). Both these VAR-based measures are the estimated residuals from a recursive VAR that includes macroeconomic variables and an indicator of oil prices, with the oil-price indicator ordered as the last variable in the VAR system. The only difference between the two is whether it is the NOPI or the log change in the PPI for crude oil that enters the VAR. In the later analysis, we refer to them as the “asymmetric” VAR-based measure and “symmetric” VAR-based measure, respectively, and explain their construction in detail in Section 3.

Figures 2A and 2B compare our market-information-based measures with the two traditional VAR-based measures on a monthly frequency. As shown in the figures, there are substantial similarities and, at the same time, significant differences between the traditional VAR-based measures and our market-information-based measures. Both kinds of oil-price shock measures capture major oil-price spikes reasonably well, for instance, during the periods March-April 1986, August-September 1990, December-February 1991, April 1999, and September-October 2004. However, the magnitudes of the shocks are somewhat different. In fact, the symmetric VAR-based measure is the most volatile series of the three, and the asymmetric VAR-based measure is the least volatile. The correlation between our “baseline” oil-shock measure (“predicting error”) and the asymmetric VAR-based measure is 24 percent, while the correlation between the “baseline” and the symmetric VAR-based measures is 23 percent. Correlation between the two VAR-based measures is 53 percent.

Such differences are not surprising. First, the two VAR-based measures are residuals from vector autoregressions that include macroeconomic variables, whereas the market-information-based measures are either residuals from an oil-price forecasting equation that does not incorporate macroeconomic variables or simply log changes of the oil price. Second, and more important, the approaches adopted to recover the oil-price shocks are completely different. We adopt an “event-study” approach and rely on market information to identify exogenous oil-price shocks, whereas the traditional VAR-based measures rely on the recursive ordering of the corresponding variables. For example, if the price of oil rises sharply following an expansion in the level of global economic activity, the traditional VAR approaches may interpret the increase in

the price of oil as a shock. Our methodology, in contrast, correctly classifies it as an increase in oil demand due to economic development (event-type 18 or 19), and will, correspondingly, exclude it from our exogenous oil-price shock measures.

3 Estimating the Effects of Oil-Price Shocks

Next, we examine the dynamic effects of oil-price shocks on the U.S. economy. For this purpose, we estimate a vector autoregressive model with exogenous variables, where the set of exogenous variables includes a deterministic time trend and, more importantly, our measure of exogenous oil-price shocks. In the econometrics literature, this type of model is sometimes referred to as a VARX model or as a rational distributed lag model (see Lütkepohl, 2005, chapter 10). Thus, our estimating system of equations is:

$$X_t = A_0 + A_1 t + A_2(L) X_{t-1} + B(L) O_t + \varepsilon_t, \quad (2)$$

where X_t is a vector that contains the log of real GDP, the log of the consumer price index (CPI), the level of the federal funds rate, and the log of the real price of oil, defined as the difference between the log of the producer price index (PPI) for crude oil and the log of the CPI.¹⁰ The variable O_t is an oil-price shock measure, and it represents the observable exogenous input variable, which is determined outside of the system in (2).¹¹ A_0 and A_1 are vectors of coefficients, while $A_2(L)$ and $B(L)$ are two finite-order polynomials in the lag operator L . Finally, t is a time trend, and ε_t is a vector of white noise and mean-zero i.i.d. error terms. The estimated dynamic responses of the endogenous variables in X_t to an oil-price shock k periods ahead are given by the point estimate of the coefficients on L^k in the expansion of the rational transfer function, $[I - A_2(L)L]^{-1} B(L)$. A similar strategy is also adopted by Christiano, Eichenbaum, and Evans (1999) and Burnside, Eichenbaum, and Fisher (2004) in estimating the effects of monetary and fiscal policy shocks, respectively.

¹⁰Our choice of the endogenous variables included in the vector X_t is very similar to that in Bernanke, Gertler, and Watson (1997), except that their VAR also includes a commodity price index to capture the effect of monetary policy shocks. As the primary focus of our study is oil-price shocks, we choose not to include the commodity price index in the vector of endogenous variables, X_t , similarly to Blanchard and Galí (2009).

¹¹Our estimation system is based on the underlying identifying assumption that, within the period, there is no contemporaneous feedback from the dependent variables to the shock measure O_t , implying, therefore, that the shock is treated as predetermined.

We estimate the model in (2) using six lags.¹² The sample consists of monthly data, with the sample period running from January 1984 to October 2007. Since the highest frequency available for real GDP is quarterly, following the work of Bernanke, Gertler, and Watson (1997), we adopt the method of Chow and Lin (1971) to obtain a monthly indicator for real GDP.¹³ In an earlier version of our work, we have also estimated a univariate distributed lag model, the same kind of strategy adopted by Ramey and Shapiro (1998) and Kilian (2009), and the estimated effects obtained there are very similar to those implied by the VARX model in (2).

To estimate the effects of oil-price shocks, we substitute the market-information-based shock measures from Section 2 (“baseline,” “broad,” or “narrow” definition of exogeneity, and shock magnitudes calculated by either “log-price change” or “predicting error” methods) for O_t , one at a time. As mentioned earlier in Section 2, for comparison we also estimate the impulse responses implied by two traditional VAR-based oil-price shock measures. The first measure is the asymmetric VAR-based measure, which is constructed using the “net oil-price increase” (NOPI) indicator proposed by Hamilton (1996). The NOPI is defined as the maximum between zero and the difference between the log of the current oil price and the maximum value of the log of the oil price during the preceding year.¹⁴ The asymmetric VAR-based measure is thus the estimated residuals from the last equation in a recursive four-variable VAR that includes, in the following order, the log of real GDP, the log of the CPI, the level of the federal funds rate, and the NOPI. The second VAR-based measure, the symmetric one, is constructed in a similar way, except that it is the change in log oil price, rather than the NOPI, that enters as the last-ordered variable. Bernanke, Gertler, and Watson (1997) and Blanchard and Galí (2009) build their VAR systems in a very similar fashion.

¹²One alternative would be to choose the lag order using standard methods such as AIC or BIC. However, this would result in the selection of different lag orders for different shock measures. For comparison purposes, we decide instead to impose the same number of lags for all shock measures.

¹³As interpolators, we use the monthly series for industrial production and total capacity utilization.

¹⁴This indicator detects increases that establish new highs relative to most recent readings and that do not reverse previous decreases. In a recent study, however, Kilian and Vigfusson (2009) challenge the use of asymmetric VAR models on the basis of little evidence against the symmetry hypothesis in response to oil-price shocks.

4 Empirical Results

4.1 Impulse Responses of Market-Information-Based Exogenous Oil-Price Shocks

Figures 3A through 3C display the estimated impulse response functions for real GDP, the CPI, the federal funds rate, and the real price of oil to an oil-price shock constructed using our market-information-based methodology. As the impulse responses of the real price of oil will be generally different when different oil-price shock measures are substituted into the system of equations (2), to facilitate the comparison of their macroeconomic effects, we normalize these oil-shock measures so that the peak response of the real price of oil is 10 percent.¹⁵ The magnitude of this normalization is roughly equivalent to 1.75 times the estimated standard deviations of the market-information-based shock measures. Statistical inference on the point estimates of the impulse responses is obtained through a traditional residual-based bootstrap method with 1,000 replications, and the resulting 95 percent standard percentile confidence intervals are denoted by a shaded area in our figures.

The estimated impulse responses fit quite well with the conventional-wisdom view about the macroeconomic effects of exogenous oil-price shocks, that following a positive oil-price shock, real GDP declines and the overall price level increases. Consider, for example, the case of the “baseline, log-price change” shock measure (Figure 3A, left column). In response to the shock, real GDP gradually declines, with the largest response (in absolute value, same below) arriving about 18 months after the shock. The response becomes statistically significant three months after the shock, and remains significant throughout the 24-month horizon. These point estimates imply a substantial impact of exogenous oil-price shocks on the real economy. Over the 24 months following the shock, the implied cumulative output loss is equivalent to 6.8 percent of a month’s real GDP, or about 0.6 percent of annual real GDP in two years. The CPI shifts up immediately on impact, and the peak response arrives three months after the shock. The price increase is both substantial and persistent, on average 14 basis points higher during the 24-month horizon following the shock, and the price level remains 10 basis points higher than

¹⁵Blanchard and Galí (2009) normalize the size of the shock so that it induces an increase in the oil price by 10 percent *on impact*. As can be seen in Figures 3Aa through 3C, our estimation implies that the responses of the real price of oil are usually hump-shaped, with the peak response arriving in the first or second month after the shock. Therefore, we choose to normalize the size of the shocks according to their largest responses rather than their impact responses. In most cases, the normalized shock sizes are very similar.

its preshock level even 24 months after the shock.

In response to the initial rise in the CPI, the federal funds rate rises by a few basis points in the first three months after the shock. However, with real GDP continuing to decline and inflation gradually decelerating, monetary policy becomes more accommodative. The 24-month cumulative decline in the federal funds rate reaches 2.6 percentage points, or about 11 basis points lower than its preshock level each month on average. The response of the real oil price to the shock is hump-shaped, with the peak arriving one month after the shock, with the oil price increase remaining statistically significant even eight months after the shock. Estimates of the impulse responses when shock sizes are calculated according to the oil-price forecasting equation (1) are fairly similar (Figure 3A, right column), with the responses remaining statistically significant and persistent.

The impulse responses to our “broad” and “narrow” measures of exogenous oil-price shocks are also similar. As shown in Figure 3B, in response to a “broadly” defined exogenous oil-price shock, real GDP declines immediately. The output loss is fairly persistent and even stronger than the one implied by our “baseline” definition, with the 24-month cumulative output loss reaching 8.4 percent, or 0.7 percent of a year’s real GDP in two years. The CPI rises on impact after the shock, and the price increase remains significant for the first six months following the shock. Again, the federal funds rate initially rises in response to the increase in the CPI, but it then declines as the output loss persists. A “narrowly” defined exogenous oil-price shock also leads to a persistent output decline and a price increase, and monetary policy is also in this case restrictive at first, switching subsequently to accommodative as the output decline continues (Figure 3C, left column).

4.2 Impulse Responses of VAR-Based Measures of Oil-Price Shocks

The estimated impulse responses displayed in Figures 3A through 3C are substantially larger and more significant than earlier estimates of the effects of oil-price shocks in the literature, in particular the output responses. For example, Bernanke, Gertler, and Watson (1997) estimate the macroeconomic effects of oil-price shocks in a VAR setting, using four alternative indicators of oil-price shocks: (1) changes in the log of the nominal PPI for crude oil, (2) the Hoover-Perez dummies for political and military events in the Middle East, scaled by the log change of the nominal PPI for crude oil, (3) the indicator proposed by Mork (1989), i.e., positive monthly changes in the log of the real price of oil, and (4) Hamilton’s NOPI measure. Bernanke,

Gertler, and Watson estimate the VAR over the period 1965-1995, and find that none of those specifications generate a statistically significant output response to an oil-price shock. Moreover, with the federal funds rate increasing persistently in response to a higher price level after the shock, they argue that it is hard to determine how much of the output decline is the direct result of the oil shock, rather than the indirect result of the tighter monetary policy. This is an evident example of the identification issues the VAR suffers from. Therefore, they conclude that “finding a measure of oil price shocks that ‘works’ in a VAR context is not straightforward.”

Hooker (1996) employs a similar VAR approach and examines two specifications, with the oil-price indicator defined as the log change in the nominal oil price and the log level of the real oil price, respectively. Interestingly, he finds that oil prices did not Granger-cause U.S. GDP or unemployment from 1973 to 1994. Rather, GDP growth exhibited a large *positive* response to an oil price increase for about four quarters and then quickly returned to its preshock level, contradicting the conventional-wisdom view on the macroeconomic effects of oil-price shocks. Bernanke, Gertler, and Watson (1997) report a similar output response when using the log change in the nominal oil price as an indicator of the state of the oil market. However, by using his “net oil-price increase” measure in a univariate autoregressive model, Hamilton (1996) finds a negative output response to an oil-price shock after 1973, although the estimated response is substantially weaker than his pre-1973 estimate and is statistically insignificant.

