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

Tao Wu
Federal Reserve Bank of Dallas

September 2006

Working Paper 2006-28
http://www.frbsf.org/publications/economics/papers/2006/wp06-28bk.pdf

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." We also find that the response of output is larger than the one implied by a conventional measure of oil-price shocks proposed in the literature.

We thank Òscar Jordà and Eric Swanson for helpful conversations and suggestions, 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 forecastability. 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). 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. This underscores the necessity of using market-based information to obtain a better measure of the exogenous oil shocks.

In this paper, we develop two measures of exogenous oil-price shocks for the period 1984 to 2006. 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 oil futures prices 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. We estimate the effects of oil-price shocks on output and prices for the U.S. economy, by including these measures, one at a time, as exogenous variables in our regression analysis. We find that the dynamic responses of output and prices 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 declines 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.

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.

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 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 define our first measure of oil-price shocks 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. In most cases, this measure 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 'absolute 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:

$$(1) \qquad P_{t+i} - P_t = \alpha^O + \sum_{j=1}^{6} \beta_j^O (P_t^j - P_t) + \varepsilon_{t+i},$$

where P_t and P_{t+i} are the spot oil price at t and $t+i$, respectively, P_t^j 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 $P_{t+i} - E_t P_{t+i}$ and we call it the 'forecasting error.' Equation (1) incorporates term structure information on futures-spot spread in forecasting future oil price. 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.

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 and oil prices 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 and oil prices for the U.S. economy. We use two procedures. First, we use a univariate autoregressive model for each of the three 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 conventional 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, and current and lagged values of prices. Identifying oil-price shocks with these residuals is equivalent to estimating a VAR including output, prices 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.

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:

$$(2) \qquad x_t = a_0 + a_1 t + \sum_{i=1}^{6} \alpha_i x_{t-i} + \sum_{i=0}^{6} \beta_i O_{t-i} + \epsilon_t,$$

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 (2) for each of the endogenous variables we are studying. Therefore, the variable x_t corresponds to output, prices and producer crude oil prices in the U.S. economy. In addition, for each one of these variables, we estimate equation (2) 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 'absolute 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 (2) 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) to estimate the dynamic responses of key macroeconomic variables to a fiscal policy shock.

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:

$$(3) \qquad X_t = A_0 + A_1 t + A_2 (L) X_{t-1} + B (L) O_t + \varepsilon_t,$$

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 (3), 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 (2), we also estimate regression (3) 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 $[I - A(L)L]^{-1} B(L)$.

The left-hand side variables that we use in (2) and that we also include in the vector X_t in (3) are the log of real GDP, the log of the producer price index for finished goods and the log of the producer price index for crude petroleum. We estimate equations (2) and (3) 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 capacity utilization.[1] Appendix A.2 describes the data used in our analysis.

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

4 Empirical results

This section presents the results we obtained following the estimation procedures outlined in the previous section. In Figures 1 through 6 we plot the estimated dynamic responses of output, prices and oil prices to a one-percent increase in the average value of the oil-price shock variable that we consider. Figures 1 through 3 illustrate the dynamic responses obtained using the estimates of the coefficients from the univariate autoregressive model in equation (2). Figures 4 through 6 illustrate the dynamic responses obtained using the estimates of the coefficients from the vector autoregressive model in equation (3). In all the figures, the top panel shows the response of output, the middle panel shows the response of prices, while the bottom panel shows the response of the price of oil. Figures 3 and 6 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 no markers to denote the point estimates of the responses to the 'absolute 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 use dashed lines to delimit 95-percent confidence intervals for the point-estimates of the dynamic responses.[2]

[2]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 (2) and (3), respectively. 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 (2) and (3) 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 confidence intervals in Figures 1 through 6 correspond to a graph of these coefficients.

Our results show that the dynamic responses of output and prices 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 declines 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 1 illustrates the dynamic responses of output, prices, and the producer price of oil obtained estimating the univariate autoregressive model of equation (1) and using the 'absolute 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, and that, after about ten months, it becomes statistically significant. After 24 months, the sum of the point estimates of the output response is equal to -0.445 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 1.255 percent. As one would normally expect, the oil-price shock leads to an increase in the producer price of oil. This increase displays a hump-shaped pattern. The point estimate is statistically significant for about 10 months after the shock has occurred. Relative to its preshock level, the producer price of oil reaches a peak of about 1.07 percent after 3 months, and it slowly reverts to zero afterwards.

Figure 2 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 an oil-price shock as measured by our 'forecasting error' variable are broadly similar to those induced when we measure the shock with our 'absolute price change' variable. Specifically, output declines, while prices and the producer price of oil increase. However, the response of output in this case is never statistically significant, and the cumulative sum of the dynamic response estimates is -0.314 percent. The price level increases following the oil-price shock, and, as in the previous case shown in Figure 1, its response becomes statistically significant two months after the shock. The sum of the point estimates of the response of prices is equal to 1.046 percent. As we obtained before, the producer 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 also happens after three months and it is equal to 0.981 percent relative to the preshock level.

