ME
7/7/07
1:55 PM

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Measuring oil-price shocks using market-based information

Michele Cavallo
Federal Reserve Bank of San Francisco
Michele.Cavallo@sf.frb.org

Tao Wu
Federal Reserve Bank of Dallas
Tao.Wu@dal.frb.org

September 2006
This version: July 2007

ABSTRACT

We develop two measures of exogenous oil-price shocks for the period 1984 to 2006 based on market commentaries on daily oil-price fluctuations. Our measures are based on exogenous events that trigger substantial fluctuations in spot oil prices and are constructed to be free of endogenous and anticipatory movements. We find that the dynamic responses of output and prices implied by these measures are 'well-behaved,' and that the response of output is larger than the one implied by a conventional measure of oil-price shocks proposed in the literature. We then present a dynamic general-equilibrium model and assess to what extent it can account for the response of key macroeconomic variables to our oil-price shocks.

KEYWORDS: Oil futures prices; Oil-price shocks; Economic fluctuations

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

We thank Michelle Alexopoulos, James Hamilton, Òscar Jordà, Monika Piazzesi, Pau Rabanal and Eric Swanson for helpful conversations and suggestions; seminar participants at the Bank of Italy, the 2007 SCE conference in Montreal, and the 2007 SED conference in Prague for valuable feedback; and Thien Nguyen for excellent research assistance. All remaining errors are our own. The views 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 Bank of San Francisco, or the Federal Reserve System.

1 Introduction

Large oil-price increases in the post-World War II period have often been followed by economic downturns in the U.S. economy. Ideally, in order to estimate with precision how much of the downturns can be accounted for by oil-price increases, one would have to isolate true oil-price shocks, that is, the exogenous and unanticipated component of oil-price changes.

Conventional measures of oil-price shocks based on oil-price changes have two obvious flaws: endogeneity and predictability. First, changes in spot oil prices may reflect shocks to other parts of the economy that create an imbalance in oil supply and demand, and such oil price changes may simply be endogenous response to other kinds of structural shocks. For instance, the oil price increases since 2002 are viewed by many as the result of 'an expanding world economy driven by gains in productivity' (Wall Street Journal, August 11, 2006). Kilian (2006) has indeed found that standard measures of oil-price shocks based on changes in spot oil prices do not represent the exogenous component of oil prices, but are likely to reflect movements driven by endogenous factors. Such endogenous movements may lead to biased estimates of the effects of oil shocks. Another problem associated with the price-based measures is that part of the observed price changes might have already been anticipated by the private agents well in advance, therefore they are hardly 'shocks.' The most commonly used oil prices in the literature are indeed the spot price or price quotes on the so-called 'front-month' contract, which is to be delivered in the month immediately after the trading day. However, when the market senses any substantial supply-demand imbalances in the future, changes on the near-horizon prices may not fully reflect such imbalances. Wu and McCallum (2005) find that oil futures prices are quite powerful in predicting the spot price movement, indicating that some of the spot price movement has been anticipated several months in advance.[1] This underscores the necessity of using market-based information to obtain a better measure of the exogenous oil-price shocks.

In this paper, we develop two measures of exogenous oil-price shocks for the period 1984 to 2006.

[1]Chinn, LeBlanc and Coibon (2005) find that, for 1999-2004 period, futures prices are unbiased predictors of future spot prices in the crude oil market.

Our measures are based on exogenous events that trigger substantial fluctuations in spot oil prices and are constructed to be free of endogenous and anticipatory movements. Specifically, we derive our measures by taking advantage of information available from two oil-industry trade journals, *Oil Daily* and *Oil & Gas Journal*. Based on market commentaries on daily oil-price fluctuations published on these journals, we identify major events that caused substantial oil-price fluctuations on a day-to-day basis, and isolate those events that are arguably exogenous. We, then, construct two measures of shocks around these days. Our first measure is the percent change in the one-month oil futures price around the day of an exogenous event. Our second measure is the unexpected change in oil prices as realized on the day immediately after the end of an exogenous event. After having derived our shock measures, we estimate their effects of oil-price shocks on output, prices and capacity utilization for the U.S. economy, by including them, one at a time, as exogenous variables in our regression analysis. We find that the dynamic responses of output, prices and utilization implied by the two oil-price shock measures that we develop are 'well-behaved.' Specifically, we find that, following an oil-price shock, output and capacity utilization decline and prices increase. In addition, in order to check the robustness of our empirical results, we compare the estimated effects of oil-price shocks implied by our measures with those implied by one conventional measure of oil-price shocks already proposed in the literature. We, therefore, include separately, as an exogenous regressor, a VAR-based measure of oil-price shocks. We find that the decline in output implied by our measures is larger than the one implied by the conventional VAR-based measure of oil-price shocks. Finally, after illustrating the estimated effects of our oil-price shocks, we present a dynamic general-equilibrium model and assess to what extent it can account for the effects of our oil-price shocks on a set of macroeconomic variables.

Several works in the literature on oil-price shocks have already relied on the identification of exogenous events associated with large oil-price increases to develop measures of oil-price shocks and study their effects on the U.S. economy. Hamilton (1985) isolated a number of dates characterized by dramatic increases in the nominal price of oil. As these large price changes reflected events that were likely exogenous to developments in the U.S. economy, Hamilton identified exogenous oil-supply shocks with dummy variables associated with these dates. Hoover and Perez (1994) worked with monthly data and extended the number of oil-shock dates originally proposed by Hamilton. Bernanke, Gertler and Watson (1997) used a modified version of the Hoover-Perez dates. In particular, they scaled the dummy variables by their relative importance, multiplying them by the log-change in the nominal price of oil over the three months centered on the event month. Finally, Hamilton (2003) combined a quantitative approach and a dummy-variable approach to get a measure of oil-price shocks. Specifically, he identified five military conflicts during the postwar period that were exogenous with respects to developments in the U.S. economy, and measured the magnitude of the drop in oil supply associated with each one of these historical episodes. He, then, used this variable as an instrument to isolate the component of oil-price movements attributable to the exogenous events he identified. In line with this literature, we view our work as an effort to develop a quantitative measure of unanticipated oil-price changes based on the identification of exogenous events behind substantial oil-price movements. In this respect, the approach in our work is also similar in spirit to that in the studies by Romer and Romer (2004, 2006), who combine a quantitative and a narrative approach to develop measures of monetary policy and tax shocks. Our approach also shares common elements with that of Kilian (2006), who relies on information about oil production shortfalls in OPEC countries to derive a measure of exogenous oil supply shocks and estimate its dynamic effects on macroeconomic variables. It also shares with the work of Anzuini, Pagano, and Pisani (2007) the use of futures price data to derive measures of oil-price shocks. These authors, however, pursue an identification strategy which is alternative to ours, following, more specifically, the methodology of Faust, Swanson and Wright (2004) to identify oil-price shocks in a vector autoregression.

The rest of the paper is organized as follows. Section 2 describes how we develop our measures of oil-price shocks. Section 3 illustrates the procedures that we follow to estimate the effects of oil-price shocks on output, prices and the producer price of oil for the U.S. economy. Section 4 presents the empirical results on the effects of oil-price shocks. Section 5 describes the model economy. Section 6 discusses the quantitative results implied by the model economy. Section 7 concludes.

2 Measures of oil-price shocks from daily data

This section illustrates how we construct our measures of oil-price shocks. We work with daily data on spot and futures prices for West Texas Intermediate light sweet crude oil at the New York Mercantile Exchange (NYMEX) from the beginning of 1984 to the end of June 2006. In Appendix A.2, we describe in detail the oil-price data that we use to derive our shock measures.

We derive our measures of oil-price shocks by incorporating information from two oil-industry trade journals, *Oil Daily* and *Oil & Gas Journal*. We start by identifying dates in which the spot price of crude oil changes substantially. Specifically, we select dates in which the spot price for crude oil changes by at least 5 percent. Although our choice of this threshold may seem somehow arbitrary, daily movements in the price of oil of this magnitude are far from trivial. It results that our choice of the 5 percent threshold is quite selective, In fact, in our sample of daily data containing 5635 observations, with each one of them corresponding to one trading day, we were able to identify 223 trading days in which the price of oil changed by more than 5 percent.

Our next step consists of identifying the reasons behind these substantial oil-price movements. Our purpose is to distinguish dates in which the price of oil was driven by arguably exogenous events from those in which it was driven by developments related to the state of the oil market. In order to carry out this distinction, we use information from *Oil Daily* and *Oil & Gas Journal*. Based on the daily market commentaries published on these journals, for each of the dates in which the oil-price movement was at least 5 percent, we track down the events that were behind these substantial price changes on a day-to-day basis. Among these events, we isolate those that are arguably exogenous to the state of the oil market and to other developments in the world economy. Overall, we were able to identify 25 dates in which the price of oil moved substantially in response to exogenous events. In Appendix A.1 we provide a list of these dates and of the events that were behind the corresponding oil-price movements. Most of such exogenous events are related to either political reasons (such as Yamani being removed from his post as Saudi Oil Minister in October 1986, or the coup in the Soviet Union in August 1991), military actions (such as the end of Iran-Iraq war, or the events related to the two Gulf wars), abrupt changes in weather (such as Hurricane Katrina), or natural disasters (such as the Chernobyl nuclear reactor disaster).

For the same purpose, we choose to exclude those kinds of events that are likely to be driven by endogenous responses to the state of the oil market. We exclude, therefore, events like OPEC meetings deliberating to change oil output and manipulate oil prices (such as the price war in 1986 and the subsequent efforts in driving up the prices in the late 1980s and early 1990s), changes in oil demand, as in most cases they are simply responses to changes in economic conditions (for example, the oil price dropped 2.4 percent on May 28, 2002 because of 'disappointing Memorial day demand'), as well as surprises in the announcements of oil inventory, as they reflect either surprises in oil demand or temporary supply or transportation disturbances that would arguably be dissolved very soon.

Having identified the days in which substantial oil-price movements were driven by arguably exogenous events, we construct two measures of oil-price shocks around these days. We derive our first measure of oil-price shocks following the work of Cochrane and Piazzesi (2002) on monetary policy shocks. We define this measure as the log-change in the one-month West Texas Intermediate (WTI) oil futures price from one day prior to the event day to one day following the end of the event:

$$(1) \qquad \log P^{1,m}_{t,d+1} - \log P^{1,m}_{t,d-1},$$

with $P^{1,m}_{t,d+1}$ and $P^{1,m}_{t,d-1}$ being the one-month futures oil price one day after and one day before the event, respectively. In most cases, the shock measure in (1) is the percent change in the oil futures price over a three-day window centered on the day of the event. We call this measure the 'one-month futures price change.'

Our second measure of oil-price shocks is based on a modified version of the forecasting equation in Wu and McCallum (2005). Specifically, we run the following forecasting regression of oil-price changes based on oil futures prices at different horizons on the day before the event:

$$(2) \qquad \log P^s_{t+1} - \log P^s_t = \alpha^O + \sum_{j=1}^{6} \beta^O_j (\log P^{j,m}_t - \log P^s_t) + \varepsilon_{t+1},$$

where P_t^s and P_{t+1}^s are the spot oil price at t and $t+1$, respectively, $P_t^{j,m}$ is the j-month oil futures price at time t, and α^O and β_j^O's are the estimation parameters. We, then, calculate the unexpected change in oil prices as realized on the day immediately after the end of the event, and we take this magnitude as our second measure of oil-price shocks. Hence, we define this measure as $(\log P_{t+1}^s - E_t \log P_{t+1}^s)$ and we call it the 'forecasting error.' Equation (2) incorporates term structure information on futures-spot spread in forecasting future oil price. Intuitively, equation (2) is similar to the bond-yield forecasting equation in the work of Cochrane and Piazzesi (2005), which also uses the information embodied in futures rates at all the available horizons to forecast one-period bond yield movements. Wu and McCallum (2005) compare the out-of-sample performance of such a 'futures-spot spread' model to that of several other kinds of models and conclude that the 'futures-spot spread' model performs the best, in particular when forecasting oil-price movement in the near future. On the other hand, we exclude price quotes on futures contracts beyond six months from the equation since futures market becomes much less liquid for those horizons, and the quoted futures prices become a much less accurate measure of oil price expectations. Wu and McCallum (2005) also find that the out-of-sample performance of 'futures-spot spread' model is much worse when the forecasting horizon goes beyond one year.

When we construct both our measures of oil-price shocks, we choose to examine oil price changes as realized on the day after the event day instead of on the event day. The motivation for our choice lies in the observation that in many cases the oil market tends either to overreact to the news on the event day and then correct on the following day, or to underreact on the event day and continue its response on the following day. For example, on April 7, 1986, the spot oil price jumped from \$12.74 to \$14.33, or 12.5 percent on the news of Norway oil platform worker strike and Iranian's rocket attack on a Saudi Arabian tanker, but it then fell back to \$12.47 the next day due to 'market corrections of the overreaction previous day' (*Oil Daily*, April 9, 1986). Similarly, oil price soared on April 19 and 20, 1989 on the British North-sea platform blast, by \$3.14 or 14.6 percent over two days, but fell back by an almost identical amount on April 21, 1989.

Figure 1 plots our two shock measures, along with the 'net oil-price increase,' a well-known measure of oil-price changes proposed by Hamilton (1996). The top, middle and bottom panels show the 'one-month futures price change,' the 'forecasting error,' and the 'net oil-price increase,' respectively. As it can be observed, our two measures are quite correlated. Indeed, their statistical correlation is 0.95.

After having obtained our two measures of unanticipated and arguably exogenous oil-price movements, we use them as shock variables to estimate their effects on output, prices, capacity utilization and the real price of oil for the U.S. economy.