To illustrate the differences between the estimated macroeconomic effects implied by our market-information-based measures and those implied by the traditional measures, we construct two VAR-based measures of shocks, the asymmetric and symmetric measures as defined in Section 3, and estimate the same VARX system as in (2). We substitute these two VAR-based measures for the exogenous input variable O_t , one at each time.¹⁶

The left column in Figure 4 displays the impulse responses to the asymmetric VAR-based shock measure, constructed in the same way as Hamilton (1996). Following the shock, real GDP declines and the CPI rises. However, the output response is no longer statistically significant and is also substantially weaker than the response implied by our market-information-based measures. For instance, the 24-month cumulative output loss implied by the asymmetric VAR-based measure is only 1.7 percent, only a quarter of the 6.8 percent output loss implied by our “baseline” measure shown in Figure 3A (left column). The responses of the CPI and the federal

¹⁶As above, the shock sizes are normalized so that the peak response of the oil price is 10 percent.

funds rate are also substantially weaker and become statistically insignificant.

The output response to an oil-price shock implied by the symmetric VAR-based measure (i.e., with the log of the real oil price entering last in the VAR), shown in the right column of Figure 4, is slightly stronger than the response implied by the asymmetric VAR-based measure. In particular, the cumulative output loss in the 24 months after the shock is 2.9 percent, although still less than half of the 6.8 percent cumulative output loss implied by our “baseline” market-information-based measure. More importantly, the response of real GDP remains statistically insignificant during most of the 24-month period following the shock. These estimates confirm the findings of earlier studies in the literature, that VAR-based identification strategies usually yield a weak and statistically insignificant output response.

It is also worth noting that the point estimates and the statistical significance of the output response reported in the right column of Figure 4 are quite close to those in Blanchard and Galí (2009). They have estimated a similar VAR using the log of the oil price as an indicator of the state of the oil market. In particular, in their second subsample period (1984:Q1 to 2005:Q4), the cumulative real GDP loss is about 1.6 percent of quarterly GDP over three years (see their Figure 6a), or about 0.13 percent of annual GDP each year on average. These estimated cumulative output losses are quantitatively very close to our estimate of a cumulative loss of 2.9 percent of monthly GDP in two years, i.e., 0.12 percent of annual GDP loss each year on average. This is not surprising. In fact, our sample period (1984:M1 to 2007:M10) overlaps with theirs considerably, the VARs are constructed in a similar fashion, and the shock size is normalized by a similar magnitude. Blanchard and Galí have also found that in the 1960s and 1970s, the response of U.S. GDP is substantially larger and statistically significant. Because of this finding, they conclude that oil-price shocks have had a smaller effect on the U.S. economy since 1984. Our estimation results paint a different scenario: even during the past two decades, exogenous oil-price shocks have continued to exert substantial and significant impacts on the U.S. economy and that the implied output losses are likely to have been substantially larger than those implied by the estimates of Blanchard and Galí.

We illustrate this point more clearly in Figure 5, where we plot the impulse responses implied by the “baseline” market-information-based measure and those implied by the two VAR-based measures, along with the corresponding 95 percent confidence intervals of the former.¹⁷ As it

¹⁷We use the “log-price change” as our “baseline” measure.

appears from the figure, the real GDP response implied by the market-information-based measure is significantly larger than the one implied by the traditional VAR-based shock measures. In fact, the point estimates of the latter lie outside the 95 percent confidence intervals for several months, in particular the output response implied by the asymmetric VAR-based measure. The responses of the CPI are more similar, which is not surprising, as part of the increase in oil prices is reflected by construction in the noncore component of the CPI.

Why are the output responses implied by the VAR-based measures so different from the responses implied by the market-information-based measures? One possible explanation is that the VAR identification strategy fails to separate the oil-price fluctuations driven by exogenous shocks from the endogenous fluctuations driven by other kinds of structural shocks. For instance, a productivity shock may lead to an economic expansion and, through a demand channel, to higher oil prices. Consequently, it will generate a positive correlation between real GDP growth and oil-price movement. Therefore, although a “pure” positive exogenous oil-price shock would lead to a substantial output decline, the VAR identification strategy may fail to separate these two kinds of shocks, thereby inducing a substantially weaker and statistically insignificant estimate of output response.

Our narrative approach provides an opportunity to examine directly this conjecture. In particular, if the above explanation is correct, then we should expect a much less negative, or possibly, even a positive, output response following an oil-price increase that is induced by gains in productivity or other kinds of endogenous oil-price increases responding to changes in oil demand. For this purpose, we construct an alternative series of oil-price shocks. We combine the shock series corresponding to event-types 18 and 19 (that is, events related to changes in oil demand due to economic development, improvement in oil usage efficiency, technology, etc.) and include the resulting shock variable as the *predetermined* input variable O_t in the system of equations (2). We plot the implied impulse responses in Figure 6, with the shock sizes calculated using both the “log-price change” and the “predicting error” approaches.

In both columns of the figure, the estimated impulse responses indicate that the narrative approach has correctly identified this kind of shock. Specifically, after a demand-driven shock, real GDP *increases* and the CPI level *declines* despite the fact that the noncore component of the CPI rises. The federal funds rate barely moves, as a potential increase driven by the higher output is partly offset by a decrease driven by the decline in the general price level. These are

exactly the qualitative responses that one would expect following a positive productivity shock. These results also confirm Kilian’s (2009) finding that an expansion in global aggregate demand leads to an increase in U.S. real GDP growth and, at the same time, in the price of oil.¹⁸

What distinguishes our findings from those of Kilian (2009) is that he has identified a global aggregate-demand shock, which drives *up* the CPI in the U.S., whereas our narrative reading has detected changes in oil demand likely reflecting productivity gains, which drives *down* the CPI in the U.S. Looking more closely at the days of our event-types also confirms this point. In fact, in our sample, such shocks occurred primarily after 2000, consistent with what many have suggested, that the oil-price increases of the past few years are the result of “an expanding world economy driven by gains in productivity” (*The Wall Street Journal*, August 11, 2006), which, among other things, were reflected in rapidly rising imports into the U.S. of inexpensive consumer goods from China and other emerging economies in the U.S.¹⁹ Over time, real GDP declines, as the adverse effect of higher oil price may eventually dominates the initial stimulating effect of the productivity shock. However, the CPI remains below its preshock level, as the price decline originating from the initial productivity shock still overweighs the inflationary pressures arising from higher oil prices even 18 months after the shock.

4.3 Other Kinds of Oil-Price Shocks Based on Market Information

Next, we examine the macroeconomic effects of other kinds of oil-price shocks identified through our narrative approach. These shocks reflect an important portion of the fluctuations observed in the global oil market during the past two and a half decades. Traditional VAR identification strategies are normally unable to separate these kinds of shocks from exogenous oil shocks.

OPEC and Non-OPEC Oil-Price “Shocks”

OPEC and non-OPEC oil exporters’ decisions to change their production plans may represent another source of endogenous oil-price fluctuations. We include these events in our event-types 13 and 14 (see Table 1). One could reasonably argue that these decisions reflect endogenous responses of oil producers to developments in the global oil market and, more generally, in

¹⁸As in Kilian (2009), the output response we obtain here is statistically insignificant, possibly reflecting the small number of observations corresponding to such shocks in our sample period.

¹⁹Two recent studies, Hamilton (2009b) and Kilian and Hicks (2009), also attribute the oil-price hikes during 2003-2008 to stronger demand led by an expansion in the world economy, in particular in China and India.

the global economy.^{20,21} As we have done with other event types, here we also construct a corresponding “shock” series and examine its effects on the U.S. economy.

As displayed in Figure 7, in response to these types of “shocks,” real GDP declines and the CPI rises. The oil price increases, with the peak response arriving three months after the “shock,” and the federal funds rate rises significantly in response to the substantial CPI increase amid only modest output decline. This pattern is different from what is induced by a typical productivity shock (Figure 6), or an exogenous oil-price shock (Figures 3A through 3C). In particular, the output loss implied by an OPEC/non-OPEC oil exporters’ “shock” is substantially weaker than the output loss induced by an exogenous oil-price shock, with the output response remaining statistically insignificant;²² however, output does not rise either as much as in Figure 6 when responding to a positive global aggregate demand shock. Accordingly, the monetary policy authority’s response also lies between its responses in those two cases, as it becomes restrictive rather than accommodative, as in Figures 3A through 3C when responding to exogenous oil shocks, but not as restrictive as in Figure 6 when responding to a demand-driven oil shock. This result is consistent with the finding in the literature that oil prices no longer Granger-cause real GDP and other macroeconomic variables after 1973, when OPEC became able to effectively influence the global oil market (see Hooker, 1996).

Oil-Market-Specific Demand Shocks

Event-types 15 through 17 reflect demand shocks that are specific to the oil market, i.e., changes in oil demand unrelated to changes in the level of global economic activity. For example, changes in the U.S. Strategic Petroleum Reserve (SPR) in the past two decades have always been associated with substantial oil price changes.²³ Changes in commercial oil and gas inventories,

²⁰See Barsky and Kilian (2002, 2004) for a detailed analysis.

²¹Nakov and Pescatori (2009) develop a model in which the global supply of oil is determined by the optimal production decisions of a dominant supplier and a fringe of competitive producers, representing, respectively, OPEC and non-OPEC oil exporters. One implication of their framework is that the price of oil is ultimately a function of OPEC’s share in the global oil market.

²²“OPEC developments” is the most frequent event type in our sample, with a sample frequency of 742 trading days. Therefore, the lack of statistical significance is unlikely due to the small number of corresponding observations.

²³For example, on January 23, 2007, the U.S. Department of Energy announced that the U.S. would double the size of its Strategic Petroleum Reserves (SPR) over the next two decades. The crude oil price responded with a 7.6 percent rally on the same day. On September 13, 2000, after the White House mentioned that it was considering a release of SPR in response to oil price increases, the price of oil dipped 1.5 percent. It dropped

including both realized changes and market expectations of future changes, have also significantly affected the oil market. These oil-market-specific demand changes often reflect changes in the precautionary demand for oil, as discussed in Kilian (2009).

An oil-market-specific demand shock drives up the real price of oil, by definition. The price increase, however, is not as persistent as the one induced by the exogenous oil-price shocks, as the real price of oil returns to its preshock level in about five months (Figure 8). In response to the shock, real GDP declines, although the decline is less persistent and quantitatively much weaker than the decline implied by an exogenous oil shock: Real GDP returns to its preshock level in about six months, and the 24-month cumulative output loss is 2.2 percent of monthly GDP, compared with a 6.8 percent cumulative output loss implied by exogenous oil shocks. Moreover, the decline in output is statistically significant at the 95 percent level only for the first three months after the shock. The CPI increases immediately after the shock, but then returns to its preshock level in four months. Because of the short-lived output and CPI responses, monetary policy accommodation is quite limited. The impulse responses induced by the oil-market-specific demand shocks identified by our narrative approach are less persistent than those induced by the precautionary demand shocks identified by Kilian (2009). One potential explanation for this difference is that, in his work, the precautionary demand shocks are identified as any real oil-price movement that cannot be explained by his measures of changes in global real economic activities or oil production, and thus may include some of the event types that are included in our exogenous shock measures, such as political development or military actions.

Oil-Price Shocks Related to Military Actions in The Middle East

Finally, we examine the effects of oil-price shocks related to military actions in the Middle East. Not surprisingly, these actions tend to drive up the real price of oil and lead to a substantial output decline (Figure 9). In particular, real GDP gradually decreases, with the largest decline arriving 11 months after the shock. The decline is statistically significant at the 95 percent level for most horizons, and the cumulative output loss over 24 months reaches 8.4 percent of monthly GDP, even larger than the 6.8 percent output loss implied by the “baseline” exogenous shock. Moreover, if one argues that U.S. military spending increases in response to these military actions and that such increases stimulate U.S. GDP, then the actual GDP decline induced by these military shocks through the oil-price channel could be even larger than the decline

another 3 percent three weeks later, on October 5, 2000, the day the SPR release was officially confirmed.

implied by our estimates. Therefore, our estimates indicate a strongly adverse effect on the U.S. economy of oil-price shocks induced by military conflicts in the Middle East, which tend to raise the market’s concerns over oil supply disruptions as well as future oil availability. The CPI also increases after the shock, although the increase is smaller and less significant than when responding to exogenous oil shocks. Monetary policy becomes even more accommodative than in the “baseline” exogenous oil shock case, with the federal funds rate decreasing by a cumulative 5.2 percentage points over 24 months after the shock, or 22 basis points lower than the preshock level on average, significantly larger than the 14-basis point decline when responding to an exogenous oil-price shock.