Figure 3 compares the estimated dynamic responses from the autoregressive model obtained using our 'absolute 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 'absolute 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 -0.352 percent, compared with a cumulative decline of -0.445 implied by the 'absolute price change' measure. This implies that the cumulative decline in output obtained with the 'absolute price change' measure is almost 30 percent larger than the cumulative decline obtained with the VAR-based measure of oil price shocks. The top panel of Figure 3 also shows that decline in output implied by the 'forecasting error' measure is comparable in magnitude with the decline implied by the VAR-based measure. In regard to the response of prices, our measures of oil-price shocks imply a cumulative increase 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 equal to 0.797 percent, lower than the 1.255 percent and 1.046 percent cumulative increases obtained with the 'absolute price change' and the 'forecasting error' shock measures, respectively. Finally, the estimated dynamic response implied by the VAR-based shock measure indicates that, in the months immediately after the shock, the producer price of oil increases by an amount larger than the one implied by our shock measures. However, the cumulative increases in the price of oil over the 24-month horizon implied by our two shock measures are larger than the one implied by the VAR-based measure. While, the VAR-based shock measure implies that the cumulative increase in the producer price of oil is equal to 11.98 percent, our two measures imply a cumulative increase of 14.38 percent, and 13.78 percent, respectively.

After having examined the dynamic responses of output, prices and the price of oil estimated using the univariate autoregressive model of equation (2), we turn to analyze the responses estimated using the vector autoregressive model of equation (3). The results from the vector autoregressive model can be broadly summarized as follows. The point estimates of the dynamic responses of output obtained using the vector autoregressive model are somehow larger than the point estimates of the dynamic responses of output obtained using the univariate autoregressive model. The opposite pattern occurs in regard to the estimated dynamic response of prices and of the producer price of oil: the point estimates of the dynamic responses obtained using the vector autoregressive model are somehow lower than the point estimates of the dynamic responses obtained using the univariate autoregressive model.

Figure 4 plots the dynamic responses estimated using the 'absolute price change' shock measure. As obtained estimating the univariate autoregressive model, the figure indicates that output starts declining two months after the date the shock occurs and deteriorates for approximately other additional ten months. The point estimate of the response of output becomes significantly different from zero six months after the shock. In addition, the estimated cumulative decline in output after 24 months relative to the preshock level is equal to -0.571 percent, compared with -0.445 percent estimated using the univariate autoregressive model. The estimated dynamic response of prices is milder than the one obtained with our autoregressive model. 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 0.932 percent compared with the 1.255 percent implied by the autoregressive 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 1.06 percent after 3 months, and it slowly reverts to zero afterwards.

Figure 5 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 the output starts declining two months after the shock, and that it keeps declining for other additional ten months. The point estimate of the response of output is never statistically different from zero over the 24-month horizon, and the cumulative decline in output is equal to -0.414 percent compared with -0.314 percent obtained estimating the univariate autoregressive model. The estimated response of prices is milder relative to the estimated response obtained with the univariate autoregressive model. In fact, the sum of the point estimates of the response of prices is 0.822 percent compared with 1.046 percent implied by the estimates of the univariate model. The 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 9 months after the shock occurs. The peak response happens three months after date of the shock and it is equal to 0.987 percent relative to its preshock level.

Finally, Figure 6 compares 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 panel shows 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 'absolute price change' and the 'forecasting error' measures. Specifically, the cumulative decline in output implied by the VAR-based shock measure is -0.274 percent, compared to -0.571 percent and -0.414 percent implied by the 'absolute price change' and the 'forecasting error' measures, respectively. This implies, in particular, that the cumulative decline in output implied by the 'absolute price change' measure is more than 100 percent larger than the cumulative output decline implied by the VAR-based shock measure. In regard to the estimated dynamic responses of prices, the VAR-based shock measure implies a cumulative increase in prices of a magnitude which is quite comparable the one implied by our two shock measures. The VAR-based measure implies a 0.859 percent cumulative increase in prices, while our two measures imply a cumulative increase of 0.932 percent and 0.822 percent, respectively. Finally, the VAR-based oil-price shock measure implies a somehow larger increase in the producer price of oil in the first few moths immediately after the shock. As a result, the estimated cumulative increase in the producer price implied by the VAR-based measure is slightly larger than that implied by our two shock measures. While, the VAR-based measure implies that the estimated cumulative increase is equal to 14.53 percent relative to its preshock level, out two shock measures imply an estimated cumulative increase of 13.2 percent and 13.05 percent, respectively, relative to the preshock level.

Summing up, our results indicate that our measures of oil-price shocks imply estimated dynamic responses of output, prices and the producer price of oil that have the expected sign. In particular, we find that output declines and that the price level increases. We also find that the decline in output implied by our 'absolute price change' measure is larger than the one implied by a conventional VAR-based measure of oil-price shocks. In particular, the univariate autoregressive model suggests that the estimated cumulative decline in output implied by our measure is about 30 percent larger relative to the one implied by a VAR-based measure. When we estimate a vector autoregressive model, the comparison is more striking, as the estimated cumulative decline in output implied by our measure is more than 100 percent larger than the one implied by a VAR-based measure.

5 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 and prices for the U.S. economy, and we found that the dynamic responses of output and prices implied by these measures are 'well-behaved.' In particular, following an oil-price shock, output declines and prices increase. We also found that the decline in output implied by one of our measures is larger than the one implied by one conventional measure of oil-price shocks proposed in the literature.

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.

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

Finished Goods PPI — Producer price index, finished goods (1982=100), Bureau of Labor Statistics, series ID: WPSSOP3000;

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

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.