3 Estimating the effects of oil-price shocks

This section describes the procedures that we follow to estimate the effects of oil-price shocks on output, prices, utilization and the real price of oil for the U.S. economy. We use two procedures. First, we use a univariate autoregressive model for each of the four variables we are studying. After that, we estimate a vector autoregressive model that includes all the variables used in the first procedure. In both the models that we use, we include our two oil-price shock variables, one at a time, as exogenous regressors. In addition, in order to check the robustness of our empirical results, we also estimate both models including separately, as an exogenous regressor, a VAR-based measure of oil-price shocks. This allows us to compare the estimated effects of our oil-price shock measures with those of one standard measure of oil-price shocks already proposed in the literature.

The VAR-based measure of oil-price shocks that we use is equal to the fitted residual series from a least squares regression of one indicator of oil-price movements on its own lagged values, current and lagged values of output, prices, and utilization. Identifying oil-price shocks with these residuals is equivalent to estimating a VAR including output, prices and utilization and the indicator of oil-price movements, ordering this oil-price variable last, and then identifying the shocks with the VAR innovations to the indicator of oil-price movements. As indicator of oil-price movements, we choose the 'net oil-price increase' variable proposed by Hamilton (1996). This choice is motivated by the well known finding that this variable has a stable relationship with macroeconomic variables. The net oil price increase is defined as follows: it is the maximum of zero and the difference between the current oil-price log-level and the maximum value of the oil-price log-level during the previous year. It indicates, in particular, the amount by which oil prices in a given period are above their peak value over the previous year. If oil prices are not above their previous peak, then this variable is equal to zero. The purpose of this indicator is that of distinguishing oil-price increases that establish new highs relative to recent experience from increases that simply reverse recent decreases.

We normalize the shock measures that we use in our analysis so that they all induce a cumulative 10-percent increase in the real price of oil over a 24-month horizon. We pursue this normalization to be able to compare the results that we obtain using different shock measures, as these measures, originally, have different distributions.

Our first step consists of estimating by least squares a univariate autoregressive model where the current value of one endogenous variable is regressed on its own lagged values, on current and lagged values of one of our oil-price shock variables, and on a constant and a time trend. The corresponding estimating regression is therefore:

$$(3) \qquad x_t = a_0 + a_1 t + \sum_{i=1}^{6} \alpha_i x_{t-i} + \sum_{i=0}^{6} \beta_i O_{t-i} + \epsilon_t,$$

where the left-hand side variable, x_t, represents the endogenous variable, t denotes a time trend variable that starts at the beginning of our sample, O_t is one oil-price shock variable, and ϵ_t is a disturbance term. In the above equation, a_0, a_1, α_i, and β_i represent the coefficients to be estimated. We estimate equation (3) for each of the endogenous variables we are studying. Therefore, the variable x_t corresponds to one element in a set of macroeconomic variables for the U.S. economy. This set includes output, the consumer price index (CPI), capacity utilization in the manufacturing sector, and the real price of oil, that is the producer price index (PPI) for crude oil prices deflated by the CPI. In addition, for each one of these variables, we estimate equation (3) three times, that is, once for each one of the measures of oil-price shocks that we consider. Two of these measures are those that we develop in this paper and that we described in Section 2, namely the 'one-month futures price change' measure and the 'forecasting error' measure. The third measure that we consider is the VAR-based measure of oil price shocks. We use the coefficients $\hat{\alpha}_i$ and $\hat{\beta}_i$ estimated from (3) to simulate the dynamic impact of an oil-price shock on the endogenous variable, x_t. These simulations, therefore, represent the estimated dynamic responses to a one-percent increase in the average value of the oil-price shock variable, where this average value is computed over the sample period we use in our estimation. This procedure is very similar in spirit to the one used by Ramey and Shapiro (1998) and Kilian (2006) to estimate the dynamic responses of key macroeconomic variables to a fiscal policy shock and to an oil-supply shock, respectively.

Our second procedure to estimate the effects of oil-price shocks consists of including our three endogenous variables in a vector autoregressive model be estimated using equation-by-equation least squares. Christiano, Eichenbaum and Evans (1999) and Burnside, Eichenbaum and Fisher (2004) used this procedure to estimate the effects of identified monetary policy shocks, and fiscal policy shocks, respectively. The corresponding estimating regression is therefore:

$$(4) \qquad X_t = A_0 + A_1 t + A_2(L) X_{t-1} + B(L) O_t + \varepsilon_t,$$

where X_t is a vector that contains as elements each one of our endogenous variables, x_t's t denotes, again, a time trend variable that starts at the beginning of our sample, O_t is one of our oil-price shock variable, and ε_t is a vector of disturbance terms. In (4), A_0 and A_1 are vectors of coefficients, while $A_2(L)$ and $B(L)$ are sixth-order vector polynomials in nonnegative power of the lag operator L. The coefficients to be estimated, therefore, are the elements of A_0 and A_1, and the coefficients in the vector polynomials $A_2(L)$ and $B(L)$. As it was the case for the estimating regression (3), we also estimate regression (4) three times, that is, once for each one of the measures of oil-price shocks that we consider. The estimated dynamic responses of the endogenous variables in X_t to an oil-price shock are then given by the estimates of the coefficients on L^k in the expansion of $[I - A_2(L) L]^{-1} B(L)$.

The left-hand side variables that we use in (3) and that we also include in the vector X_t in (4) are the following variables multiplied by one hundred: the log of real GDP, the log of the consumer price index, the log of capacity utilization in the manufacturing sector, and the log of the producer real price of oil. We estimate equations (3) and (4) using monthly data from 1984:01 to 2006:06.

Our sample starts from the beginning of 1984 because we could find issues of *Oil Daily* and *Oil & Gas Journal* starting only from that date. We obtain a monthly series for real GDP by interpolating with the method of Chow and Lin (1971) the available quarterly series. As interpolators, we use the monthly series for industrial production and total capacity utilization.[2] Appendix A.2 describes the data used in our analysis.

4 Empirical results

This section presents the results we obtained following the estimation procedures outlined in the previous section. In Figures 2 through 8 we plot the estimated dynamic responses of output, CPI, utilization, and the real price of oil prices to a one-percent increase in the average value of the oil-price shock variable that we consider. Figures 2 through 4 illustrate the dynamic responses obtained using the estimates of the coefficients from the univariate autoregressive model in equation (3). Figures 5 through 8 illustrate the dynamic responses obtained using the estimates of the coefficients from the vector autoregressive model in equation (4). In all the figures, the top left and right panels show the response of output and the CPI, respectively, while the bottom left and right panels show the response of capacity utilization and of the real producer price of oil. Figures 4, 7, and 8 compare the estimated dynamic responses obtained using the two oil-price shock measures that we develop with those obtained using the VAR-based measure of shocks. We use solid lines with square markers to denote the point estimates of the responses to the 'one-month futures price change' measure of oil-price shocks. We use solid lines with cross markers to denote the point estimates of the responses to the 'forecasting error' measure of oil-price shocks. We use solid lines with diamond markers to denote the point estimates of the responses to the VAR-based measure of oil-price shocks. Finally, we denote with shaded areas the 95-percent confidence intervals for the point-estimates of the dynamic responses.[3]

[2] The corresponding high-frequency correlations in levels and in year-on-year growth rates are 0.994 and 0.736, respectively.

[3] We computed the confidence intervals following a bootstrap Monte Carlo procedure. Specifically, with T being the length of the sample period that we consider in the empirical analysis, we computed 500 artificial time series of length T on the variable x_t as follows. We constructed 500 new time series of residuals $\{\hat{u}_t(j)\}_{t=1}^{T}$, $j = 1, \ldots, 500$, by drawing randomly with replacement from the vector of fitted residuals $\{\hat{u}_t\}_{t=1}^{T}$ from equations (3) and (4), respectively.

Our results show that the dynamic responses of output, prices, utilization and the real price of oil implied by the two oil-price shock measures that we develop are 'well-behaved.' In particular, we find that, following an oil-price shock, output and utilization decline and prices increase. In line with the findings of the literature on oil-price shocks, we also find that the dynamic effects of oil-price shocks on output are not quantitatively large.

Figure 2 illustrates the dynamic responses of output, the CPI, capacity utilization, and the producer real price of oil obtained estimating the univariate autoregressive model of equation (2) and using the 'one-month futures price change' measure of oil-price shock as a shock variable. This figure shows, that following an oil-price shock, output starts declining about two months after the date of the shock. The figure also shows that the point estimate of the dynamic response of output is increasingly negative. After 24 months, the sum of the point estimates of the output response is equal to -0.3 percent. The price level increases following an oil-price shock. The point estimate relative to the dynamic response of this variable becomes statistically significant after two months, and, over the 24-month horizon, the cumulative increase in the price level is equal to 0.29 percent. As it happens with output, capacity utilization also declines following the oil-price shock. The negative response of this variable becomes statistically significant after about nine months, and its cumulative decline after 24 months is -0.7 percent. As one would normally expect, the oil-price shock leads to an increase in the producer real price of oil. This increase displays a hump-shaped pattern. The point estimate is statistically significant for about 7 months after the shock has occurred. Relative to its preshock level, the producer price of oil reaches a peak three months after the shock, and then it reverts to zero.

Figure 3 shows the estimated dynamic responses from the univariate autoregressive model using the 'forecasting error' measure as a shock variable. The figure shows that the responses induced by

For each constructed series of new residuals, we computed an artificial time series $\{\hat{x}_t(j)\}_{t=1}^{T}$, $j = 1, \ldots, 500$, using the estimated equation and the historical initial conditions on x_t. We then reestimated equations (3) and (4) using $\{\hat{x}_t(j)\}_{t=1}^{T}$ and we calculated the implied dynamic response function for $j = 1, \ldots, 500$. For each fixed lag, we computed the 12^{th} lowest and the 487^{th} highest values of the corresponding dynamic response coefficients across all 500 artificial dynamic response functions. The boundaries of the shaded areas in Figures 2 through 8 correspond to a graph of these coefficients.

an oil-price shock as measured by our 'forecasting error' variable are broadly similar to those induced when we measure the shock with our 'one-month futures price change' variable. Specifically, output and utilization decline, while the CPI and the producer real price of oil increase. The cumulative response of output over 24 months in this case is -0.22 percent. The CPI increases following the oil-price shock, and, as in the previous case shown in Figure 2, its response becomes statistically significant two months after the shock. The sum of the point estimates of the response of prices is equal to 0.24 percent. Capacity utilization declines, and it 24-months cumulative drop is -0.53 percent. As we obtained before, the producer real price of oil increases in a hump-shaped way, and the point estimates of the response are generally statistically different from zero over the 24-month horizon. The peak response in this variable happens three months after the shock also in this case.

Figure 4 compares the estimated dynamic responses from the autoregressive model obtained using our 'one-month futures price change' and 'forecasting error' measures of oil-price shocks with those obtained using a more conventional VAR-based measure of shocks. The figure shows that the decline in output implied by the 'one-month futures price change' measure of oil-price shocks is larger than the output decline implied by the VAR-based measure of shocks. In particular, the cumulative decline in output implied by the VAR-based measure is equal to -0.15 percent, compared with a cumulative decline of -0.3 implied by the 'one-month futures price change' measure. This implies that the cumulative decline in output obtained with the 'one-month futures price change' measure is twice larger than the cumulative decline obtained with the VAR-based measure of oil price shocks. The top left panel of Figure 4 also shows that decline in output implied by the 'forecasting error' measure is larger in magnitude than the decline implied by the VAR-based measure. In regard to the response of the CPI, our measures of oil-price shocks imply a cumulative increase which is larger than the one implied by the VAR-based measure. In fact, the cumulative increase in prices obtained with the VAR-based shock measure is virtually zero, lower than the 0.29 percent and 0.22 percent cumulative increases obtained with the 'one-month futures price change' and the 'forecasting error' shock measures, respectively. The cumulative responses of utilization induced by our shock measures are also larger in magnitude than the cumulative response induced by the

VAR-based measures. While our first and second measure produce a cumulative 24-month decline of -0.70 and -0.53 percent, respectively, the VAR-based measure produces a decline of -0.44 percent. Finally, the estimated dynamic response implied by the VAR-based shock measure indicates that, on impact and in the months immediately after the shock, the producer real price of oil increases by an amount larger than the one implied by our shock measures. Our measures, however, induce an hump-shaped pattern for real price of oil and a response which, six months after the shock, is larger than the one obtained with the VAR-based measure.

After having examined the dynamic responses of output, the CPI, utilization and the real price of oil estimated using the univariate autoregressive model of equation (3), we turn to analyze the responses estimated using the vector autoregressive model of equation (4). The results from the vector autoregressive model can be broadly summarized as follows. The point estimates of the dynamic responses of output and capacity utilization obtained using the vector autoregressive model are larger in absolute value than the corresponding point estimates obtained using the univariate autoregressive model. In addition, these estimates become statistically different from zero about after only a few months following the shock. In regard to the estimated dynamic response of prices, the point estimates obtained using the vector autoregressive model are somehow larger than the corresponding point estimates obtained using the univariate autoregressive model.

Figure 5 plots the dynamic responses estimated using the 'one-month futures price change' shock measure. Differently from what obtained estimating the univariate autoregressive model, the figure indicates that output starts declining right after the date the shock occurs and deteriorates for several months thereafter. The point estimate of the response of output becomes significantly different from zero five months after the shock. In addition, the estimated cumulative decline in output after 24 months relative to the preshock level is equal to -0.77 percent, compared with -0.3 percent estimated using the univariate autoregressive model. The estimated dynamic response of prices in the earlier periods after the shock is similar in magnitude to the one obtained with the univariate autoregressive model. However, the estimated response obtained from the autoregressive model declines towards its preshock level more rapidly. Over the 24-month horizon, the point estimate of the dynamic response is statistically significant only in the first few months following the shock, and the cumulative increase in prices is equal to 0.17 percent compared with the 0.29 percent implied by the univariate autoregressive model. The response of capacity utilization is larger than the one obtained from the univariate model. The lower left panel of Figure 5 shows that utilization declines in a hump-shaped way reaching a through about one year after the shock. The cumulative decline in this variable over 24 months is -0.99 compared with -0.70 obtained from the univariate model. The producer price of oil increases following the shock, and its dynamic response is quite similar to the one obtained using the univariate autoregressive model. Relative to its preshock level, the producer price of oil reaches a peak of about 1 percent after 3 months, and it slowly reverts to zero afterwards.