4.4 Robustness and Stability

We perform a number of robustness and stability checks to verify the validity of our main conclusions. In particular, we consider two subsample periods—January 1984 to December 1994 and January 1995 to October 2007—to check the consistency of the estimated responses across different periods.²⁴

Table 2 summarizes the estimation results from our robustness check. As above, to facilitate the comparison of the effects of different oil-price shock measures, the shock sizes are normalized so that the maximal response of real oil price is 10 percent. Under the heading “output,” we report the 24-month cumulative output loss, which is the sum of impulse response coefficients for output. Under the heading “price,” we report the sum of impulse response coefficients over 24 months for the CPI, divided by 24, which can be interpreted as the average increment in the CPI over the two years after the shock. Finally, under the heading “interest rate,” we report the sum of impulse response coefficients for the federal funds rate, divided by 24, which measures the average monetary policy response during the two years after the shock.

When the 24-year sample period is split into two subsample periods (1984-1994, 1995-2007), the overall response pattern changes significantly. In particular, the responses implied by the two VAR-based measures differ greatly across the two subsample periods, with the directions of the real GDP and CPI responses bearing exactly the opposite signs. For instance, for the first subsample period, the estimates implied by the asymmetric VAR-based measure indicate

²⁴In an earlier version of our work, we have also obtained results for alternative measures of price level (e.g., the personal consumption expenditure deflator), as well as for a univariate, distributed lag model. None of those experiments have had important effects on the robustness of our conclusions and are thus omitted here.

that following an oil-price shock, real GDP declines, with the 24-month cumulative output loss reaching 4.2 percent. In stark contrast, for the second subsample period, the corresponding estimates indicate that real GDP increases, with a cumulative output gain of 3.6 percent in the 24 months after the shock (Figure 10). With regard to the symmetric VAR-based measure, in the first subsample period, real GDP barely declines after the shock, with the 24-month cumulative output gaining 0.5 percent, consistent with the earlier literature's findings that the symmetric oil-price measure cannot generate a positive output response following the oil-price declines in the mid-1980s (Hamilton, 1996). The estimates from the second subsample, however, are more consistent with the conventional wisdom. Real GDP declines after the oil-price shock, with a cumulative output loss of 2.7 percent over the 24 months after the shock. However, estimates using the market-information-based measures suggest that the output response has not changed substantially across these two subsample periods, with real GDP declining significantly after the exogenous oil-price shocks in both periods. The positive output response implied by the asymmetric VAR-based measure for the post-1995 subsample may simply reflect a predominance of demand-driven oil-price shocks since the late 1990s, likely originating from a global economic expansion (see Kilian, 2009), rather than genuine exogenous oil-price shocks.

5 Concluding Remarks

This paper combines narrative and quantitative approaches to examine the dynamic effects of oil-price shocks on the U.S. economy. To correctly identify exogenous oil shocks, we first collect oil-market related information from a number of oil-industry trade journals and government publications, and compile a database identifying all major events that have significantly affected the global oil market on a daily basis since 1984. Based on such information, we are able to isolate events that are exogenous to the U.S. economy and construct corresponding measures of exogenous oil-price shocks. Furthermore, shock magnitudes are calculated by running an oil-price forecasting model incorporating oil futures prices. These procedures help alleviate the endogeneity and predictability problems that have pestered the traditional VAR identification strategies in the literature.

One contribution of our work is the thorough examination of all kinds of oil-related events in the past two and a half decades, more comprehensive than just focusing on geopolitical or military events, as most of the earlier literature has done so far. Moreover, in constructing the database, we have preserved as much primitive information on the oil-market developments as

possible, with the hope of facilitating possible future studies by other researchers on the nature and implications of these events.

After deriving our measures of various kinds of oil shocks, we go on to examine their dynamic macroeconomic effects. We find that exogenous oil-price shocks have had substantial and statistically significant impacts on the U.S. economy during the past two and a half decades. In contrast, traditional VAR identification strategies imply a substantially weaker and insignificant real effect for the same period. Further analysis reveals that this discrepancy stems from the inability of VAR-based approaches to separate exogenous oil-supply shocks from endogenous oil-price fluctuations driven by changes in oil demand. Notably, our study also suggests that the U.S. economy may not have become as insulated from external oil shocks during the last 25 years as earlier studies have suggested. To fully examine how the oil price-macroeconomy relationship has evolved during the whole postwar period, a thorough study along the same narrative and quantitative approach for the period prior to the “Great Moderation” is called for. This will be the topic for future research.

A Data Description

This appendix describes the data series used in our paper.

Output — Real gross domestic product (billions of chained 2000 dollars), Bureau of Economic Analysis, National Income and Products Accounts, Table 1.1.6, Line 1;

Industrial production — Industrial production, total index (2002=100), Federal Reserve Board, statistical release G.17, Haver Analytics mnemonic: IP@USECON.

Capacity utilization — Capacity utilization, total industry (percent of capacity), Federal Reserve Board, statistical release G.17, Haver Analytics mnemonic: CUT@USECON.

Federal funds rate — Federal funds effective rate (percent per annum), Federal Reserve Board, statistical release H.15, Haver Analytics mnemonic: FFED@USECON.

Headline CPI — Consumer price index, all urban consumers, U.S. city average, all items (1982-84=100), Bureau of Labor Statistics, series ID: CUUR0000SA0;

Crude Petroleum PPI: Producer price index - Crude petroleum (domestic production, 1982=100), Bureau of Labor Statistics, Series ID: WPU0561;

Oil prices — We use daily spot and futures market prices (dollars per barrel) at the New York Mercantile Exchange (NYMEX) of West Texas Intermediate (WTI) light sweet crude oil for delivery at Cushing, Oklahoma.

Spot price — spot market price, Wall Street Journal, Haver Analytics mnemonic: PZTEXA@daily.

One-month futures price — First-expiring contract settlement (Contract 1, near month), Wall Street Journal and Department of Energy, Haver Analytics mnemonic: PZTEXF1@daily.

Two-month futures price — 2-month Contract Settlement (Contract 2), Department of Energy, Haver Analytics mnemonic: PZTEXF2@daily.

Three-month futures price — 3-month Contract Settlement (Contract 3), Wall Street Journal and Department of Energy, Haver Analytics mnemonic: PZTEXF3@daily.

Four-month futures price — 4-month Contract Settlement (Contract 4), Department of Energy, Haver Analytics mnemonic: PZTEXF4@daily.

Six-month futures price — 6-month Contract Settlement, Wall Street Journal, Haver Analytics mnemonic: PZTEXF6@daily.

References

- [1] Alexopoulos, Michelle, 2008. "Read All About It! What Happens Following a Technology Shocks," working paper, University of Toronto.
- [2] Alexopoulos, Michelle, and Jon Cohen, 2009. "Measuring Our Ignorance, One Book at a Time: New Indicators of Technological Change, 1909–1949," *Journal of Monetary Economics* 56, 450-470.
- [3] Alquist, Ron and Lutz Kilian, 2009. "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, forthcoming.
- [4] Barsky, Robert B., Lutz Kilian, 2002. "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative." In: Bernanke, Ben S. and Kenneth Rogoff (Eds.), *NBER Macroeconomics Annual 2001*. MIT Press, Cambridge, 137-183.
- [5] Barsky, Robert B., Lutz Kilian, 2004. "Oil and the Macroeconomy since the 1970s," *Journal of Economic Perspectives* 18, 115-134.
- [6] Bernanke, Ben S., Mark Gertler, and Mark Watson, 1997. "Systematic Monetary Policy and the Effects of Oil Price Shocks." *Brookings Papers on Economic Activity*, 1997:1, 91-157.
- [7] Blanchard Olivier J. and Jordi Galí, 2009. "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so Different from the 1970s?" In: *International Dimensions of Monetary Policy*, Jordi Galí and Mark Gertler (Eds.). University of Chicago Press: Chicago, IL, forthcoming.
- [8] Burnside, Craig, Martin Eichenbaum, and Jonas D.M. Fisher, 2004. "Fiscal Shocks and Their Consequences," *Journal of Economic Theory* 115, 89-117.
- [9] Chinn, Menzie D., Michael LeBlanc, and Olivier Coibon, 2005. "The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline," NBER Working Paper No. 11033.
- [10] Chow, Gregory C. and An-loh Lin, 1971. "Best Linear Unbiased Interpolation, Distribution and Extrapolation of Time Series by Related Series," *Review of Economics and Statistics* 53, 372-375.
- [11] Christiano, Larry J., Martin Eichenbaum, and Charles L. Evans, 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In: Taylor, John B. and Michael Woodford (Eds.), *Handbook of Macroeconomics*, Vol. 1A. North-Holland, Amsterdam, 65-148.
- [12] Cochrane, John H. and Monika Piazzesi, 2002. "The Fed and Interest Rates: A High-frequency Identification," *American Economic Review* 92, 90-95.
- [13] Cochrane, John H. and Monika Piazzesi, 2005. "Bond Risk Premia," *American Economic Review* 95, 138-160.
- [14] Hamilton, James D., 1983. "Oil and the Macroeconomy Since World War II," *Journal of Political Economy* 91, 228-248.

- [15] Hamilton, James D., 1985. "Historical Causes of Postwar Oil Shocks and Recessions," *Energy Journal* 6, 97-116.
- [16] Hamilton, James D., 1996. "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics* 38, 215-220.
- [17] Hamilton, James D., 2003. "What is an Oil Shock?" *Journal of Econometrics* 113, 363-398.
- [18] Hamilton, James D., 2009. "Understanding Crude Oil Prices," *Energy Journal* 30, 179-206.
- [19] Hamilton, James D., 2009. "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, Spring 2009.
- [20] Hooker, Mark A., 1996. "What Happened to the Oil Price-Macroeconomy Relationship?" *Journal of Monetary Economics* 38, 195-213.
- [21] Hoover, Kevin D. and Stephen J. Perez, 1994. "Post Hoc Ergo Propter Hoc Once More: An Evaluation of 'Does Monetary Policy Matter?' in the Spirit of James Tobin," *Journal of Monetary Economics* 34, 89-99.
- [22] Kassarijian, Harold H., 1977. "Content Analysis in Consumer Research," *Journal of Consumer Research* 4, 8-18.
- [23] Kilian, Lutz, 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much do They Matter for the U.S. Economy?" *Review of Economics and Statistics* 90, 216-240.
- [24] Kilian, Lutz, 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053-1069.
- [25] Kilian, Lutz and Bruce Hicks, 2009. "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?" working paper, University of Michigan.
- [26] Kilian, Lutz and Robert J. Vigfusson, 2009. Pitfalls in Estimating Asymmetric Effects of Energy Price Shocks. International Finance Discussion Paper no. 2009-970, Federal Reserve Board.
- [27] Levy, Daniel, Shantanu Dutta, and Mark Bergen, 2002. "Heterogeneity in Price Rigidity: Evidence from a Case Study Using Microlevel Data," *Journal of Money, Credit, and Banking* 34, 197-220.
- [28] Lütkepohl, Helmut, 2005. *New Introduction to Multiple Time Series Analysis*. Springer-Verlag: Berlin, Germany.
- [29] Nakov, Anton and Andrea Pescatori, 2009. "Monetary Policy Tradeoffs with a Dominant Oil Producer," *Journal of Money, Credit, and Banking*, forthcoming.
- [30] Ramey, Valerie A., 2009. "Identifying Government Spending Shocks: It's All in the Timing," NBER Working Paper No. 15464.
- [31] Ramey, Valerie A. and Matthew D. Shapiro, 1998. "Costly Capital Reallocation and the Effects of Government Spending," *Carnegie-Rochester Conference Series on Public Policy* 48, 145-194.

- [32] Romer, Christina D. and David H. Romer, 2004. "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review* 94, 1055-1084.
- [33] Romer, Christina D. and David H. Romer, 2009. "The Macroeconomic Effect of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, forthcoming.
- [34] Wu, Tao and Andrew McCallum, 2005. "Do Oil Futures Prices Help Predict Future Oil Prices?" *Economic Letter* no. 2005-38, Federal Reserve Bank of San Francisco.