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.

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.

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.

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

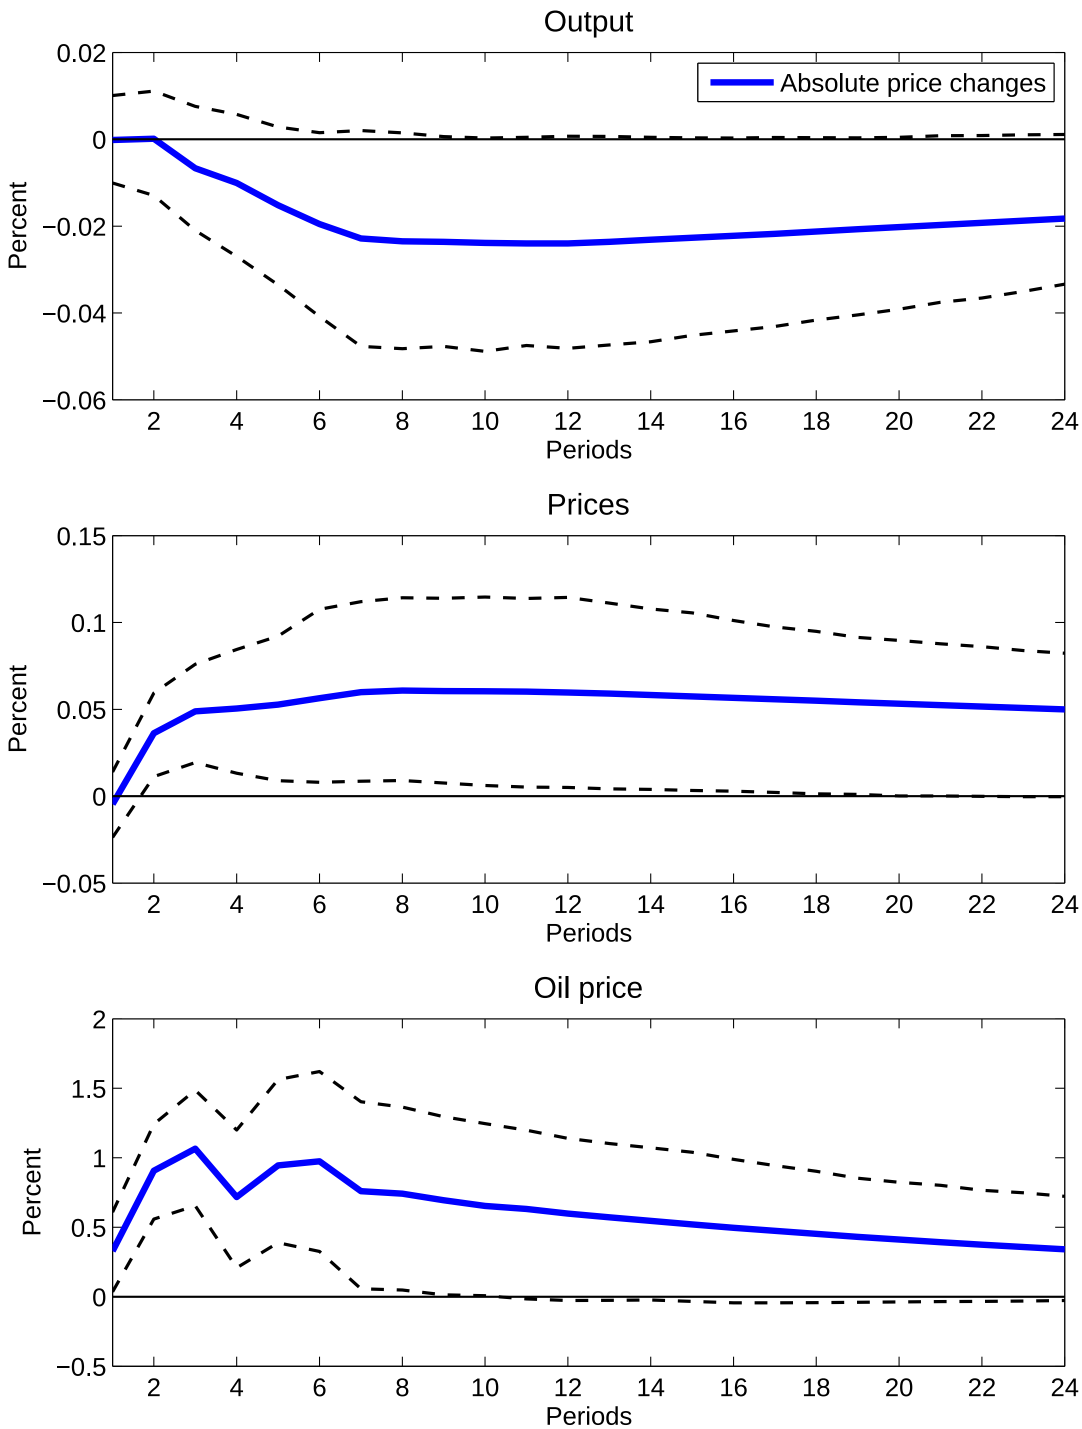

Note: Solid lines denote point estimates of the dynamic responses to the 'absolute price change' measure of oil-price shocks. Dashed lines delimit a 95-percent confidence interval for the point estimates of the dynamic responses.