Figure 6 shows the dynamic responses estimated using the vector autoregressive model and the 'forecasting error' measure of oil-price shocks. As in the univariate case, the figure shows that output starts declining two months after the shock, and that it keeps declining for several additional months. The point estimate of the response of output becomes statistically different six months after the shock, and the cumulative decline in output is equal to -0.61 percent compared with -0.22 percent obtained estimating the univariate autoregressive model. The estimated response of prices is somehow larger in the earlier periods after the shock relative to the estimated response obtained with the univariate autoregressive model. However, it is less persistent and it reverts back to its preshock level after reaching a peak six months after the shock. The sum of the point estimates of the response of prices is 0.16 percent compared with 0.24 percent implied by the estimates of the univariate model. The response of capacity utilization obtained with the vector autoregressive model is negative and larger in absolute value than the corresponding response using the univariate model. It becomes statistically different from zero approximately six months after the shock and it reaches a through one year after the shock. The cumulative response of utilization is -0.85 compared with -0.53 obtained with the univariate model. The real producer price of oil increases in a hump-shaped pattern following the shock, and the point estimates of the response are statistically significant for the first six months after the shock occurs. The peak response happens three months after date of the shock and it is equal to 0.95 percent relative to its preshock level.

Finally, Figures 7 and 8 compare the estimated dynamic responses from the vector autoregressive model obtained using our two measure of oil-price shocks with those obtained using the VAR-based shock measure. The top left panels in these figures show that the point estimates of the decline in output after an oil-price shock implied by the VAR-based measure are lower than the point estimates implied by the 'one-month futures price change' and the 'forecasting error' measures. Specifically, the cumulative decline in output implied by the VAR-based shock measure is -0.14 percent, compared to -0.77 percent and -0.61 percent implied by the 'absolute price change' and the 'forecasting error' measures, respectively. In addition, for several periods after the shock, the response of output implied by the VAR-based measure are outside the 95-percent confidence bands relative to the estimated responses of output implied by our two-measures. In regard to the estimated dynamic responses of prices, the VAR-based shock measure implies a smaller response. In cumulative terms, the estimated response of prices obtained using the VAR-based measure is essentially zero, compared with -0.17 and -0.16 percent obtained with our two measures. The responses of capacity utilization, plotted in the left bottom panels indicate the VAR-based measure imply an estimated response smaller than the response obtained our two shock measures. Figures 7 and 8 also show that, for a few periods, the response implied by the VAR-based measure lie outside the upper bounds of the 95-percent confidence bands relative to the estimated responses implied by our two shock measures. In cumulative terms, the decline in capacity utilization implied by the VAR-based measure is -0.2 percent, compared with -0.99 and -.084 percent implied by the 'one-month futures price change' and the 'forecasting error,' respectively. Finally, the VAR-based oil-price shock measure induces an immediate increase on impact in the real producer price of oil, while our two shock measures induce a hump-shaped estimated response of the real price of oil after the shock.

Summing up, our results indicate that our measures of oil-price shocks imply estimated dynamic responses of output, prices, capacity utilization and the producer price of oil that have the expected sign. In particular, we find that output and capacity utilization decline and that the price level increases. We also find that the declines in output and capacity utilization implied by our measures are substantially larger than the corresponding responses implied by a conventional VAR-based measure of oil-price shocks. When we estimate a vector autoregressive model, the comparison is more striking, as the estimated responses of output and utilization obtained using our measures are significantly different form the corresponding responses obtained using a VAR-based measure.

5 The model economy

This section describes the dynamic general-equilibrium model economy that we use to assess the effects of oil-price shocks. The model incorporates a number of features that are present in the studies by Christiano, Eichenbaum and Evans (2005), Leduc and Sill (2004), and Levin, Onatski, Williams and Williams (2006). It is, essentially, a monetary model with monopolistic competition, nominal rigidities, and in which the use of energy is essential for capital utilization. The set of agents includes final-good firms, intermediate-goods firms, and households, all of whom are infinitely-lived. It also includes a monetary authority. Finally, time is discrete and each time period, representing one month, is indexed by the subscript t.

Final-good firms. There is large number of perfectly competitive identical firms. The representative firm produces the final good, Y_t, by combining a continuum of differentiated intermediate goods according to the following CES aggregator function:

$$(5) \qquad Y_t = \left[\int_0^1 Y_t\left(i\right)^{1/(1+\lambda_p)} di \right]^{1+\lambda_p},$$

where $Y_t\left(i\right)$, with $i \in [0,1]$, represents the intermediate-good input of type i, and where the parameter λ_p is a strictly positive constant. The aggregate output of the final good can be used for both consumption and investment.

Intermediate-goods firms. There is a continuum of intermediate-goods firms, with each of these firms and the goods they produce indexed by $i \in [0,1]$. Each typical firm i produces its own differentiated good, $Y_t\left(i\right)$, according to the following Cobb-Douglas production function:

$$(6) \qquad Y_t\left(i\right) = \left(Z_t\left(i\right) K_t\left(i\right)\right)^{1-\alpha} N\left(i\right)^{\alpha},$$

where $Z_t(i)$, $K_t(i)$, and $N(i)$ denote firm i's capital utilization rate, beginning-of-period capital stock, and labor input, respectively, with $\alpha \in [0,1]$ representing the labor-output share.

Intermediate-goods firms own the economy's capital stock and make investment decisions. As in Greenwood, Hercowitz and Huffman (1988) and Leduc and Sill (2004), the depreciation rate of capital, $\delta(\cdot)$ is a convex and increasing function of the capital utilization rate $Z_t(i)$, with $\delta'(\cdot)$ and $\delta''(\cdot)$ strictly positive. Consistent with the empirical findings of Basu and Kimball (1997), we assume the following functional form for the depreciation rate function:

$$\delta(Z_t(i)) = \delta + \Psi_0 \frac{Z_t(i)^{1+\Psi_1} - 1}{1+\Psi_1}, \tag{7}$$

where δ, Ψ_0 and Ψ_1 are strictly positive constants, with $\Psi_1 = Z_t(i)\,\delta''(\cdot)/\delta'(\cdot)$ representing the elasticity of marginal depreciation, $\delta'(\cdot)$, with respect to the capital utilization rate, $Z_t(i)$.

We assume that there are adjustment costs to changing the level of investment, as in Christiano, Eichenbaum and Evans (2005). These authors argue that introducing a specification of this kind helps generate a hump-shaped response of aggregate investment to shocks. We choose to consider this specification to improve the ability of the model in generating an inertial response of output to an oil-price shock, consistent with the empirical results of Section 4.

Within this setup, therefore, firm i's capital stock evolves according to:

$$K_{t+1}(i) = [1 - \delta(Z_t(i))]\,K_t(i) + I_t(i)\left[1 - S\left(\frac{I_t(i)}{I_{t-1}(i)}\right)\right], \tag{8}$$

with $K_{t+1}(i)$ and $I_t(i)$ denoting firm i's end-of-period capital and investment expenditure, respectively, and with the function $S(\cdot)$ representing investment adjustment costs. This function is such that only the dynamics of the model are affected by the presence of investment adjustment costs, but not its steady state. Specifically, the function S satisfies the following properties: $S(1) = S'(1) = 0$, and $S''(1) = s > 0$, with the parameter s reflecting the magnitude of investment adjustment costs.

As in Finn (1995, 2000) and Leduc and Sill (2004), we assume that capital utilization requires the use of energy. Each firm i must therefore purchase energy, whose price, P_t^e, we treat as exogenous. For the typical firm i capital utilization is related to use of energy per unit of capital as follows:

$$\frac{e_t(i)}{K_t(i)} = a(Z_t(i)), \tag{9}$$

with $a'(\cdot)$ and $a''(\cdot)$ strictly positive, and with the following functional form for the utilization function:

$$(10) \quad a(Z_t(i)) = \frac{Z_t(i)^{\upsilon}}{\upsilon},$$

where $e_t(i)$ denotes firm i's use of energy, and where the parameter υ is strictly larger than one.

This specification captures the notion that the use of energy is essential to obtain a positive rate of capital utilization, and that a more intense capital utilization requires a larger use of energy, at an increasing rate. More importantly, as discussed in Finn (2000), this specification allows increases in the price of energy to have contractionary effects on output through a number of channels. When energy becomes more expensive, the marginal cost of utilization increases and firms reduce the degree of capital utilization. This reduction, in turn, has a negative impact on the magnitude of capital services, and, therefore, on output.[4] In addition, the decrease in the degree of capital utilization causes a reduction in labor productivity and in the real wage, which, in turn, affects negatively labor supply and output. Finally, the increase in the price of energy, together with the resulting reductions in utilization and labor supply induce a reduction in the marginal return to capital, and, therefore, in the marginal return to investment. This factor works through reducing the level of investment expenditure and of the stock of capital, producing ultimately a negative output effect.

We also consider, as an additional feature of the model economy, adjustment costs to capital utilization, as in Jaimovich and Rebelo (2006). Within a framework where the degree of capital utilization depends on energy, an increase in the price of energy causes a drop on impact in capital utilization. Ideally, introducing adjustment costs to utilization should help the model capture the estimated response of utilization to the oil-price shock by dampening the impact response of utilization to an increase in the price of energy. Adjustment costs to utilization are represented by $\Phi(Z_t/Z_{t-1})Z_t$, with $\Phi(1) = \Phi'(1) = 0$, and with $\Phi''(\cdot) = \phi > 0$ reflecting the magnitude of utilization adjustment costs. Similarly to the function representing investment adjustment costs,

[4] A slightly different mechanism through which the degree of capital utilization helps propagate the effects of oil-price shocks is present in Aguiar-Conraria and Wen (2007). Their mechanism is based on the notion of externalities among firms and increasing returns to scale.

the function $\Phi(\cdot)$ affects only the dynamics of the model but not its steady state.

Each intermediate-goods firm i is a monopolistically competitive producer of its own differentiated good. It faces, therefore, a downward-sloping demand function and is able to set the price for its own product of type i. We assume that firms set their price according to a Calvo-style mechanism. In particular, each period the typical firm i faces a constant probability $1-\xi_p$ of resetting its price, with $\xi_P \in [0,1]$. As in Levin et al. (2006), we also allow for partial indexation. Specifically, firms that are not able to reoptimize can partially adjust their price according to the lagged rate of price inflation, with $\gamma_p \in [0,1]$ representing the degree of price indexation to lagged inflation. This implies that the aggregate price index, P_t, evolves according to:

$$(11) \quad P_t = \left[\xi_p P_{t-1}^{-1/\lambda_p}\left(\frac{P_{t-1}}{P_{t-2}}\right)^{-\gamma_p/\lambda_p} + \left(1-\xi_p\right)\tilde{W}_t^{-1/\lambda_p}\right]^{-\lambda_p},$$

where $\tilde{P}_t$ represents the optimal wage chosen by the firms that can reoptimize in period t.

Households. There is a continuum of infinitely-lived households, with each household indexed by $h \in [0,1]$. The preferences of the typical household h are described by the following utility function:

$$(12) \quad U(h) = E_0 \sum_{t=0}^{\infty} \beta^t u_t(h),$$

where $\beta \in (0,1)$ is the subjective discount factor, and $u_t(h)$ is the period utility function. In each period t, households obtain utility from consuming final goods and from holding real cash balances. They also derive disutility from supplying labor hours to intermediate-goods firms. The period utility function is therefore:

$$(13) \quad u_t(h) = \log\left(C_t(h) - bC_{t-1}(h)\right) + \frac{\theta_N}{1-\gamma_N}\left(1-N_t(h)\right)^{1-\gamma_N} + \frac{\theta_M}{1-\gamma_M}\left(M_t(h)\right)^{1-\gamma_M},$$

where $C_t(h)$, $N_t(h)$, and $M_t(h)$ denote household h's final goods' consumption, labor hours, and real cash balances, respectively, and where the preference parameters θ_N, γ_N, θ_M, and γ_M are all strictly positive constants. The period utility function in (13) includes habit formation in consumption, with the parameter $b \in [0,1]$ reflecting the degree of habit persistence. This is an additional element that helps the model generate an inertial response of output to an oil-price shock. The period utility function also implies that the Frisch elasticity of labor supply is $\eta_N = (1-N)/(\gamma_N N)$, with N being the typical household's steady-state amount of labor hours. It also implies that the marginal rate of substitution between consumption and labor hours is given by:

$$(14) \quad MRS_t(h) = \theta_N (1 - N_t(h))^{-\gamma_N} C_t(h).$$

As in Erceg, Henderson and Levin (2000), each household h is a monopolistically competitive supplier of its own differentiated labor service, $N_t(h)$. The labor service of household h is allocated across each intermediate-goods firm i:

$$(15) \quad N_t(h) = \int_0^1 N_t(i,h)\, di,$$

where $N_t(i,h)$ represents the labor input of type h for the intermediate-goods firm i. In turn, the various differentiated labor services of type h, $N_t(i,h)$, are combined into firm i's labor input, $N_t(i)$, through the following CES aggregator function:

$$(16) \quad N_t(i) = \left[\int_0^1 N_t(i,h)^{1/(1+\lambda_w)}\, dh\right]^{1+\lambda_w},$$

where the parameter λ_w is a strictly positive constant. Each household, therefore, faces a downward-sloping labor demand function and is able to set its own nominal wage. Following the work of Erceg, Henderson and Levin (2000) and Christiano, Eichenbaum and Evans (2005), we assume that households set their wage according to a Calvo-style mechanism. In particular, each period the typical household h faces a constant probability $1 - \xi_w$ of resetting its wage, with $\xi_w \in [0,1]$. Following the work of Levin et al. (2006), we also allow for partial wage indexation. Specifically, households that are not able to reoptimize can partially adjust their wage according to the lagged rate of price inflation, with $\gamma_w \in [0,1]$ representing the degree of wage indexation to lagged inflation. This implies that the aggregate nominal wage index, W_t, evolves according to:

$$(17) \quad W_t = \left[\xi_w W_{t-1}^{-1/\lambda_w} \left(\frac{P_{t-1}}{P_{t-2}}\right)^{-\gamma_w/\lambda_w} + (1-\xi_w)\, \tilde{W}_t^{-1/\lambda_w}\right]^{-\lambda_w},$$

where $\tilde{W}_t$ represents the optimal wage chosen by the households that can reoptimize in period t.

