Macroeconomic Dynamics, **15** (Supplement 3), 2011, 379–395. Printed in the United States of America.
doi:10.1017/S1365100511000630

VOLATILITY IN OIL PRICES AND MANUFACTURING ACTIVITY: AN INVESTIGATION OF REAL OPTIONS

JOHN ELDER
Colorado State University

APOSTOLOS SERLETIS
University of Calgary

Previous research shows that volatility in oil prices has tended to depress output, as measured by nonresidential investment and GDP. This is interpreted as evidence in support of the theory of real options in capital budgeting decisions, which predicts that uncertainty about, for example, commodity prices will cause firms to delay production and investment. We continue that investigation by analyzing the effect of oil price uncertainty on monthly measures of U.S. firm production related to industries in mining, manufacturing, and utilities. We use a more general specification, an updated sample that includes the increased oil price volatility since 2008, and we control for other nonlinear measures of oil prices. We find additional empirical evidence in support of the predictions of real options theory, and our results indicate that the extreme volatility in oil prices observed in 2008 and 2009 contributed to the severity of the decline in manufacturing activity.

Keywords: Oil Volatility, Uncertainty, Multivariate GARCH VAR, Real Options

1. INTRODUCTION

Elder and Serletis (2010) find that volatility in oil prices has tended to depress some components of aggregate investment. The motivation for their analysis is based on "real option" models, also known as investment under uncertainty. That is, real option theory predicts that some firm expenditures—expenditures that have uncertain future return, that are costly to reverse, and for which there is flexibility in timing—may tend to be delayed or abandoned as uncertainty about that reward increases. See, for example, Bernanke (1983), Brennan and Schwartz (1985), Madj and Pindyck (1987), Brennan (1990), Gibson and Schwartz (1990), and Dixit and Pindyck (1994).

Many firm expenditures fall in this category, in addition to capital budgeting decisions regarding fixed investment in large manufacturing facilities, such as investment in an automobile plant. For example, firm decisions about the level

The authors thank Ana Maria Herrera and Kiseok Lee for comments that greatly improved the paper. Address correspondence to: John Elder, Department of Finance and Real Estate, Colorado State University, 1272 Campus Delivery, Fort Collins, CO 80523-1272, USA; e-mail: john.elder@colostate.edu.

 1365-1005/11

of production may involve nonrecoverable costs associated with the hiring and training of labor. These decisions may also involve expenses for equipment that does not have a well-functioning secondary market, because of, for example, the inability to observe quality perfectly. In addition, firms might choose to lower production, rather than completely abandon an existing investment, if abandonment would make the investment unusable or obsolete in the future. For more details, see Dixit and Pindyck (1994) and the summary in Bredin et al. (2010). Because oil prices tend to be highly correlated with other energy prices, we might therefore expect volatility in oil prices to dampen current production, particularly in manufacturing-related industries, which tend to be energy-intensive.

At the micro level in the oil industry, these effects are estimated precisely by Kellog (2010), who finds that oil drilling activity varies with uncertainty about energy prices in a manner consistent with the theory of investment under uncertainty. It is worthwhile to note that Kellog's data set includes very small investments in drilling equipment (with some wells costing as little as a few thousand dollars) and that his measure of irreversible investments includes labor and capital. In addition, the time horizon for the drilling investments considered by Kellog is very short, being measured in days, with well depletion beginning, on average, at about seven months.

Interest in the potential negative effects of oil price volatility surged as oil prices fluctuated in 2008 from \$90 per barrel to more than \$130 per barrel, before dropping to \$40. Consider, for example, the following headlines: "Uncertainty Clouds Outlook for Oil Sector" (*New York Times*, May 18, 2009), "We Must Address Oil-Market Volatility" (by Gordon Brown, the former prime minister of the United Kingdom, and Nicolas Sarkozy, the president of France, in the *Wall Street Journal*, July 8, 2009) and "Coping with Oil Price Volatility" (*The World Bank*, August 2008). Motivated by the increased volatility in oil prices, the Chicago Board Options Exchange recently created an index of oil price volatility derived from options on oil futures, and started trading futures contracts on this volatility index.

In this paper, we continue the empirical investigation of Elder and Serletis (2010) by extending their analysis in several aspects. First, we use higher-frequency data that are more focused on manufacturing production and the production of durables goods. In particular, we use the index of industrial production compiled by the Federal Reserve Board, which concentrates on production in manufacturing, mining, and utilities. This measure of output excludes many services, and so may tend to be more closely linked to energy prices than broader measures of output, because they either use or produce relatively large amounts of energy. Compiled at a monthly frequency, this index tends to be relatively responsive to current economic conditions.