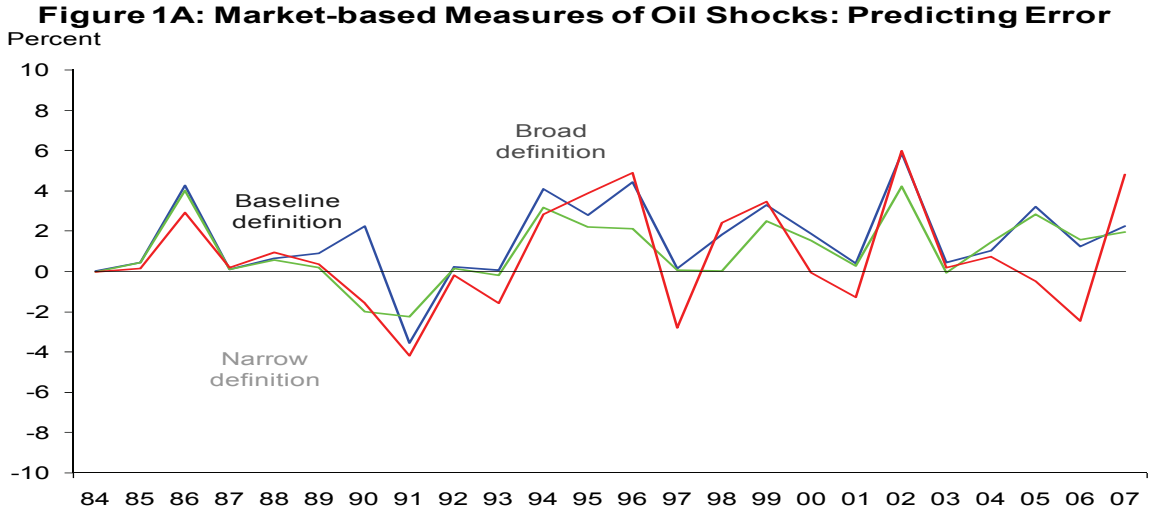
Table 1: Classification of oil-related events

CODE	EVENT-TYPES	NUMBER OF DAYS	RELATIVE FREQUENCY
1.	U.S. weather changes	127	2.12%
2.	Non-U.S. weather changes	6	0.10%
3.	Oil production/transportation disruptions in the U.S., e.g., refinery explosions, etc.	248	4.15%
4.	Oil production/transportation disruptions not in the U.S., natural disasters, etc.	238	3.98%
5.	New oil field discoveries in the U.S.	0	0%
6.	New oil field discoveries outside the U.S.	0	0%
7.	Political developments in the Middle East	476	7.98%
8.	Political developments in non-Middle East oil-exporting countries	47	0.78%
9.	Political developments in other regions	0	0%
10.	Military actions in the Middle East	110	1.85%
11.	Military actions in non-Middle East oil-exporting countries	2	0.03%
12.	Military actions in other regions	0	0%
13.	OPEC developments on oil production plan, e.g., proposals to cut back oil production, etc.	741	12.41%
14.	Non-OPEC oil exporters' developments on oil production plan	121	2.03%
15.	U.S. oil and gas inventory announcements, SPR (strategic petroleum reserve) changes, etc.	730	12.23%
16.	Non-US oil inventory announcements, etc.	53	0.89%
17.	Changes in expectations of U.S. oil and gas inventories	206	3.44%
18.	Changes in U.S. oil demand originating from economic development, improvements in oil usage efficiency, technology, etc.	31	0.51%
19.	Changes in non-U.S. oil demand from economic development, improvements in oil usage efficiency, technology, etc.	30	0.51%
20.	Other oil product price movements (gasoline, heating oil, etc)	118	1.98
21.	Technical reasons, speculations on the oil market, e.g., short covering for certain contracts, etc.	826	13.83%
22.	No particular reason	941	15.75%
23.	Information not available	921	15.42%
TOTAL		5,971	100%

Table 2: Robustness and stability check

	Market-information based			Symmetric VAR-based			Asymmetric VAR-based		
	Output	Price	Interest rate	Output	Price	Interest rate	Output	Price	Interest rate
Baseline estimates	-6.75	0.14	-0.11	-2.91	0.05	-0.07	-1.66	0.04	-0.04
Subsamples									
1984-1994	-3.33	-0.11	-0.10	0.48	-0.04	0.01	-4.22	-0.24	-0.21
1995-2007	-2.44	0.14	0.17	-2.68	0.07	-0.01	3.61	0.10	0.34

Note: Table 2 reports the 24-month cumulative effects of different oil-price shock measures. For comparison, the shock magnitudes are normalized so that the largest response of the real oil price is 10 percent. The columns under the heading “Output” display the 24-month cumulative output loss; The columns under the “Price” and “Interest rate” headings display the 24-month cumulative changes in the CPI and the federal funds rate divided by 24, which capture, respectively, the average increment in the CPI and the average monetary policy response over the two years after the shock.



Note: Monthly shock series averaged to annual frequency.