Monetary authority. We assume that the monetary authority conducts monetary policy with an interest-rate rule. Specifically, it sets the short-term nominal interest rate, i_t, in response to the lagged interest rate and to deviations of aggregate price inflation and aggregate output from their steady-state levels. It adjust, therefore, i_t, according to the following rule:

$$(18) \quad i_t = (i_{t-1})^{\rho_i} \left[\left(\frac{P_t}{P_{t-1}} \right)^{\rho_\pi} \left(\frac{Y_t}{Y} \right)^{\rho_y} \right]^{1-\rho_i},$$

where Y denotes the steady-state level of aggregate output.

The real price of energy. The real price of energy, P_t^e/P_t, is the only exogenous driving variable, and it is taken as given by all firms and households. We describe the evolution of this variable as follows:

$$(19) \quad \frac{P_t^e}{P_t} = \frac{P^e}{P} + h_e(L)\,\varepsilon_t,$$

where P^e/P is the steady-state real price of energy, $h_e(L)$ is an ordered polynomial in nonnegative powers of the lag operator L, and ε_t is a zero-mean innovation to the price of energy.

Market clearing and equilibrium. In equilibrium, all inputs, intermediate-goods, and final-good markets clear. This requirement, along with the optimality conditions for all firms and households, implies the following aggregate resource constraint:

$$(20) \quad Y_t = C_t + I_t + P_t^e e_t + \Phi\left(\frac{Z_t}{Z_{t-1}} \right) Z_t.$$

This constraint indicates that aggregate final-good output is allocated between aggregate consumption, investment, energy expenditure and utilization adjustment costs.

6 Quantitative results and implications

This section discusses the quantitative implications for how the model economy presented in the previous section reacts to an oil-price shock. In our model simulations, the exogenous driving process is represented by the estimated dynamic response of the real price of oil to one of our two oil-price shock measures, in turn. To obtain our results, we log-linearize the equilibrium conditions of the model economy around its steady state, and we assign numerical values to the parameters of the log-linear system. Before turning to the quantitative results, we describe these numerical values.

We set the subjective discount factor β equal to 0.9946, consistent with an annual real interest rate equal to 6.5 percent. We set N, the steady-state amount of hours, equal to 0.2 and the Frisch elasticity of labor supply, η_N, equal to 1. The steady-state depreciation rate, δ, is such that, on an annual basis, capital depreciates by 10 percent. The labor share of income, α, is equal to 0.33. We set Ψ_1, the elasticity of marginal depreciation with respect to the capital utilization rate, equal to 1, consistent with the point estimate in Basu and Kimball (1997). We set the value of the parameter υ equal to 1.7, which is broadly consistent with the numerical values used by Finn (1995, 2000) and Leduc and Sill (2004). The share of energy expenditure on GDP is equal to 4 percent. We set the values of the habit parameter, b, and of the investment adjustment cost parameter, s, equal to 0.65 and to 2.5, respectively. These are the values of the point estimates obtained by Christiano, Eichenbaum and Evans (2005). The value of the utilization adjustment cost parameter, ϕ, is equal to 0.4. In regard to the parameters relative to price and wage determination, we use the estimates of Levin, Onatski, Williams and Williams (2006). The Calvo price-setting and wage-setting parameters, ξ_p and ξ_w, are equal to 0.83 and 0.79, respectively. The price and wage indexation parameters, γ_p and γ_w, are equal to 0.08 and 0.79, respectively, and the wage mark-up parameter, λ_w, is equal to 0.2. The values of the parameters in the monetary policy rule, ρ_i, ρ_π, and ρ_y, are standard. We set these values equal to 0.9, 1.5, and 0.5, respectively. Finally, we set the values of the coefficients in the polynomial $h_e(L)$ equal to the point estimates of the corresponding 24 coefficients in the dynamic response of the real price of oil to our oil-price shock measures, as estimated through equation (4) in Section 3, and shown in Figures 5 and 6.

After having described the parameters values that we use to derive our model simulations, we turn to examine the dynamic effects of oil-price shocks on the same set of macroeconomic variables that we considered in our empirical analysis in Sections 3 and 4. These variables are output, the consumer price index, as indicated in equation (11), and capital utilization. Differently from the normalization we pursued in Section 3, in our simulations we normalize our shock measures, so that they induce on impact a 1-percent increase in the real price of oil. We display our model simulations in Figures 9 and 10 and we compare them to the estimated dynamic responses that we obtained from our empirical analysis. Figure 9 shows the dynamic responses of the macroeconomic variables to our 'one-month futures price change' measure of oil-price shocks. The solid lines with no markers denote the simulated dynamic responses obtained from the model economy, whereas the solid line with square markers denote the dynamic responses estimated using the vector autoregressive model in equation 4. Figure 10 shows the dynamic responses of the macroeconomic variables to our 'forecasting error' measure of oil-price shocks. As in Figure 9, the solid lines with no markers denote the simulated dynamic responses obtained from the model economy. The solid line with cross markers denote the dynamic responses estimated using the vector autoregressive model in equation (4). In both Figures 9 and 10, shaded areas denote 95-percent confidence intervals.

Both Figures 9 and 10 show that, following an oil-price shock, output and capital utilization decline and prices increase. These responses are qualitatively consistent with their empirical counterparts estimated in Section 4. Capital utilization declines mainly because the marginal energy cost of utilization increases. As discussed in the previous section, output declines because capital utilization decreases, but also because the amounts of the labor and capital inputs decline. Output also falls for an additional reason. When the price of energy increases, real marginal cost increases relative to its steady-state level, and this leads to an increase in inflation. The monetary authority reacts to the increase in inflation by increasing the nominal interest rate according to the monetary policy rule described in equation (18). With nominal rigidities, an increase in the short-term nominal interest rate translates into an increase in the real interest rate. This increase in the real interest rate, in turn, discourages consumption and investment expenditure, providing ad additional channel through which an increase in the price of energy leads to a decline in output.

Considering, first, the effects of an oil-price shock as represented by our 'one-month futures price change' measure, Figure 9 shows that the model that we presented in the previous section has only a limited ability in accounting for the effects on output and capital utilization, while it does a good job in accounting for the effects on the price level. In regard to the response of output, as it can be seen from the top left panel of Figure 9, the model predicts that output falls on impact and that it reaches a trough of -0.08 percent relative to its preshock level after five periods. Figure 9 also shows the model has a limited ability in capturing the size and the persistence of the estimated response of output to our measure of oil-price shocks. In fact, the estimated response indicates that output barely moves on impact after the shock, and that it reaches its lowest level after 18 months, eventually reverting to its preshock level more slowly than what predicted by the model. These differences between the simulated and the estimated responses are reflected in the implied cumulative decline in output. The simulated response implies that output falls by 1 percent on a cumulative basis over a 24-month horizon, whereas the corresponding decline for the estimated response is -2.44 percent. In regard to the response of the consumer price level, the top right panel of Figure 9 show that the model does quite a good job at accounting for the hump-shaped response of this variable as estimated from the data. The simulated response of the CPI comes quite close to its estimated counterpart and it lies within its 95-percent confidence interval. In terms of the cumulative increase in the CPI induced by the oil-price shock, the model predicts that the CPI increase by 0.75 percent on a cumulative basis over 24 months, compared with 0.55 implied by the corresponding estimated response. In regard to the response of capacity utilization, as it can be seen from the lower left panel of Figure 9, the model misses the timing of the estimated response of this variable to our oil-price shock. Specifically, the simulated response indicates that capacity utilization falls on impact and that it reaches through of -0.15 percent after five periods. The estimated response, on the other hand, indicates that utilization barely responds on impact and that it reaches a comparable through of -0.2 percent after 12 months. On a cumulative basis, the simulated and the estimated responses of utilization imply a decline of -1.97 percent and -3.15 percent, respectively.

Considering, next, the effects of an oil-price shock as represented by our 'forecasting error' measure, Figure 10 shows simulated responses that are quite similar to those illustrated in Figure 9. Specifically, the top left panel shows that the response of output implied by the model cannot capture the timing and the persistence of the corresponding estimated response. While the estimated response indicates that output essentially does not react on impact and that it reaches its lower level about 18 months, the simulated response shows that output falls on impact following the oil-price shock and that it reaches its lowest level after 6 months. On a cumulative basis, the drop in output implied by the simulated model response over 24 months is -0.96, compared with a decline of -

1.83 percent implied by the estimated response. Turning to the response of prices, the top right panel of Figure 10 shows that the model does a good job at capturing the estimated hump-shaped response of the CPI, similarly to what we found in the previous case of the 'one-month future price change' measure. In terms of cumulative increase, the model predicts that the CPI increase by 0.78 percent over a 24-month horizon, compared with an increase of 0.49 percent implied by the estimated response. Finally, also for the case of the 'forecasting error' measure, the model has a limited ability of capturing the timing of the estimated response of capacity utilization. In particular, while the model predicts that utilization falls on impact and that it reaches a through of -0.14 percent after 6 months, the estimated response indicates that utilization essentially does not react on impact and that it reaches a comparable trough of -0.16 percent after 12 months. On a cumulative basis, the decline in utilization implied by the simulated and estimated responses are -1.86 percent and -2.53 percent, respectively.

Summing up, we have considered a monetary model with nominal and real frictions, and in which the use of energy is essential for capital utilization. Within such a framework, an increase in the price of energy leads to a decline in output and in capital utilization and to an increase in the consumption price index. We have examined how the simulated dynamic responses of these macroeconomic variables implied by our oil-price shock measures compare to the corresponding estimated responses. In doing this, we have considered the estimated response of the real price of oil to our shocks as the exogenous driving variable. We have found that while the model does a good job at capturing the response of the CPI, it has a limited ability of accounting for both the timing and the persistence of the estimated responses of output and utilization.

7 Conclusions

In this paper, we developed two measures of exogenous oil-price shocks for the period 1984 to 2006. These measures are constructed to be free of endogenous and anticipatory movements. We derived our measures by incorporating information from two oil-industry trade journals, *Oil Daily* and *Oil & Gas Journal*. Based on market commentaries published on these journals, we identified major events that caused substantial oil-price movements on a day-to-day basis, and isolate those events that are arguably exogenous. We, then, construct two measures of oil-price shocks around these days.

We used these measures to estimate the effects of oil-price shocks on output, capacity utilization, and prices for the U.S. economy, and we found that the dynamic responses of these macroeconomic variables implied by our measures are 'well-behaved.' In particular, following an oil-price shock, output and capacity utilization decline, while prices increase. We also found that the decline in output and utilization implied by our measures is larger than the one implied by one conventional measure of oil-price shocks proposed in the literature.

We also presented a dynamic general-equilibrium model and assessed to what extent it can account for the response of the macroeconomic variables to our oil-price shocks. We have found that the model economy we considered does a good job at accounting for the response of prices, but it only has a limited ability in accounting for the timing and the persistence of the estimated responses of output and utilization to our oil-price shock measures. These findings pose a challenge to the empirical plausibility of models in which capacity utilization and output respond directly on impact to unanticipated changes in the price of oil.

Appendix

A.1 Dates

1986, May 2nd — Closure of some Soviet nuclear reactors in wake of Chernobyl disaster

1986, October 30th — Yamani ousted as Saudi Oil Minister;

1988, July 18th — Iran accepts UN calls for cease fire;

1989, January 23rd — Unexpected warm weather in the Northeast;

1989, December 18th — Frigid temperatures in the U.S.;

1989, December 20th — U.S. invasion of Panama;

1990, August 2nd — Iraq invasion of Kuwait; U.S.-led oil boycott;

1990, September to December — Middle East tensions;

1990, January — First Gulf war;

1991, August 19th — Soviet coup;

1996, February 13th — Freezing temperatures in the U.S. northeast and in northern Europe;

1996, February 23rd — Iraq accepted UN resolution 986: exchange of oil for food;

1996, June 17th — UN-Iraq weapons inspection standoff; Many believe that the oil-sale deal may be in jeopardy;

1996, September 3rd — U.S. bombing on southern Iraq;

1996, December 16th — Frigid weather across the U.S.;

1998, January 26th — U.S. comments that patience with Iraq is running out;

1998, September 3rd — Disruption to Russian and Nigerian crude oil supplies; U.S.-Iraq tension on weapon inspection;

1998, December 16th — UN weapons inspectors withdraw from Iraq, a military strike in Iraq may be possible; However, despite the air strike, Iraqi oil continues to flow;

2002, January 2nd — Cold weather in the U.S.;

2002, December 16th — Strikes in Venezuela continue;

2002, December 23rd — Ongoing general strike in Venezuela; Potential war against Iraq;

2003, March — Second Gulf war; U.S. invades Iraq; Traders expected a relatively short war in Iraq with minimal damage to oil installations, but the war looks tougher; British and US military officials say that it will take months before oil from Iraq's southern fields is again ready to be exported; ongoing civil unrest in Nigeria, where approximately 800,000 barrels per day of oil is shut.