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

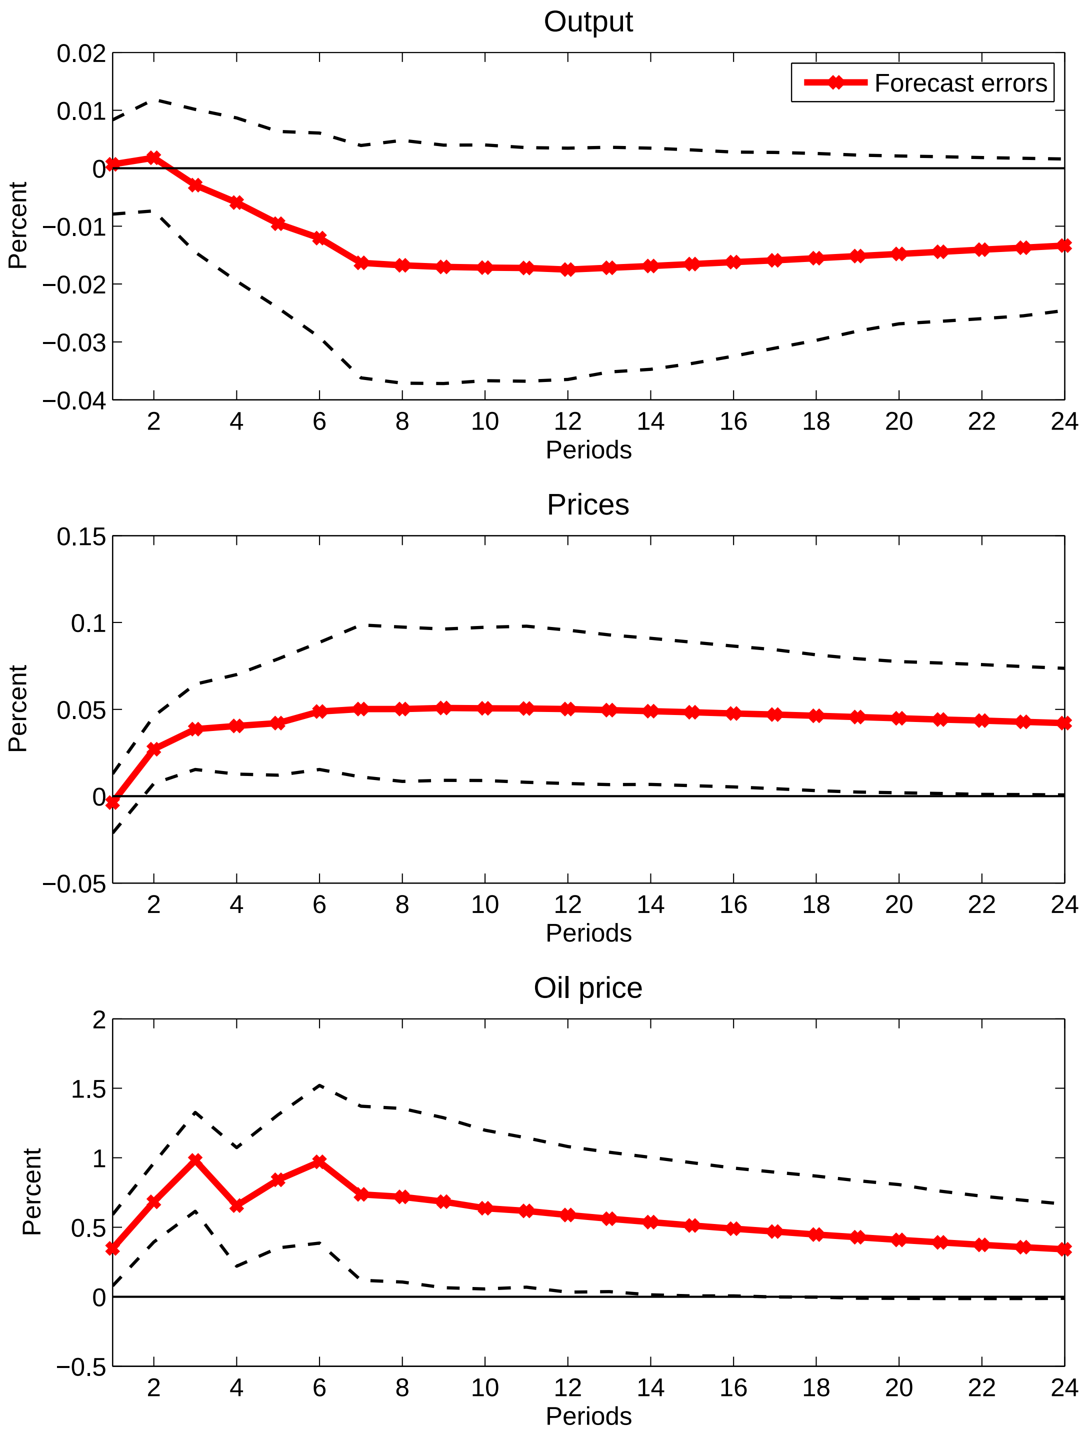

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Dashed lines delimit a 95-percent confidence interval for the point estimates of the dynamic responses.