Figure 2A: Alternative Measures of Exogenous Oil Shocks: Predicting Error




Figure 2B: Alternative Measures of Exogenous Oil Shocks: Log-Price Change




Figure 3A: Baseline Definition of Exogeneity

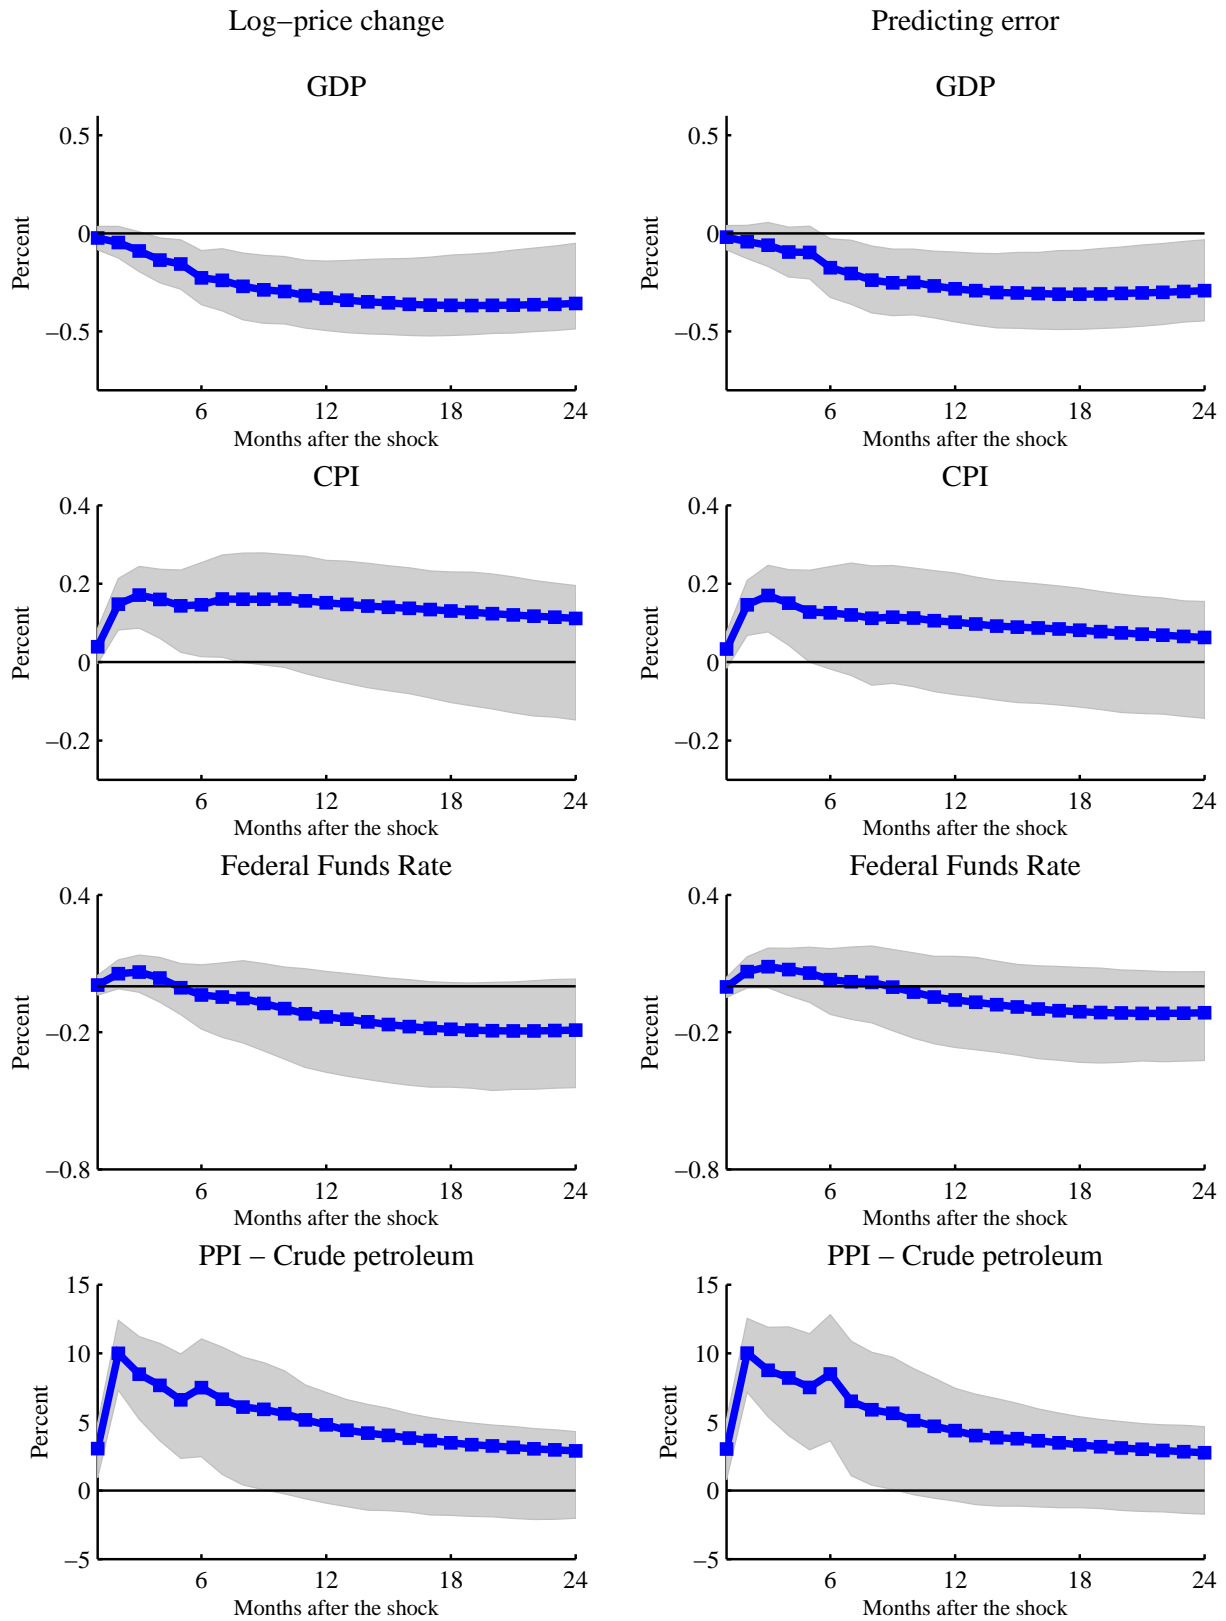


Figure 3B: Broad Definition of Exogeneity