2003, July 22nd — Saddam's two sons die at the hands of U.S. troops;

2003, August 1st — Pipeline fire in Iraq, suspected to be caused by sabotage; Heightened concerns about the situation in Iraq;

2003, August 23rd and 24th — Concerns over Tropical Storm Jose and another suspension of Basrah oil loadings in Iraq supported oil prices; New forecasts for a storm (Katrina) hitting the US Gulf Coast and another hefty withdrawal in gasoline stocks pushed crude futures on the New York Mercantile Exchange (NYMEX) to a new record.

A.2 Data

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.

Capacity utilization — Capacity utilization, manufacturing (percent of capacity), Federal Reserve Board, statistical release G.17, Haver Analytics mnemonic: CUMFG@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.

One-year futures price — 1-year Contract Settlement, Wall Street Journal, Haver Analytics mnemonic: PZTEXFY@daily.

References

Aguiar-Conraria, L., Wen, Y., 2007. Understanding the large negative impact of oil shocks. Journal of Money, Credit and Banking 39, 925-944.

Anzuini, A., Pagano, P., Pisani, M., 2007. Oil supply news in a VAR: Information from financial markets. Manuscript, Bank of Italy.

Basu, S., Kimball, M.S., 1997. Cyclical productivity with unobserved input variation. NBER Working Paper no. 5915.

Bernanke, B.S., Gertler, M., Watson, M., 1997. Systematic monetary policy and the effects of oil price shocks. Brookings Papers on Economic Activity, 1:1997, 91-157.

Burnside, C., Eichenbaum, M., Fisher J.D.M., 2004. Fiscal shocks and their consequences, Journal of Economic Theory 115, 89-117.

Chinn, M.D., LeBlanc, M., Coibon, O., 2005. The predictive content of energy futures: An update on petroleum, natural gas, heating oil and gasoline. NBER Working Paper No. 11033.

Chow, G.C. and A. Lin, 1971. Best linear unbiased interpolation, distribution and extrapolation of time series by related series. Review of Economics and Statistics 53, 372-375.

Christiano, L.J., Eichenbaum, M., Evans, C.L., 1999. Monetary policy shocks: What have we learned and to what end? In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, Vol. 1A. North-Holland, Amsterdam, pp. 65-148.

Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1-45.

Cochrane, J.H., Piazzesi, M., 2002. The Fed and interest rates: A high-frequency identification. American Economic Review 92, 90-95.

Cochrane, J.H., Piazzesi, M., 2005. Bond risk premia. American Economic Review 95, 138-160.

Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal monetary policy with staggered wage and price contracts. Journal of Monetary Economics 46, 281-313.

Faust, J., Swanson, E.T., Wright, J.H., 2004. Identifying VARs on high frequency futures data. Journal of Monetary Economics 51, 1107-1131.

Finn, M.G., 1995. Variance properties of Solow's productivity residual and their cyclical implications. Journal of Economic Dynamics and Control 19, 1249-1281.

Finn, M.G., 2000. Perfect competition and the effects of energy price increases on economic activity. Journal of Money, Credit and Banking 32, 400-416.

Greenwood, J., Hercowitz, Z., Huffman, G., 1988. Investment, capacity utilization and the real business cycle. American Economic Review 78, 402-417.

Hamilton, J.D., 1985. Historical causes of postwar oil shocks and recessions. Energy Journal 6, 97-116.

Hamilton, J.D., 1996. This is what happened to the oil price-macroeconomy relationship. Journal of Monetary Economics 38, 215-220.

Hamilton, J.D., 2003. What is an oil shock? Journal of Econometrics 113, 363-398.

Hoover, K.D., Perez, S.J., 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.

Jaimovich, N., Rebelo, S., 2006. Can news about the future drive the business cycle? NBER Working Paper No. 12537.

Kilian, L., 2006. Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? Review of Economics and Statistics, forthcoming.

Leduc, S., Sill, K., 2004. A quantitative analysis of oil-price shocks, systematic monetary policy and economic downturns. Journal of Monetary Economics 51, 781-808.

Levin, A.T., Onatski A., Williams, J.C., Williams, N., 2006. Monetary policy under uncertainty in micro-founded macroeconometric models. In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2005. MIT Press, Cambridge, pp. 229-287.

Ramey, V.A., Shapiro, M., 1998. Costly capital reallocation and the effects of government spending. Carnegie-Rochester Conference Series on Public Policy 48, 145-194.

Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: Derivation and implications. American Economic Review 94, 1055-1084.

Romer, C.D., Romer, D.H., 2006. The macroeconomic effect of tax changes: Estimates based on a new measure of fiscal shocks. Manuscript, University of California, Berkeley.

Wu, T., McCallum, A., 2005. Do oil futures prices help predict future oil prices? Federal Reserve Bank of San Francisco Economic Letter no. 2005-38.

Figure 1: Oil-price shock measures.

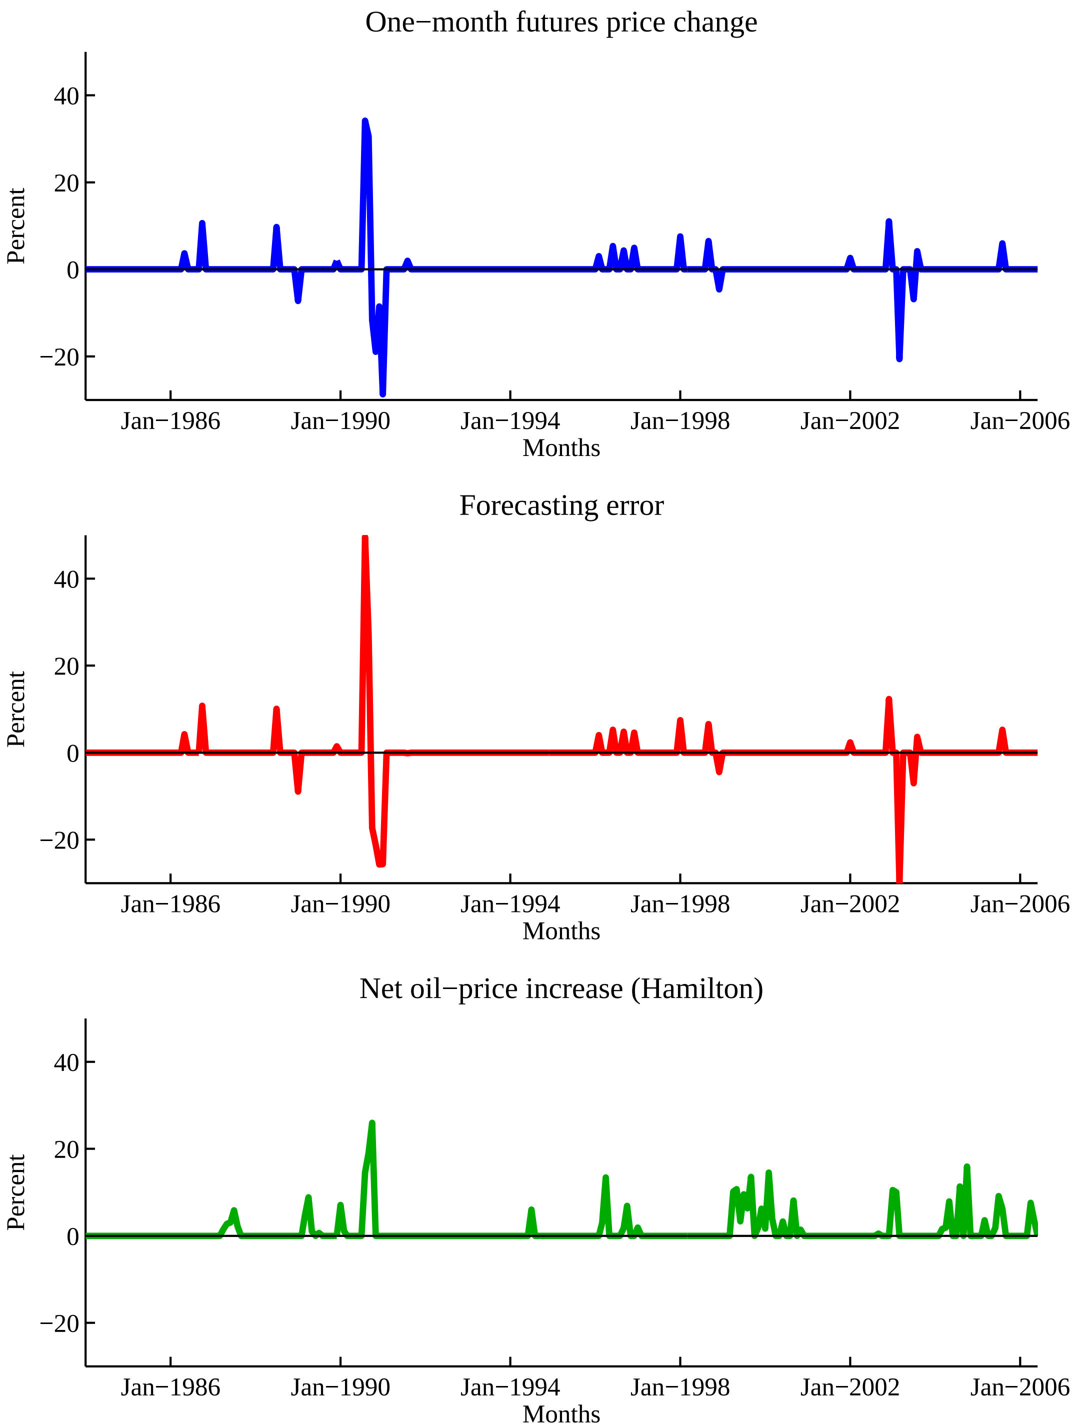

Figure 2: Dynamic responses to an oil-price shock estimated using the univariate autoregressive model in equation (3).

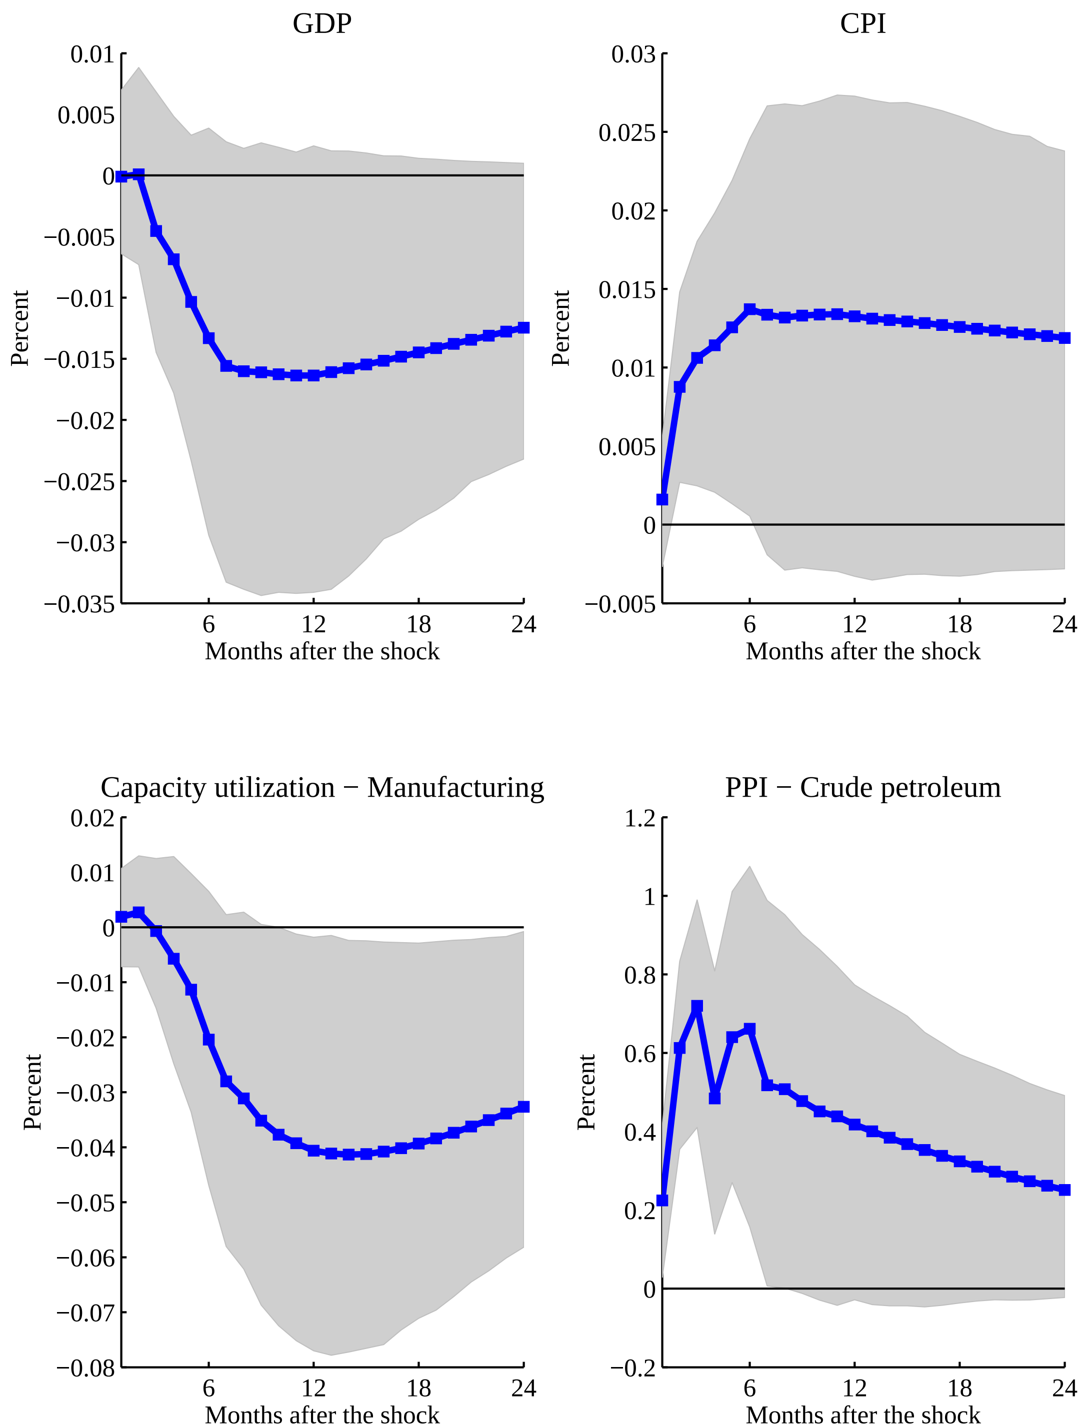

Note: Solid lines with square markers denote point estimates of the dynamic responses to the 'one-month futures price change' measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 3: Dynamic responses to an oil-price shock estimated using the univariate autoregressive model in equation (3).