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

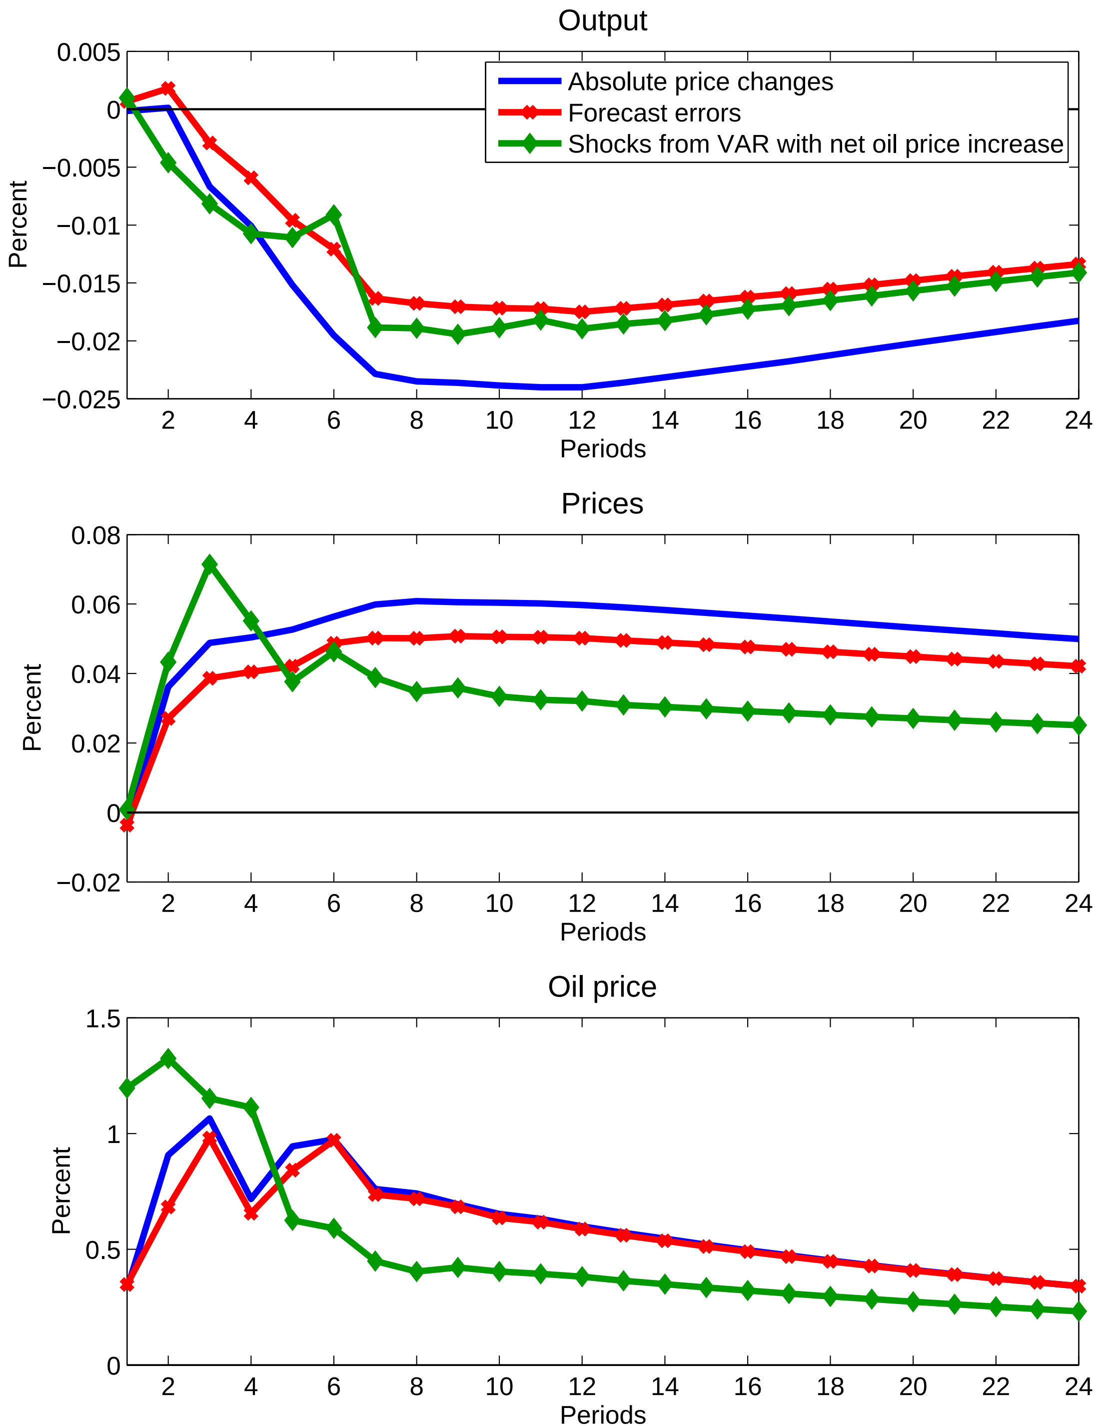

Note: Solid lines denote point estimates of the dynamic responses to the 'absolute price change' measure of oil-price shocks. Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Solid lines with diamond markers denote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