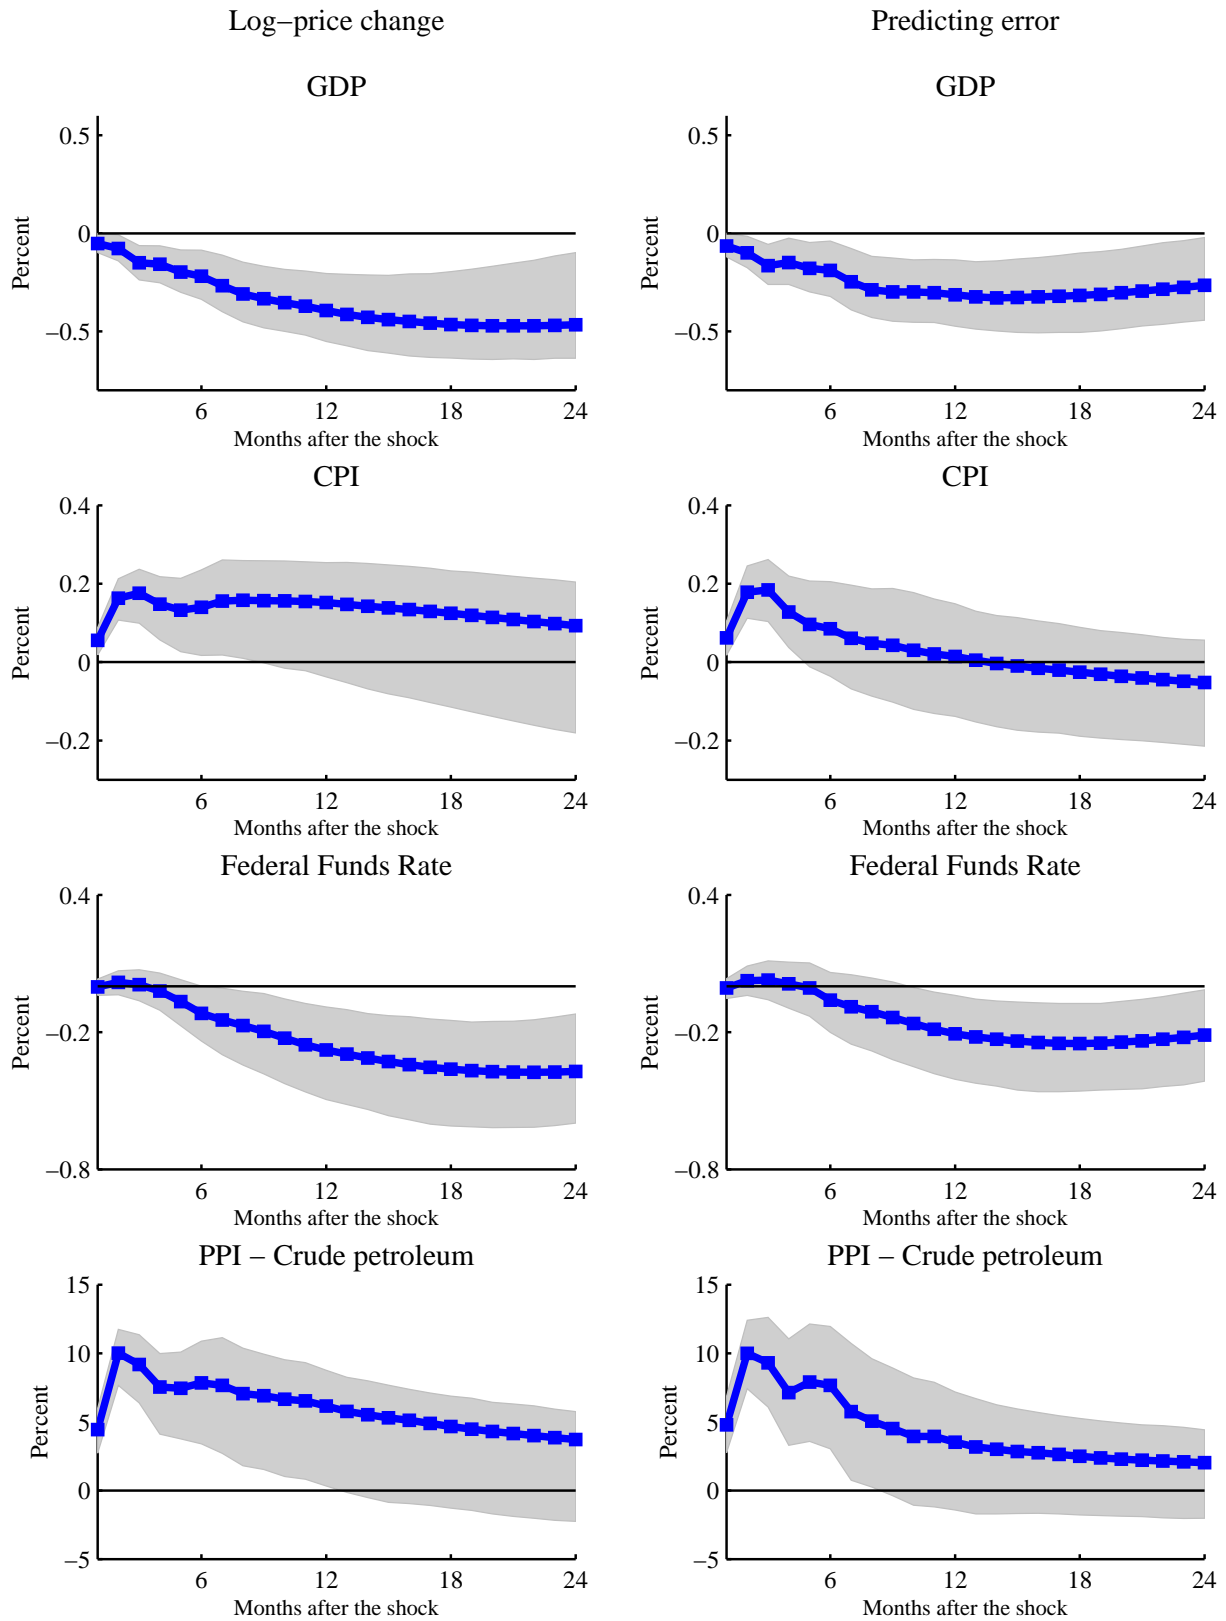


Figure 3C: Narrow Definition of Exogeneity

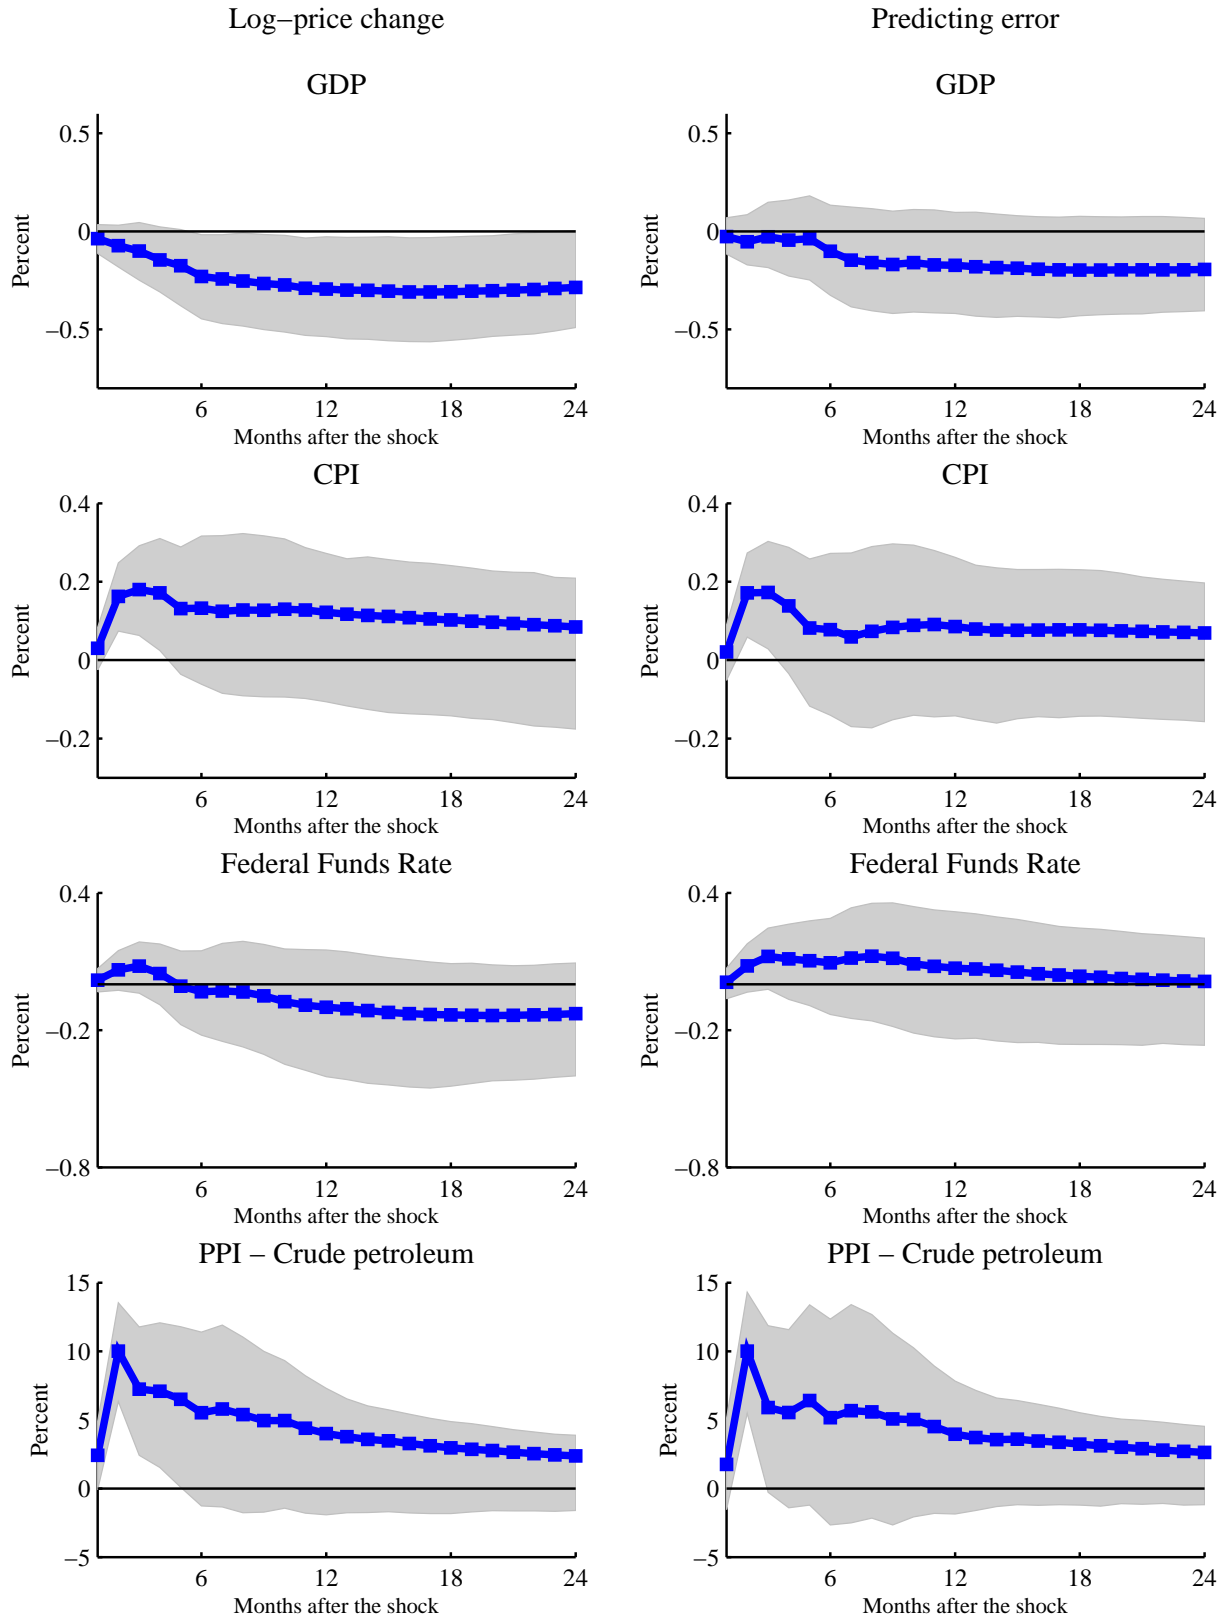


Figure 4: VAR–Based Measures

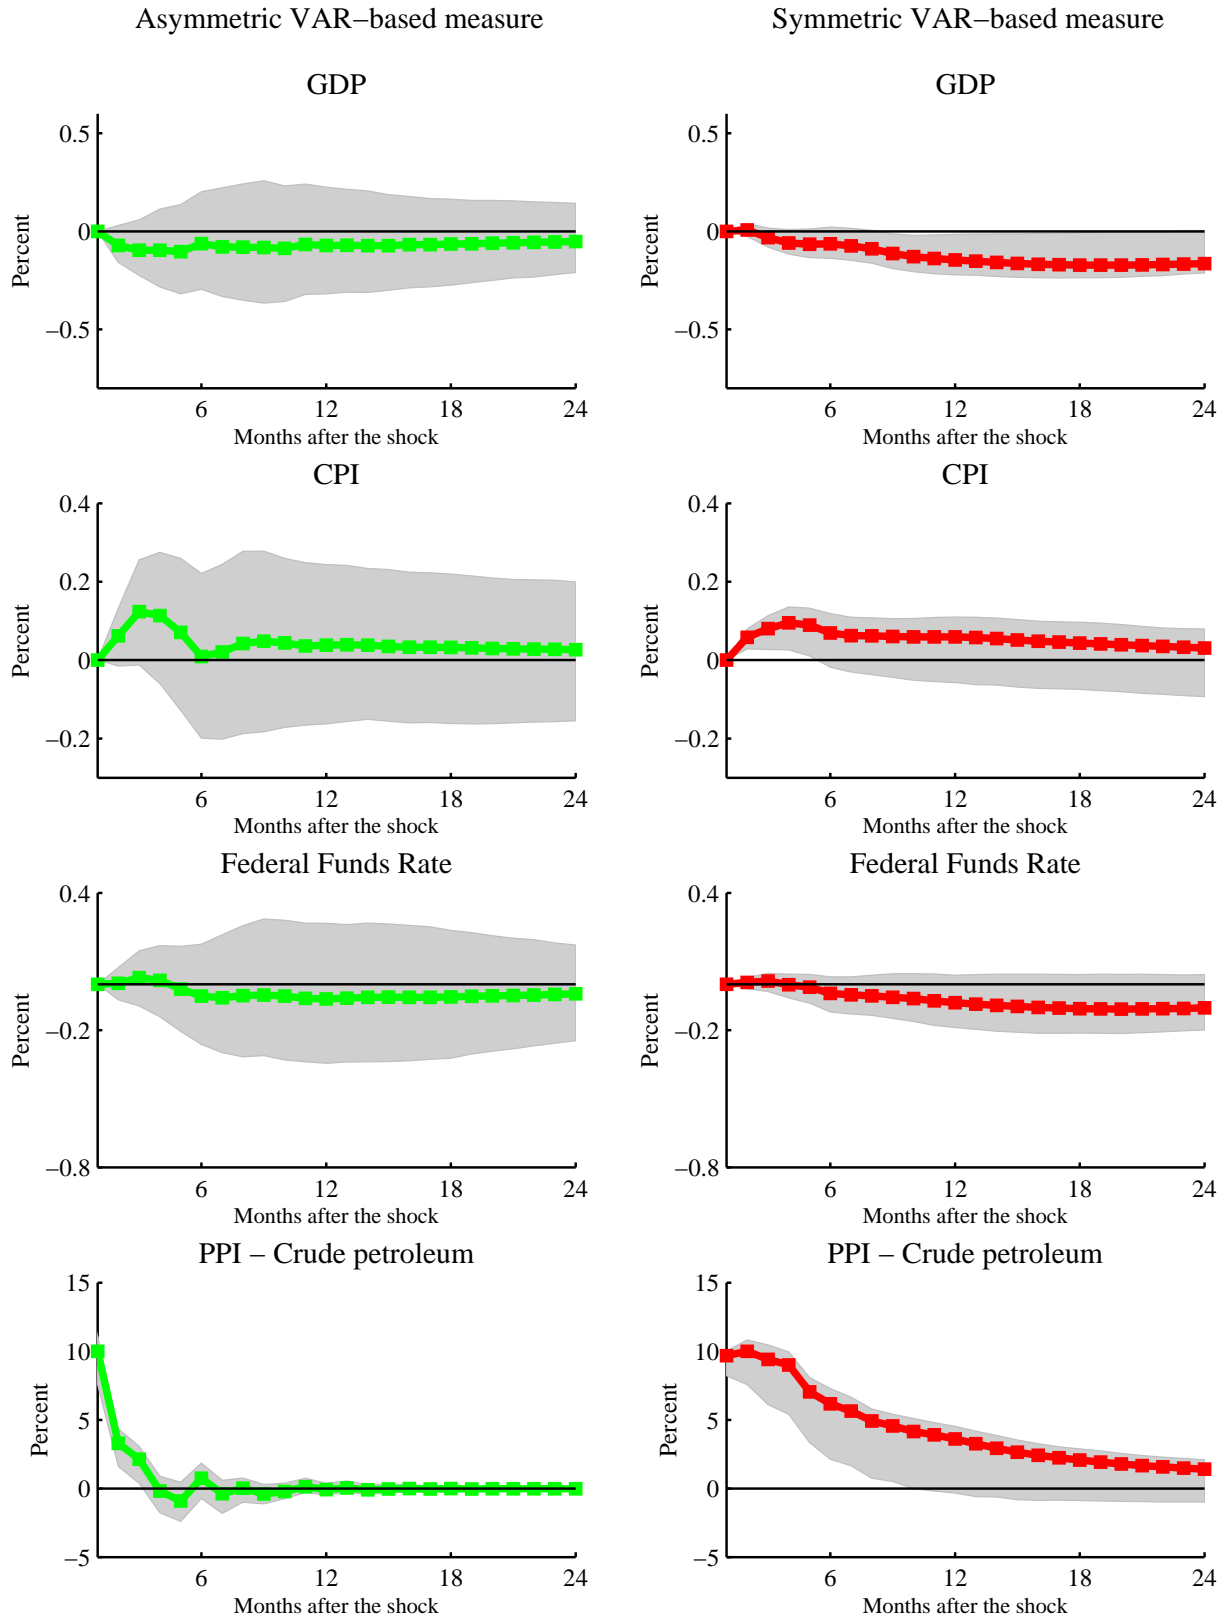


Figure 5: Alternative Estimates of Impulse Responses
 Log-price change Predicting error