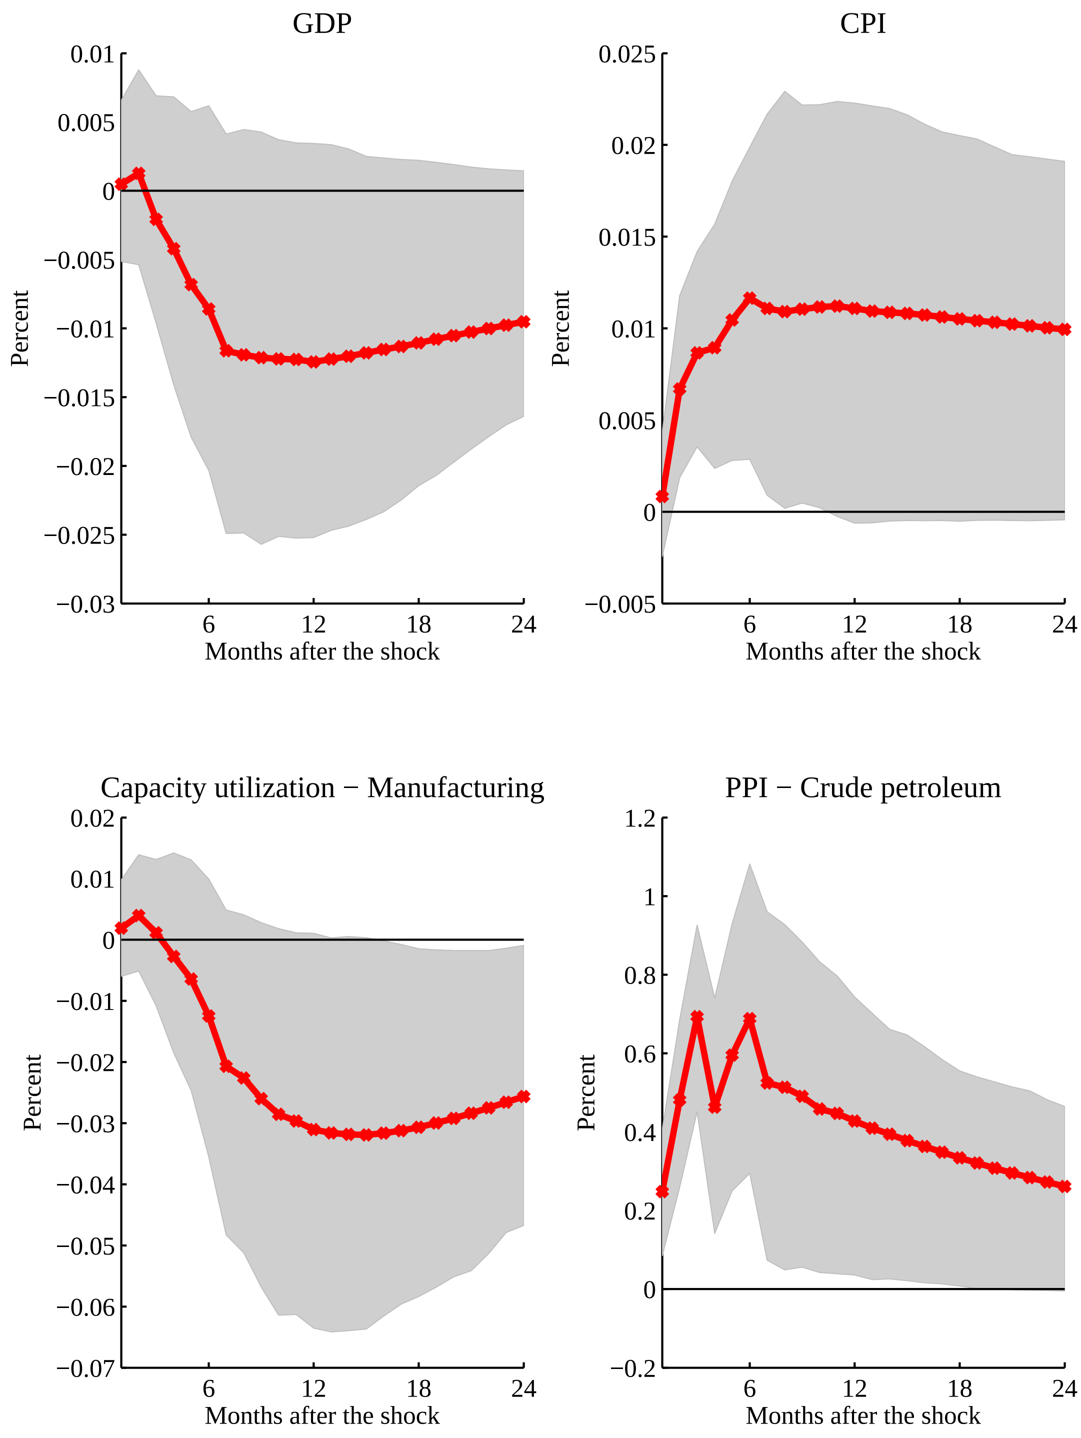

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 4: Dynamic responses to an oil-price shock estimated using the univariate autoregressive model in equation (3).

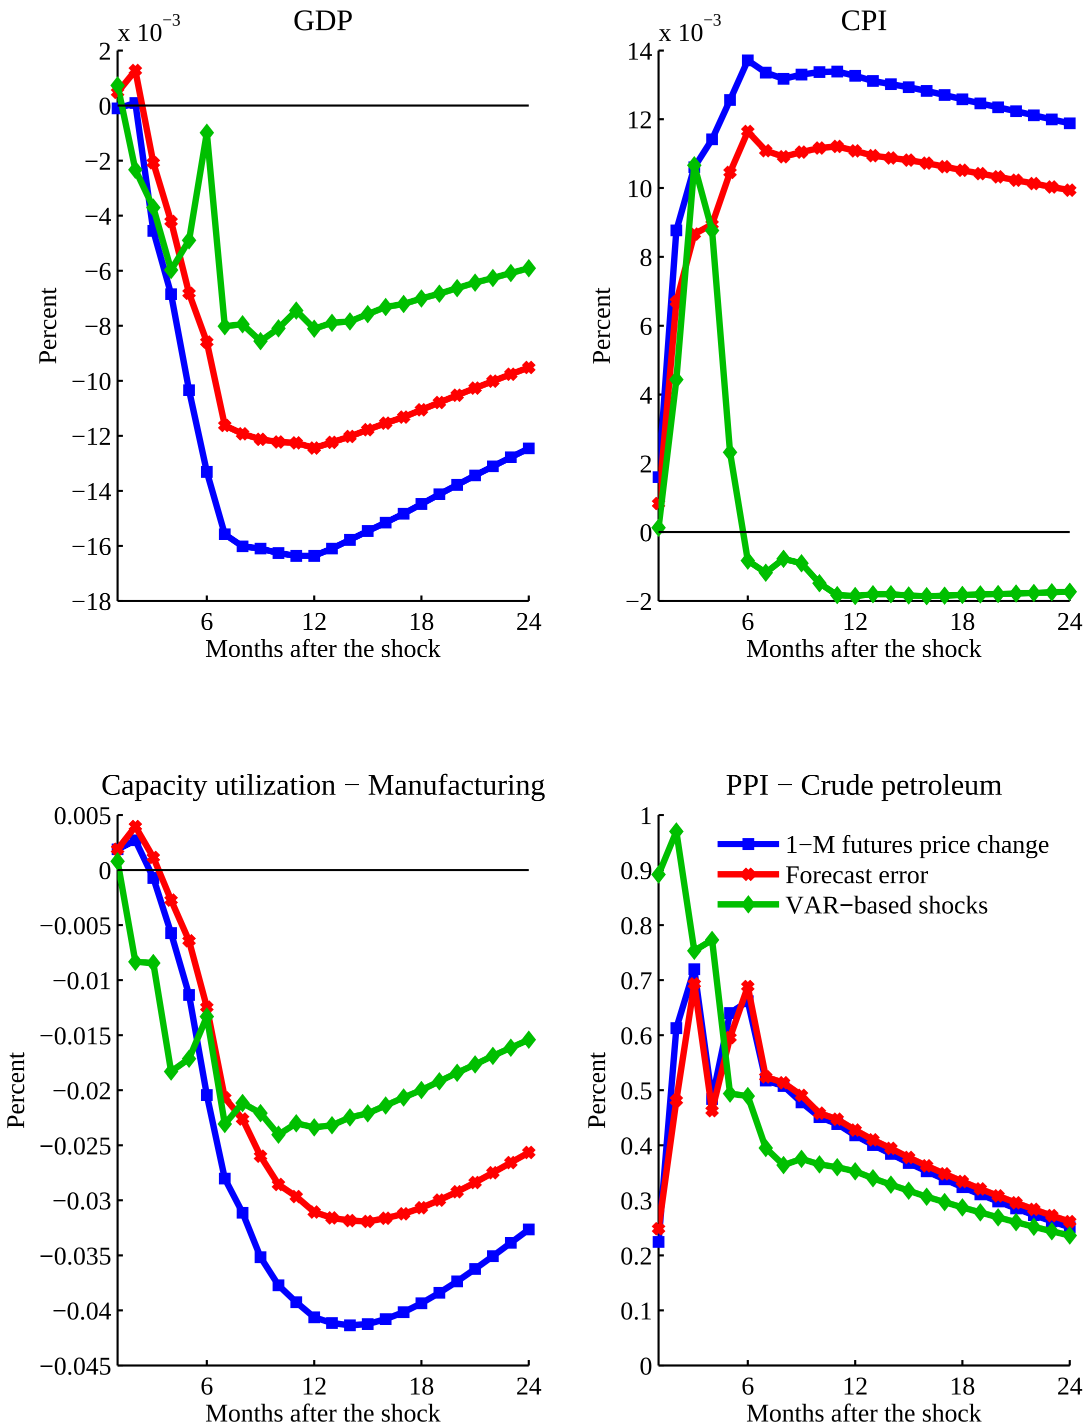

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'one-month futures price change' measure of oil-price shocks. Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Solid lines with diamond markers denote point estimates of the dynamic responses to the VAR-based measure of oil-price shocks.

Figure 5: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (4).