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

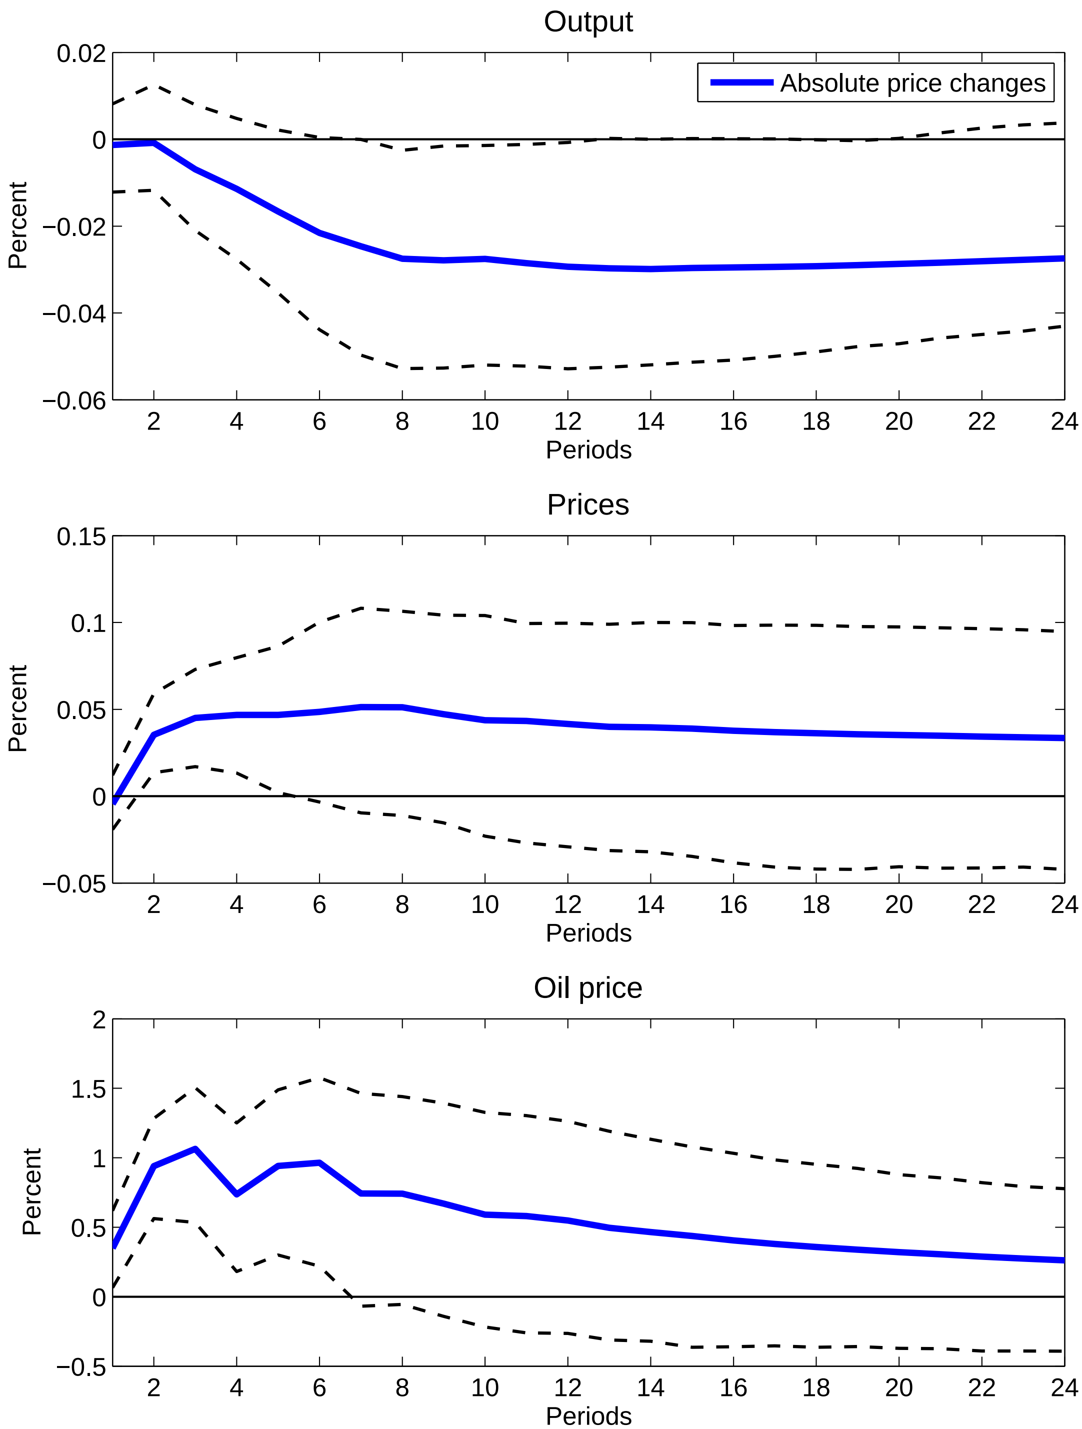

Note: Solid lines denote point estimates of the dynamic responses to the 'absolute price change' measure of oil-price shocks. Dashed lines delimit a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 5: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (3).