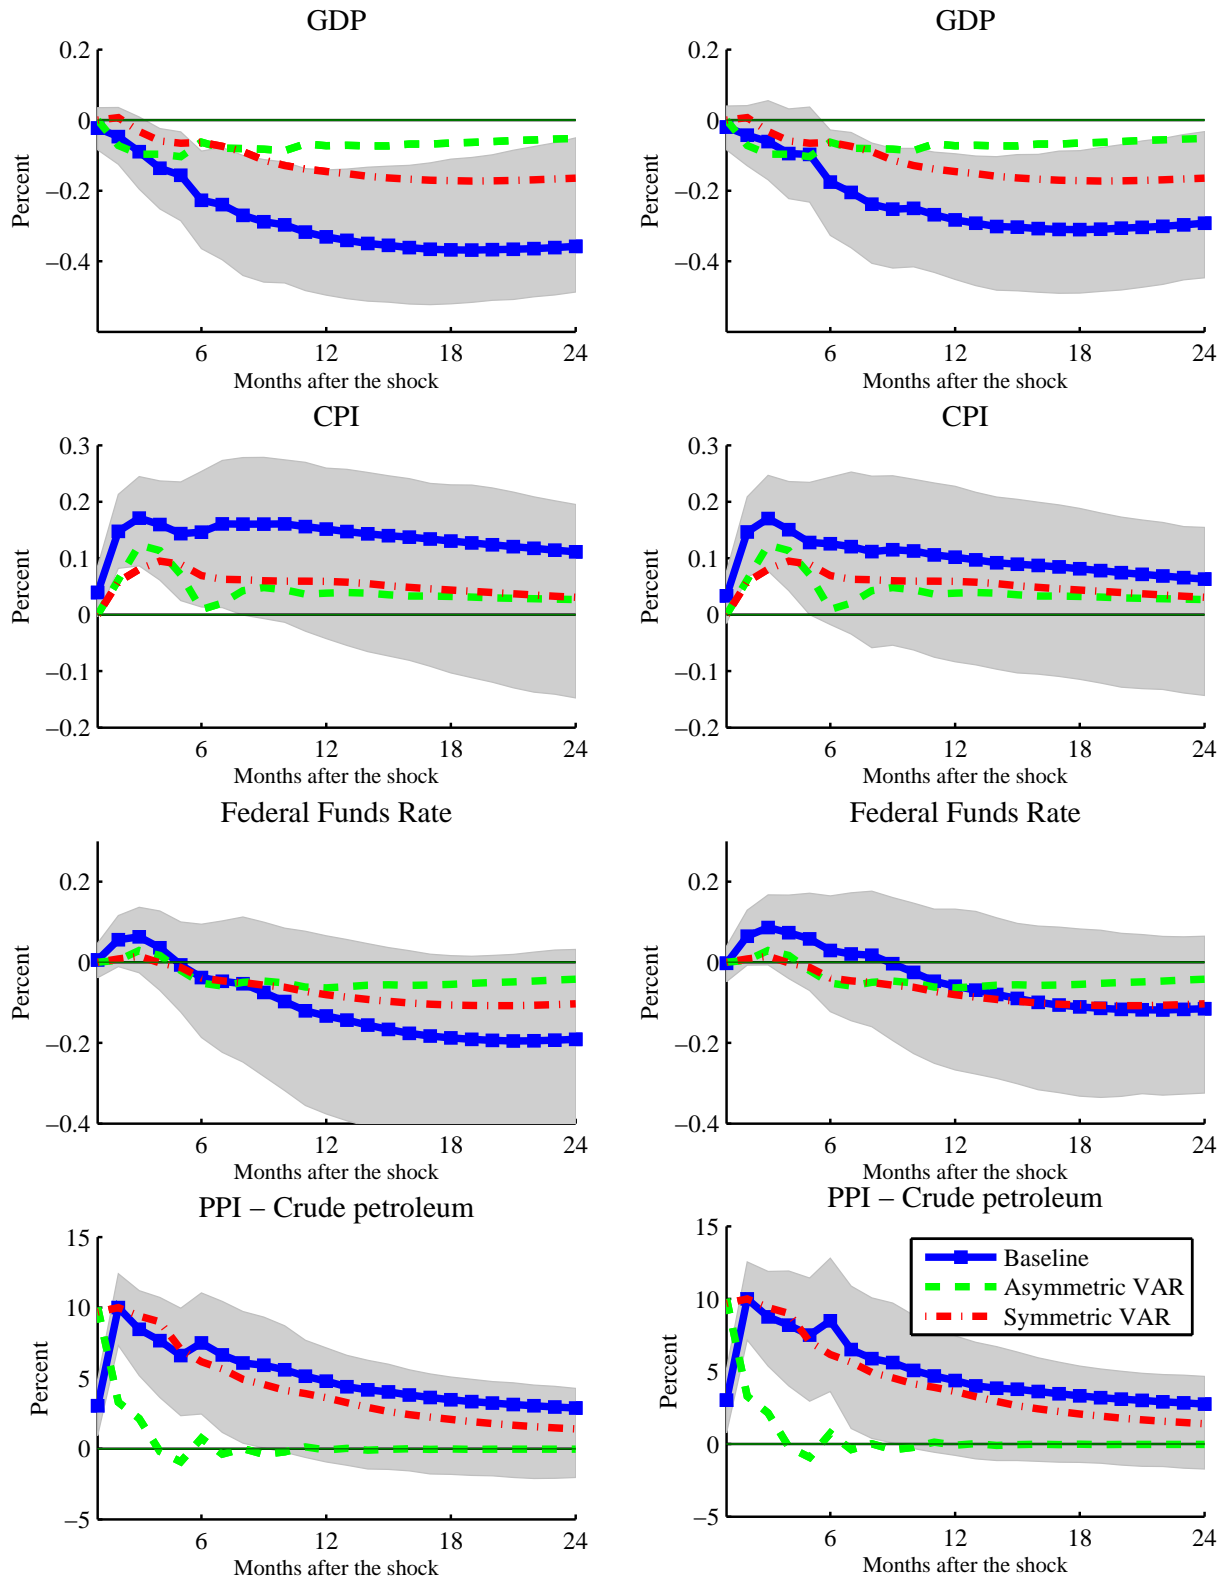


Figure 6: Impulse Responses to Oil Demand Changes

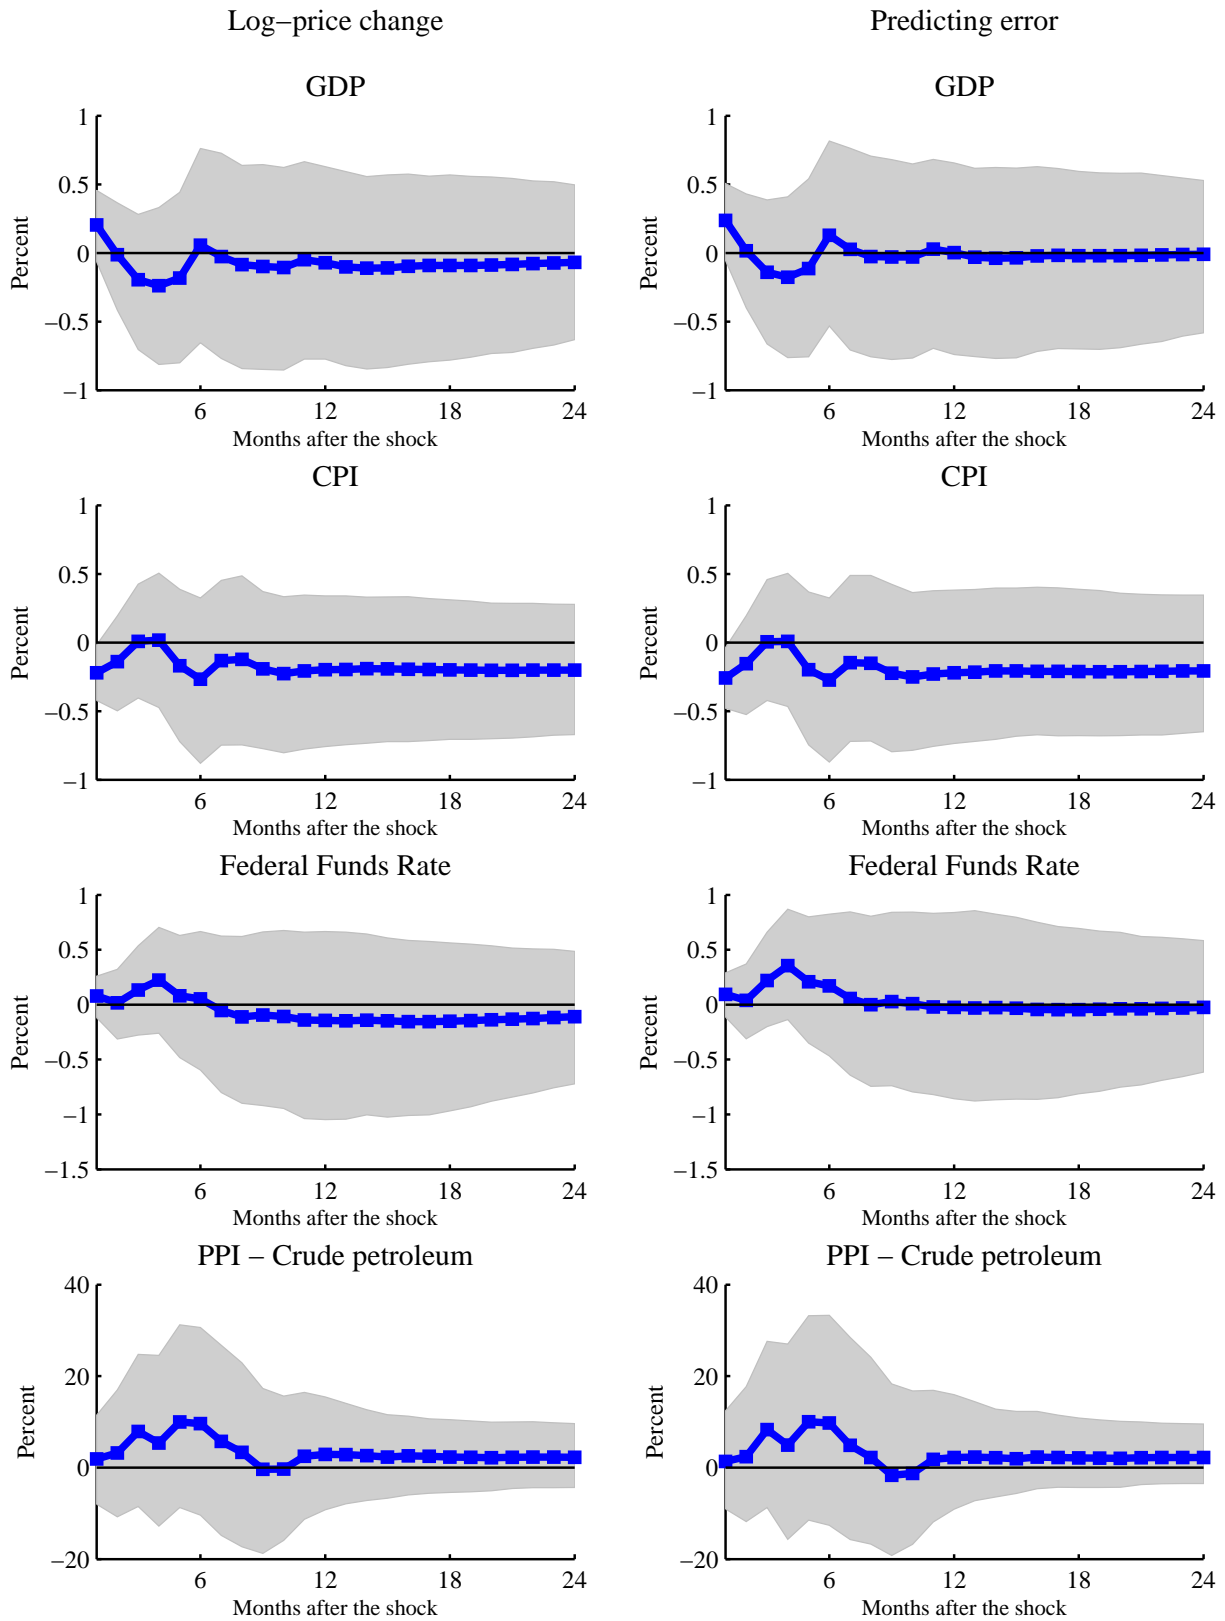


Figure 7: Impulse Responses to OPEC/Non-OPEC Moves

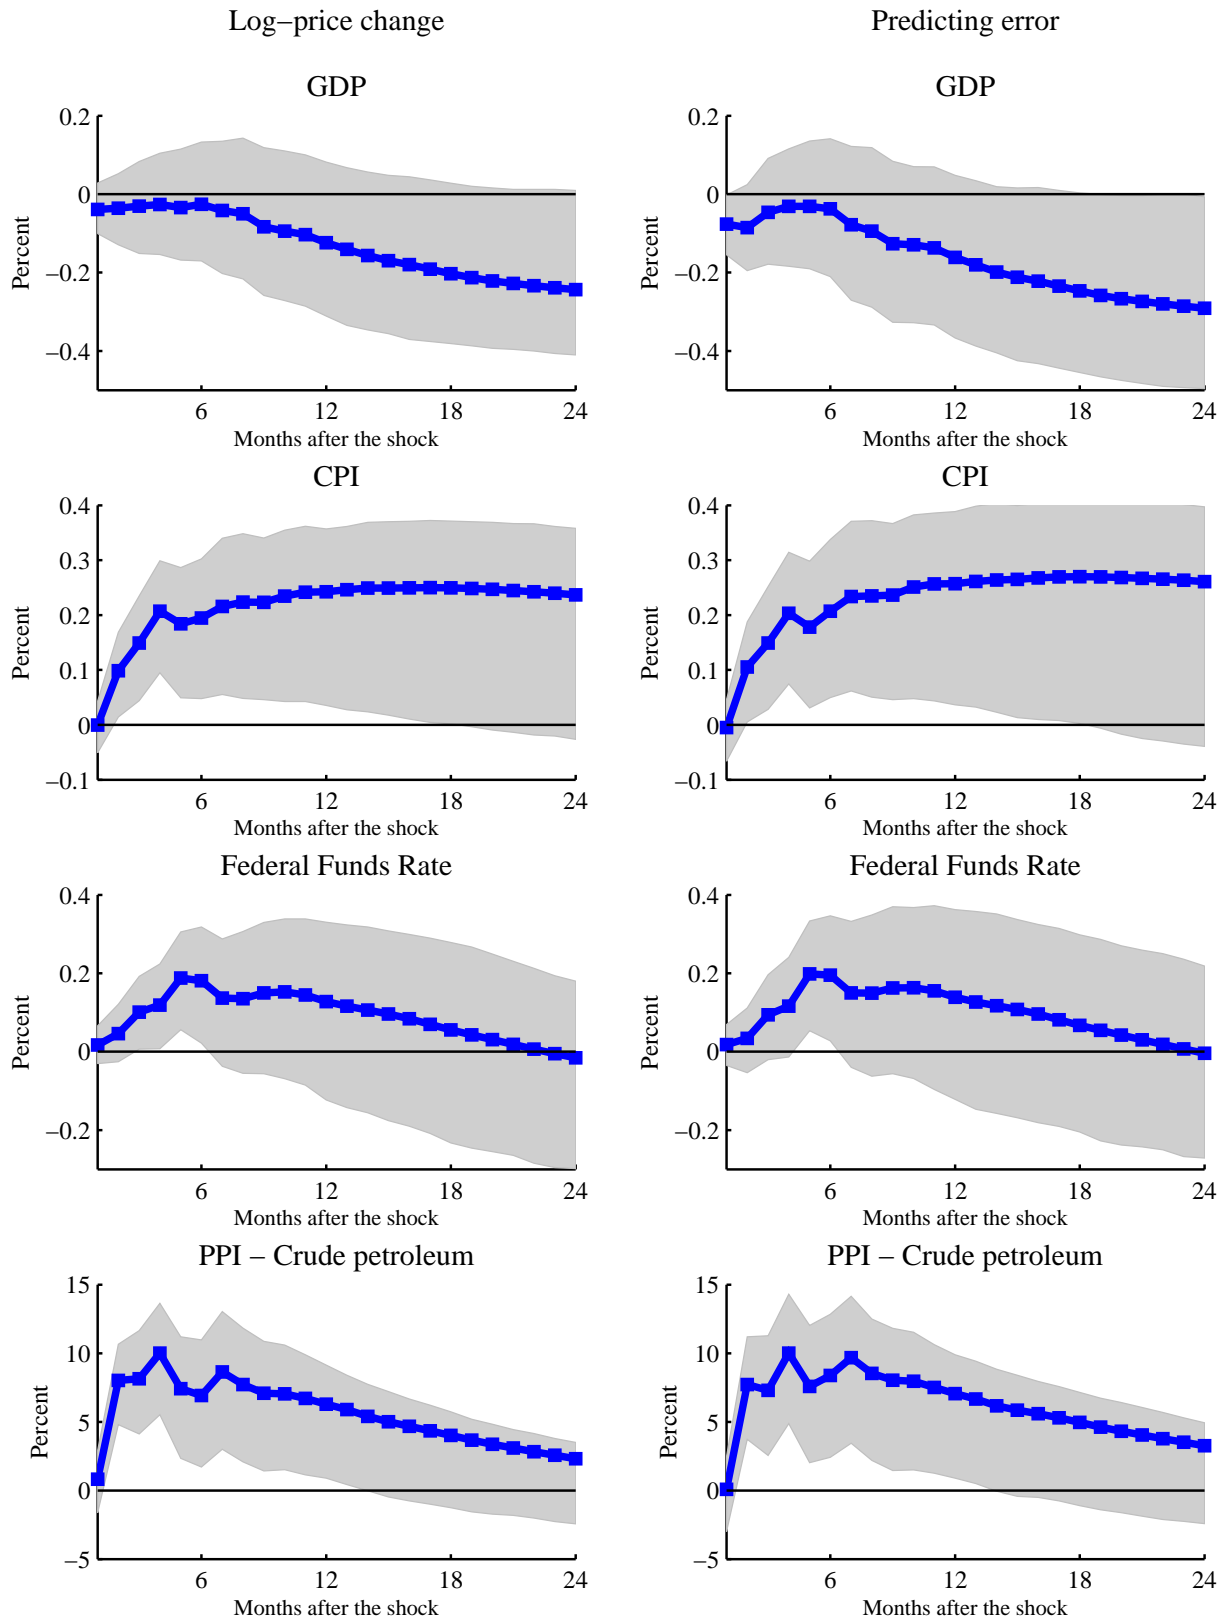