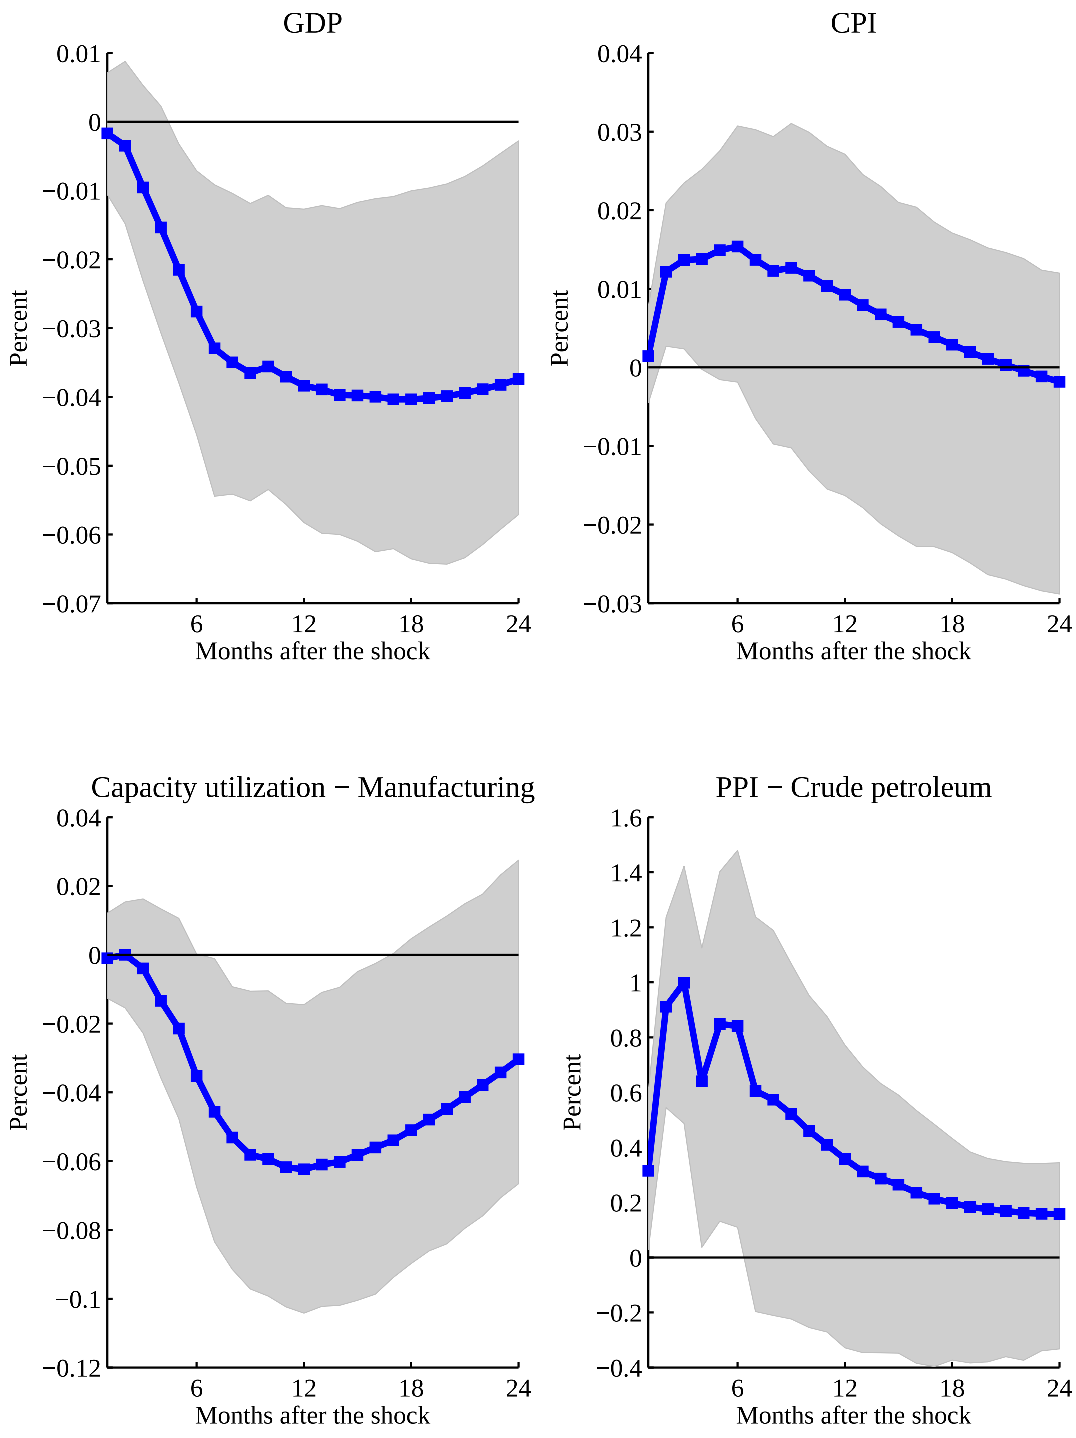

Note: Solid lines with square markers denote point estimates of the dynamic responses to the ‘one-month futures price change’ measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 6: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (4).

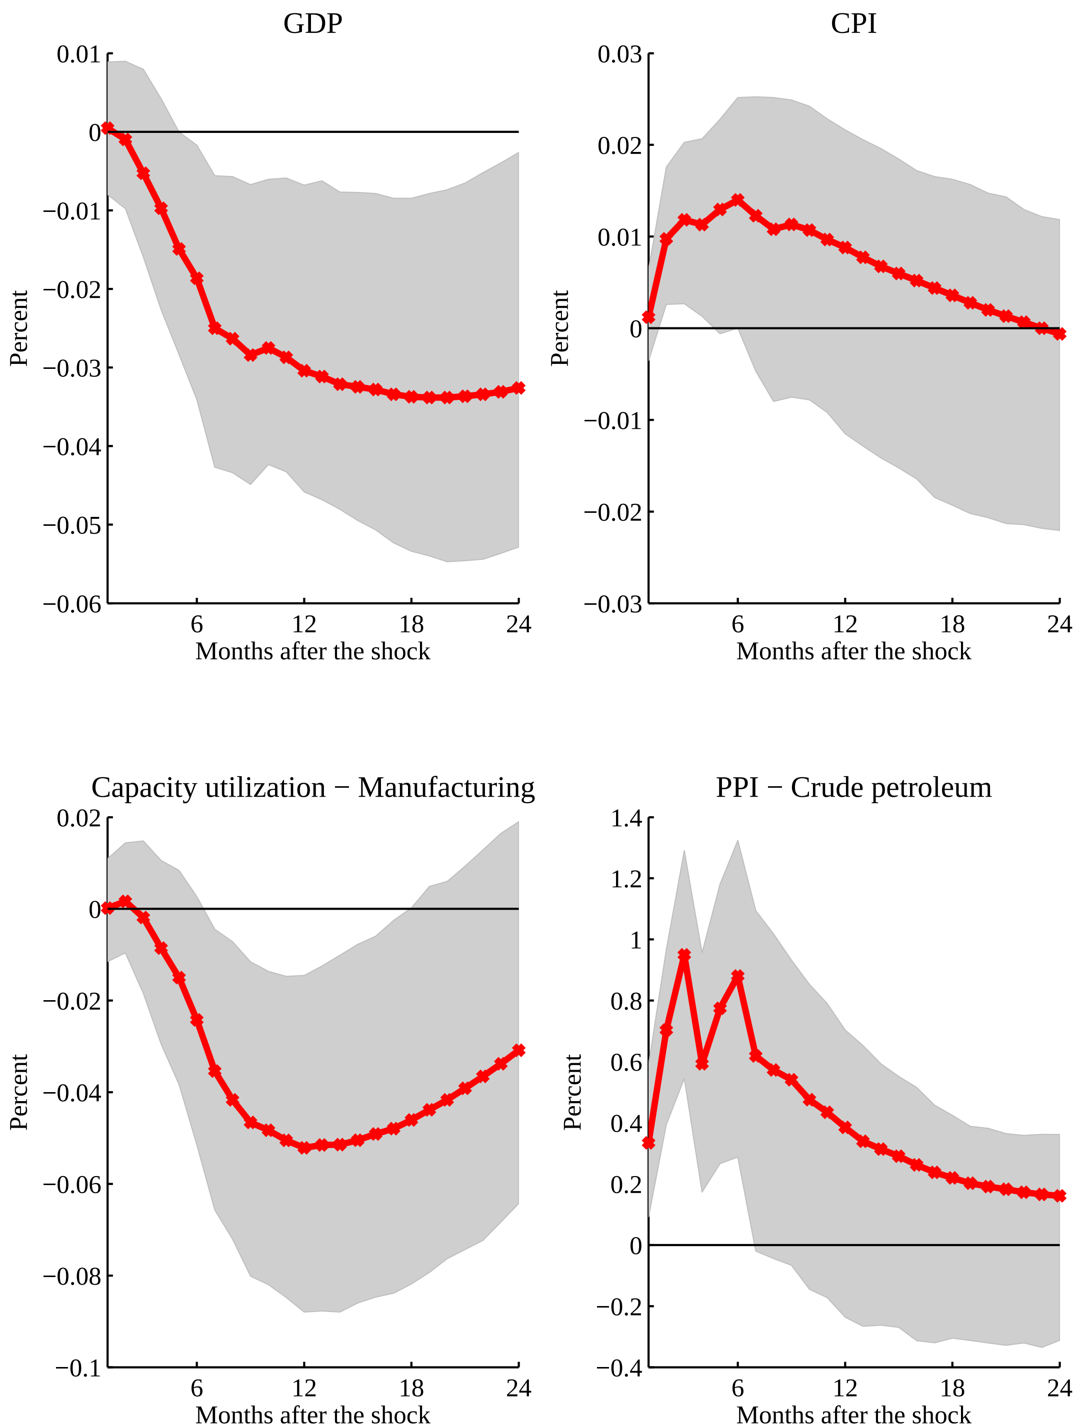

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 7: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (4).

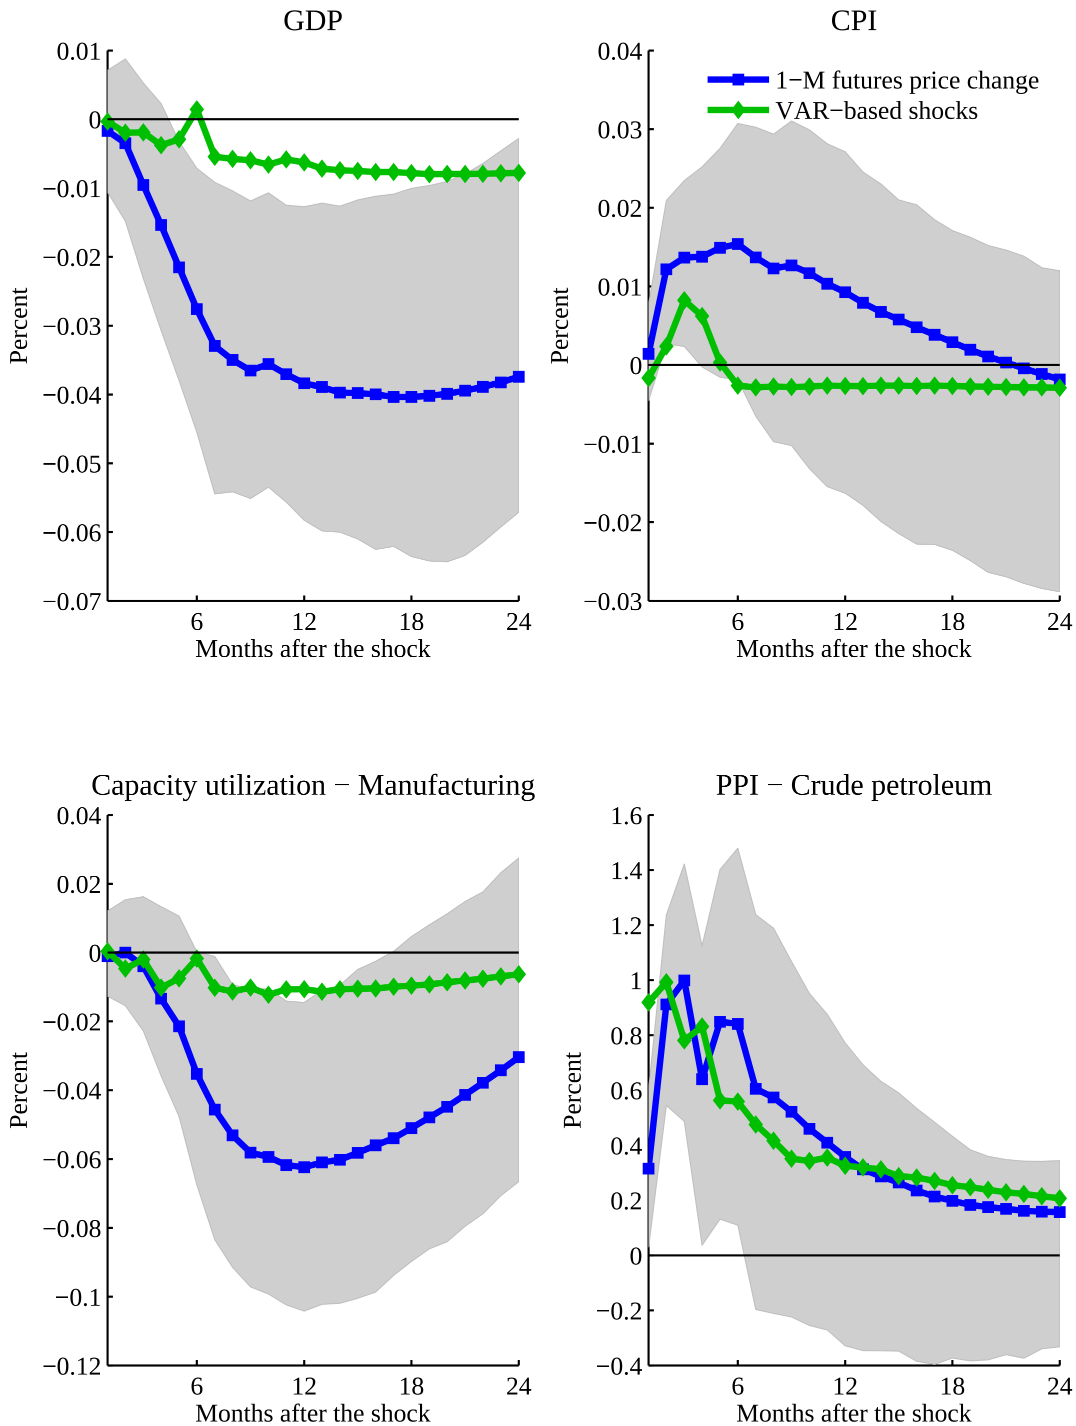

Note: Solid lines with square markers denote point estimates of the dynamic responses to the 'one-month futures price change' measure of oil-price shocks. Solid lines with diamond markers denote point estimates of the dynamic responses to the VAR-based measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 8: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (4).