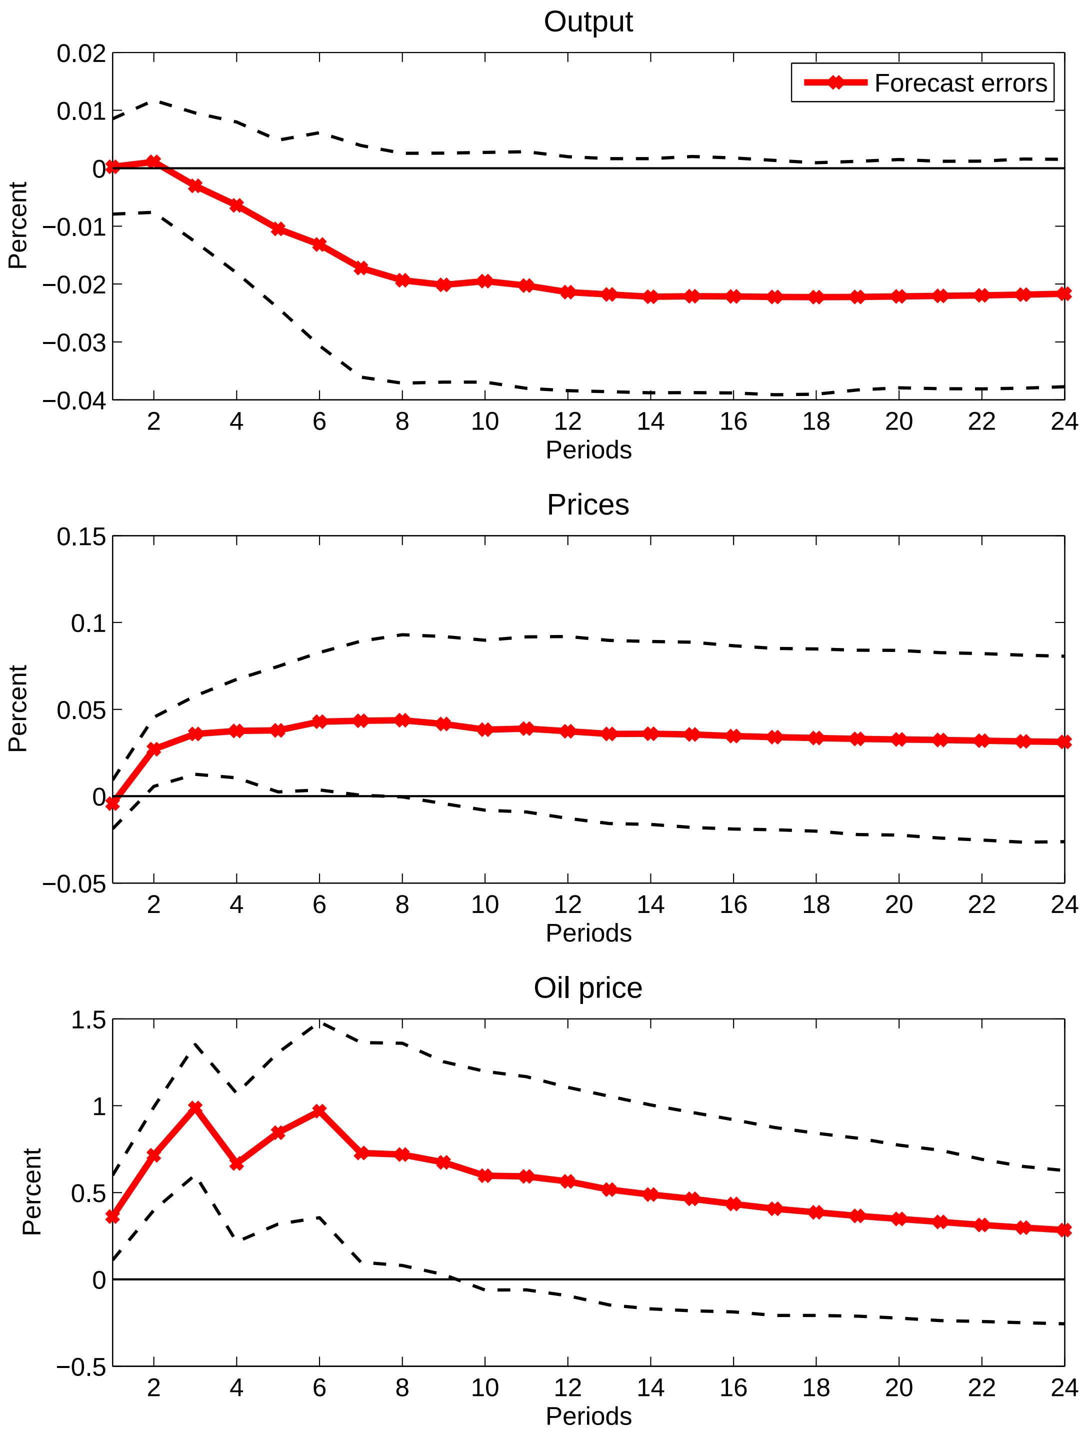

Note: Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Dashed lines delimit a 95-percent confidence interval for the point estimates of the dynamic responses.

Figure 6: Dynamic responses to an oil-price shock estimated using the vector autoregressive model in equation (3).