Figure 8: Impulse Responses to Oil Inventory Changes

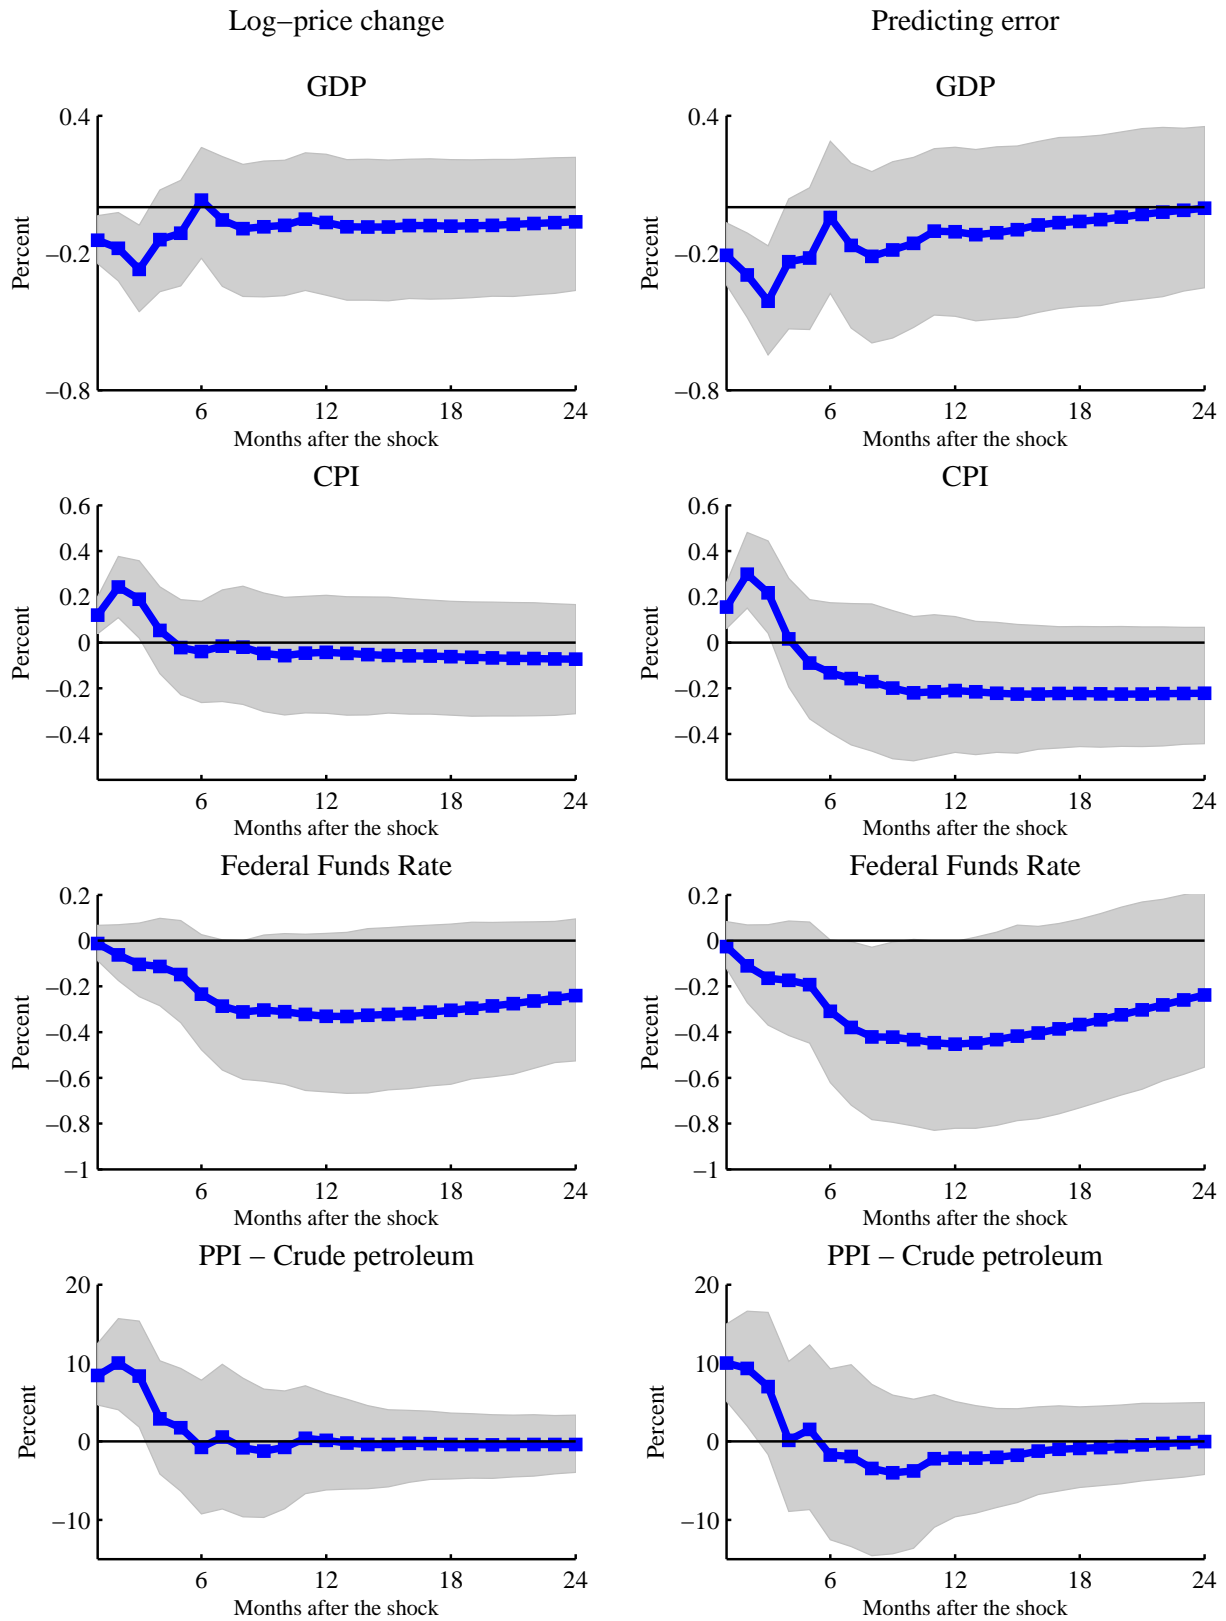


Figure 9: Impulse Responses to Military Conflicts

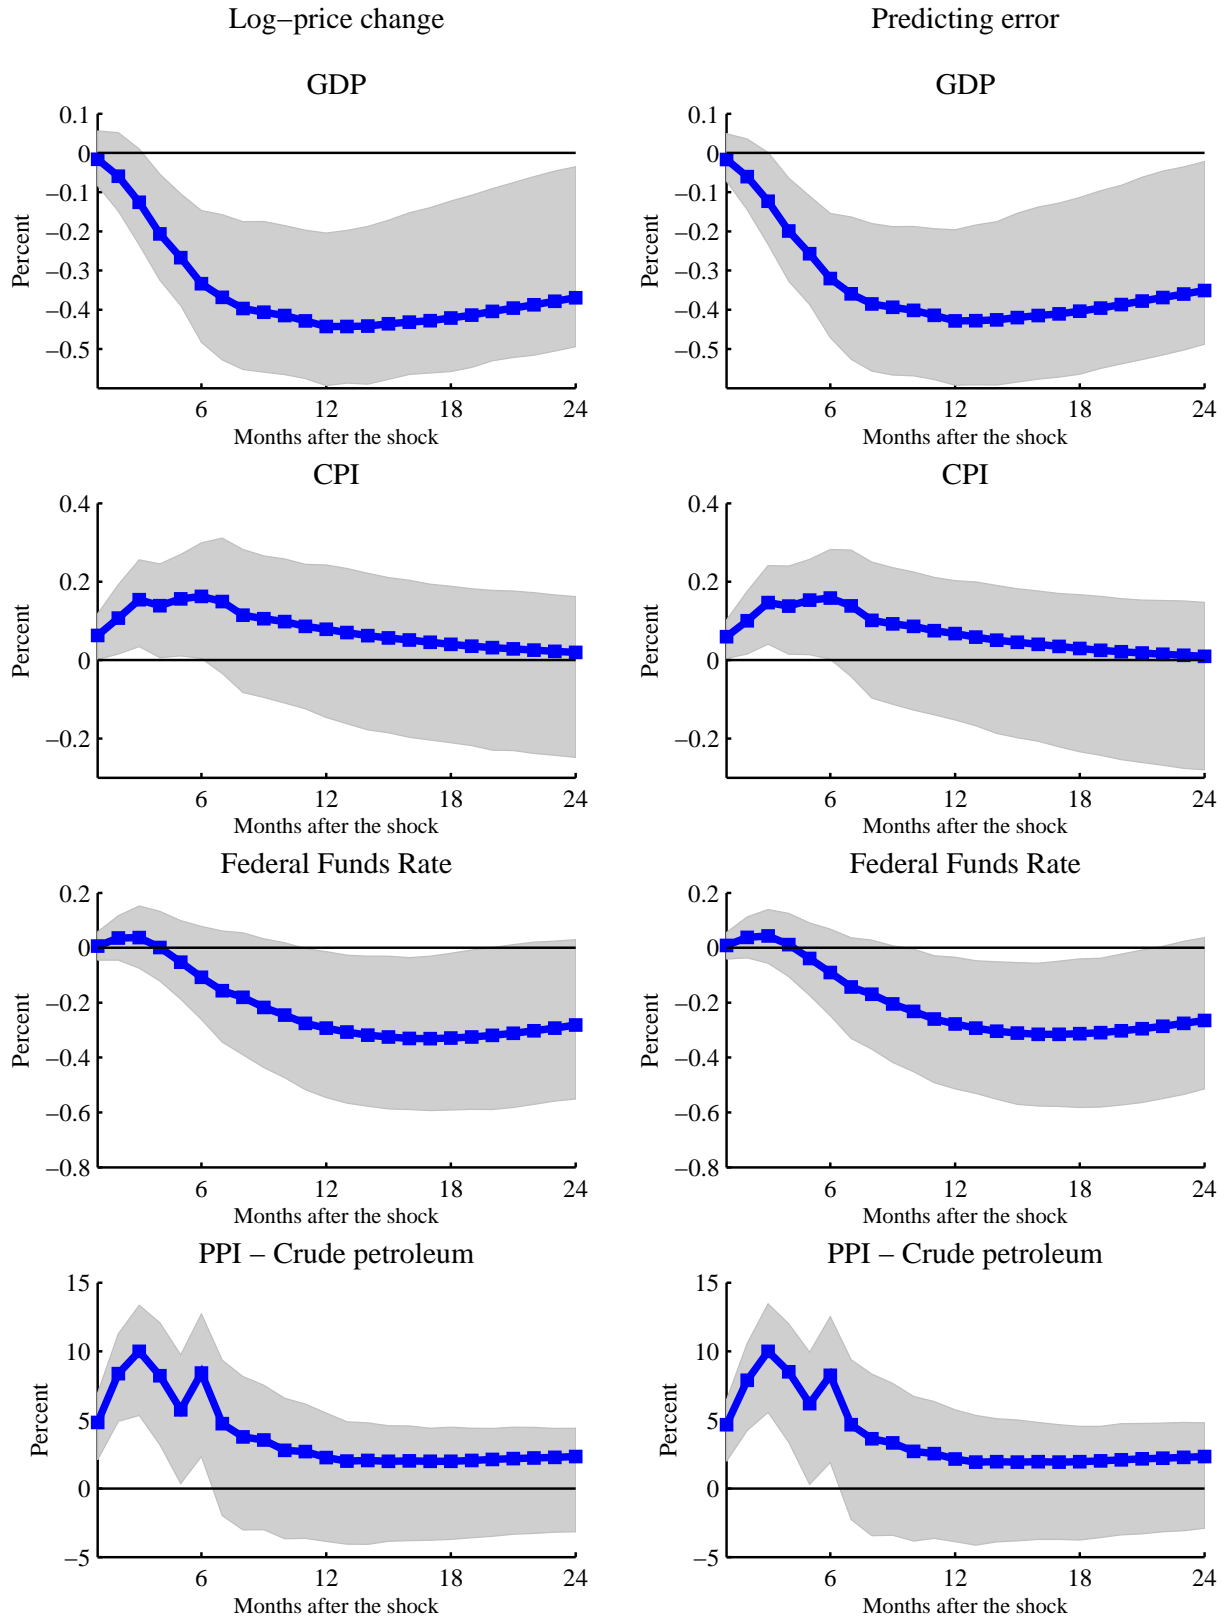


Figure 10: Impulse Responses in Sub-sample Periods