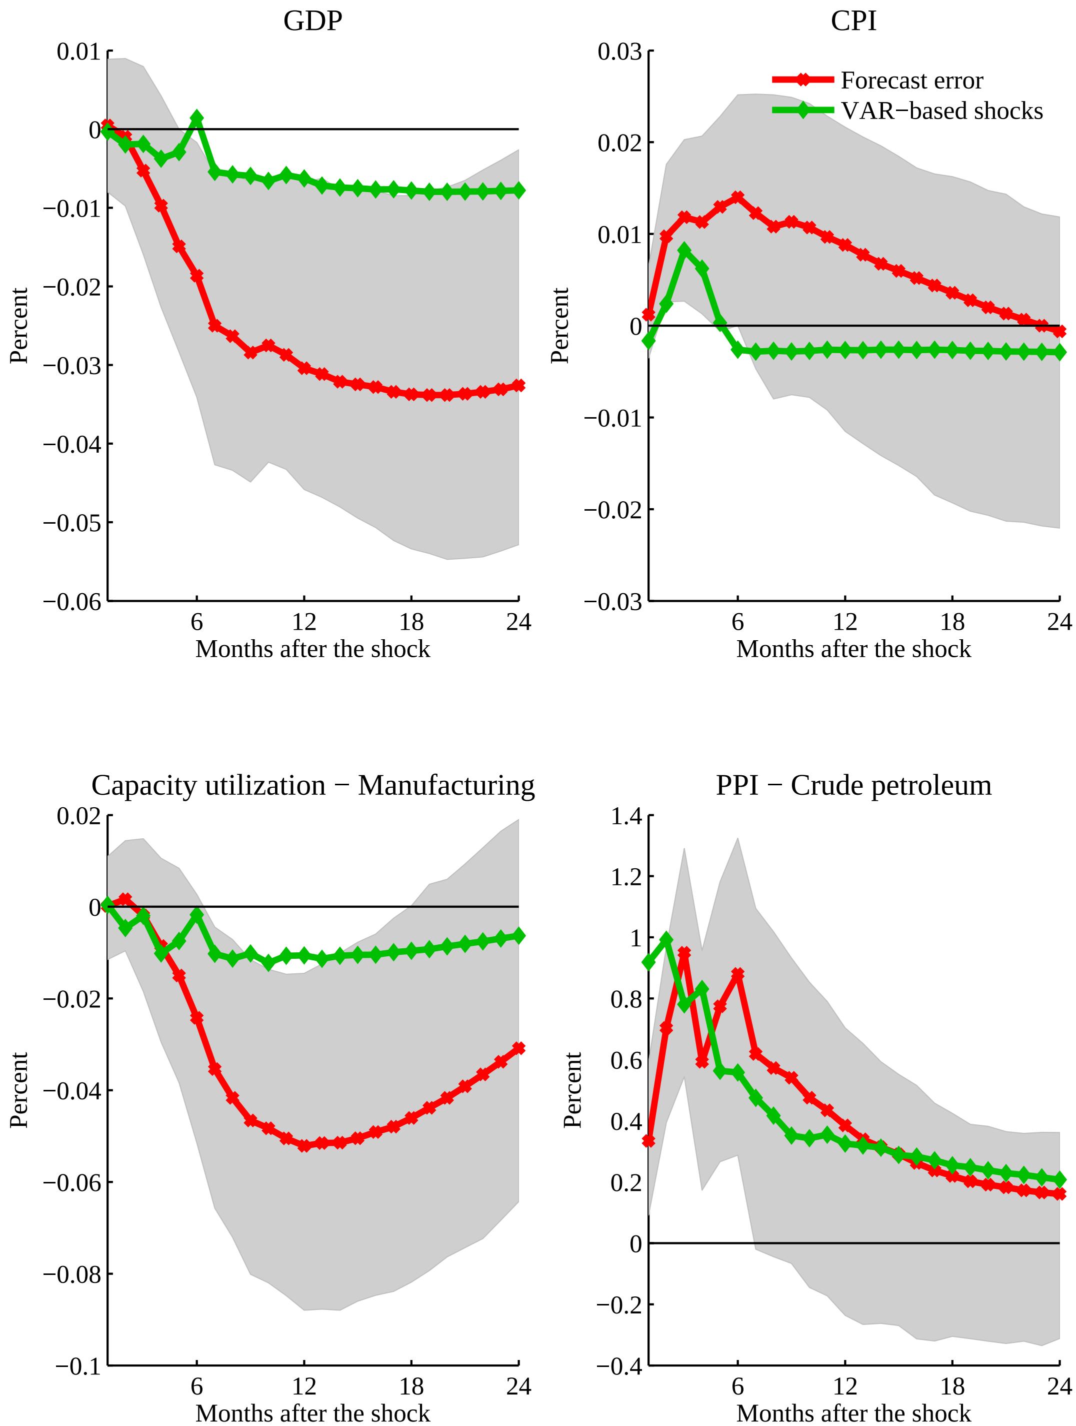

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Solid lines with diamond markers denote point estimates of the dynamic responses to the VAR-based measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 9: Simulated dynamic responses to the ‘one-month futures price change’ measure of oil-price shocks

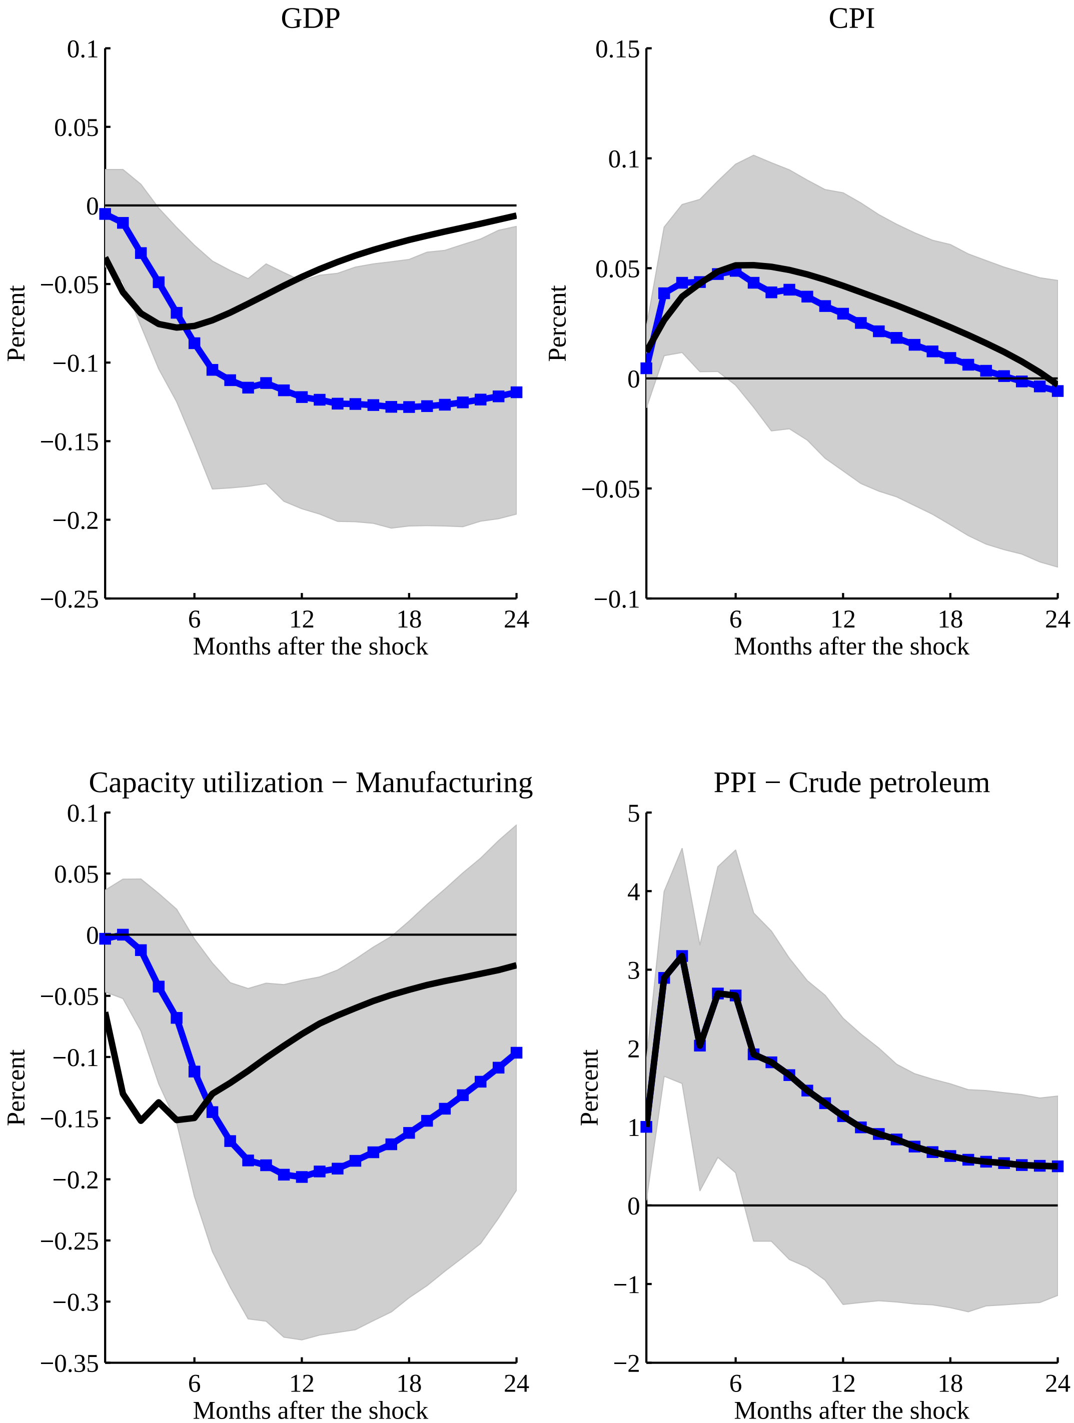

Note: Solid lines without markers denote the simulated dynamic responses to the ‘one-month futures price change’ measure of oil-price shocks. Solid lines with cross markers denote point estimates of the dynamic responses to the ‘one-month futures price change’ measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 10: Simulated dynamic responses to the 'forecasting error' measure of oil-price shocks

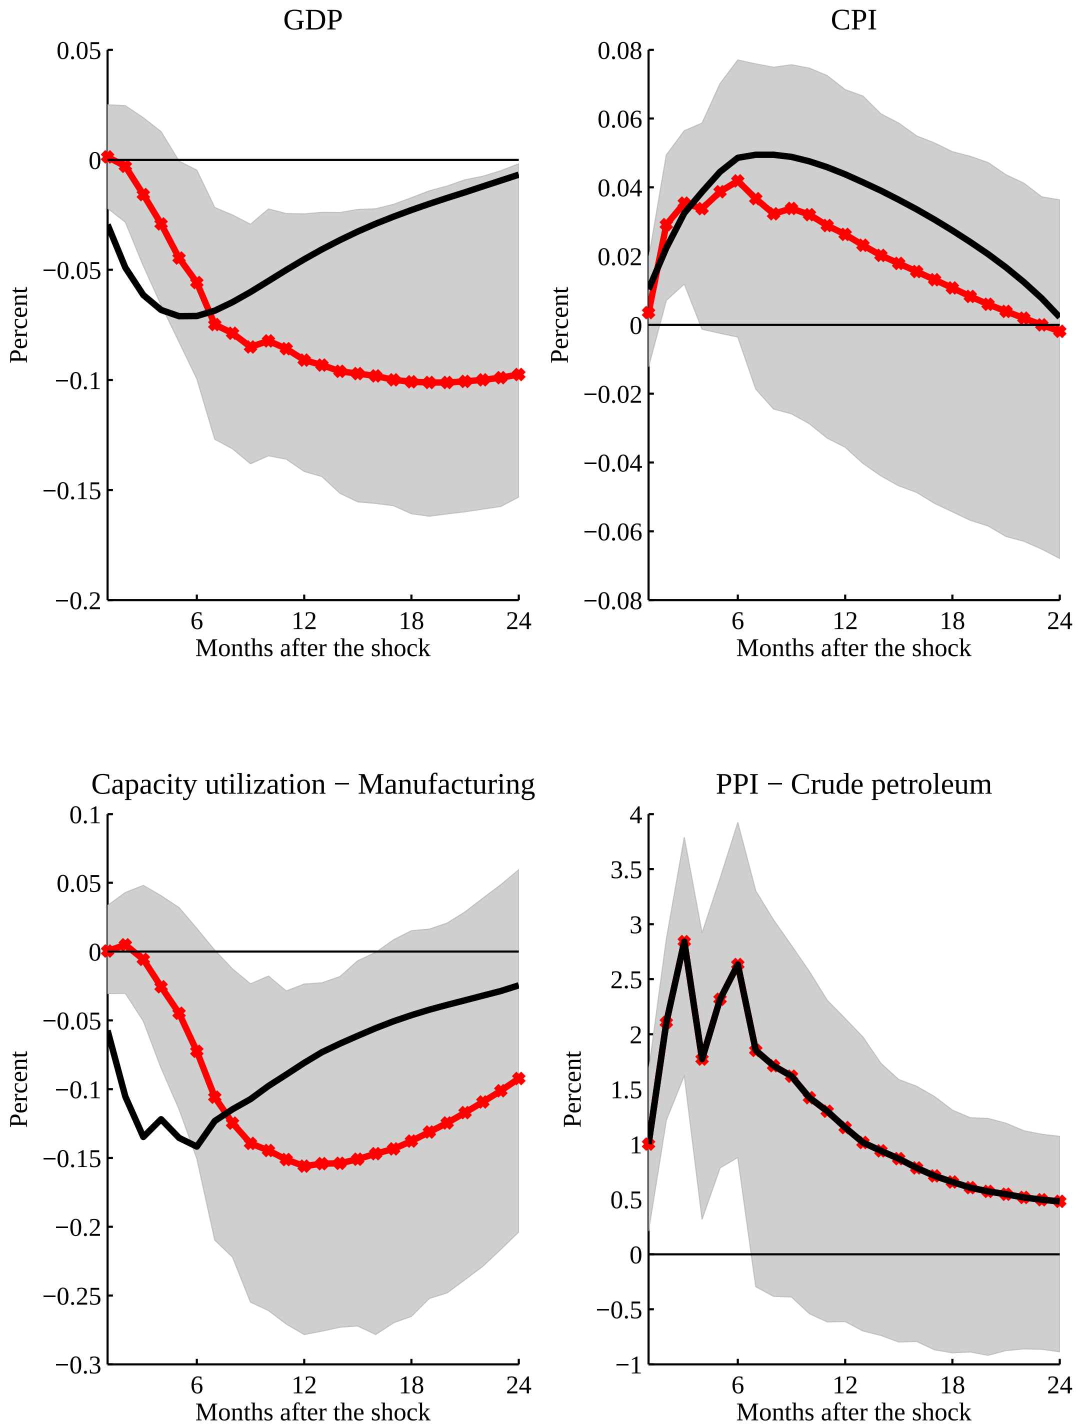

Note: Solid lines without markers denote the simulated dynamic responses to the 'forecasting error' measure of oil-price shocks. Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Shaded areas denote a 95-percent confidence interval for the point estimates of the dynamic responses.