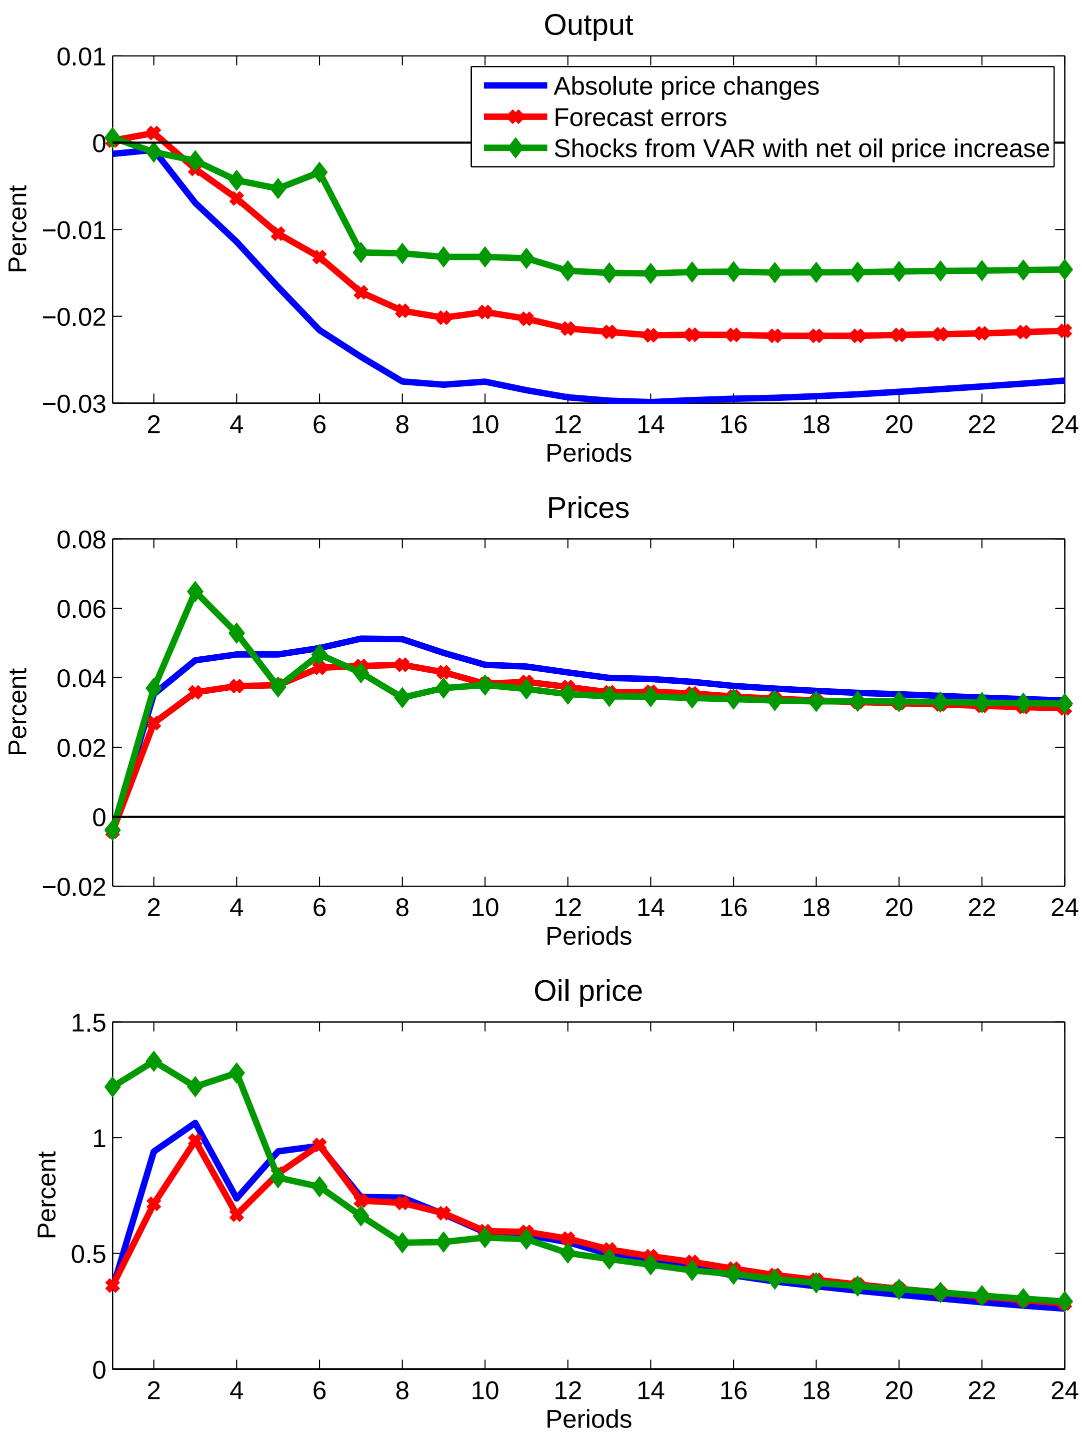

Note: Solid lines denote point estimates of the dynamic responses to the 'absolute price change' measure of oil-price shocks. Solid lines with cross markers denote point estimates of the dynamic responses to the 'forecasting error' measure of oil-price shocks. Solid lines with diamond markers denote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