Second, we use a more fully specified four-variable model rather than a bivariate model. The four-variable model is more general, and is less likely to suffer from bias due to omitted variables and more likely to produce realistic estimates of uncertainty about oil prices. The four-variable model also isolates the aggregate price level from nominal oil prices, reducing measurement and contamination issues associated with combining the two variables. We should note, however, that some authors would clearly argue that the four-variable system is not preferred in all circumstances. Our main results, however, persist in a bivariate version of our model.

Third, we use an updated sample that includes the extreme increase in the volatility of oil prices associated with the recession of 2008.

Importantly, we also examine whether oil price uncertainty explains variation in industrial production after controlling for other nonlinear variables that have previously been shown to affect output. For example, Hamilton (2003, 2009) finds that a measure he calls the "net oil price increase (NOPI)" captures the tendency of output to respond to sustained increases in oil prices. Lee et al. (1995) find that the price of oil deflated by its current standard deviation (in order to shrink shocks that occur in periods of higher volatility) also captures this same effect. By including these variables as regressors in our empirical analysis, we assess the marginal predictive power of volatility in oil prices.

We find that the tendency of volatility in oil prices to suppress manufacturing output is stronger in monthly industrial production data than in quarterly data, that the effect is stronger for durables production, that the effect persists after controlling for popular nonlinear measures of oil prices, and that the effect persists in an updated sample that covers the period of extreme volatility in oil prices since 2008.

In addition to providing evidence in support of the predictions of real options theory, our results also have a bearing on the current debate about whether the response of output to higher oil prices is symmetric with respect to the response of output to lower oil prices. Our results indicate that volatility in oil prices increases when oil prices either rise or fall; and that oil price volatility tends to significantly depress output. Therefore, oil price volatility will tend to exacerbate the decline in production associated with higher oil prices, and dampen any immediate increase in production associated with lower oil prices.

2. STRUCTURAL VAR WITH MULTIVARIATE GARCH-IN-MEAN

The primary empirical model is developed in Elder (2004), and is based on the structural VAR with modifications for conditional heteroskedasticity in the parametric form of multivariate GARCH-in-Mean. This section reviews the empirical model briefly, as more details are available in Elder (2004) and Elder and Serletis (2010). The basic assumption is that the structural system can be represented as

$$\boldsymbol{B}\boldsymbol{z}_t = \boldsymbol{C} + \boldsymbol{\Gamma}_1\boldsymbol{z}_{t-1} + \boldsymbol{\Gamma}_2\boldsymbol{z}_{t-2} + \cdots + \boldsymbol{\Gamma}_p\boldsymbol{z}_{t-p} + \boldsymbol{\Lambda}\sqrt{\boldsymbol{H}_t} + \boldsymbol{e}_t, \tag{1}$$

where $\dim(\boldsymbol{B}) = \dim(\boldsymbol{\Gamma}_i) = \dim(\boldsymbol{\Lambda}) = (n \times n)$, $\boldsymbol{e}_t \mid \Omega_{t-1} \sim$ i.i.d. $N(\boldsymbol{0}, \boldsymbol{H}_t)$, $\boldsymbol{0}$ is the null vector, and Ω_{t-1} denotes the information set at time $t-1$, which includes variables dated $t-1$ and earlier. This specification allows the matrix of conditional standard deviations, denoted $\sqrt{\boldsymbol{H}_t}$, to affect the conditional mean, so we test whether oil price volatility affects real economic activity by examining the appropriate element of $\boldsymbol{\Lambda}$.

Elder (2004) suggests imposing the common identifying assumption that the structural disturbances are contemporaneously (and conditionally) uncorrelated, which implies that the conditional variance matrix $\boldsymbol{H}_t$ is then diagonal. In a dynamic setting, the assumption that the errors are conditionally uncorrelated is stronger than necessary, but the result is a considerable simplification in the multivariate variance function, because the conditional covariances do not need to be modeled. Of course, the reduced form errors will, in general, be correlated, but we can avoid modeling these correlations by estimating the structural parameters of the model directly. Following this reasoning, the multivariate GARCH variance function is simply

$$\boldsymbol{h}_t = \operatorname{diag}(\boldsymbol{H}_t) = \boldsymbol{C}_v + \sum_{j=1}^{J} \boldsymbol{F}_j \operatorname{diag}\left(\boldsymbol{e}_{t-j}\boldsymbol{e}'_{t-j}\right) + \sum_{i=1}^{I} \boldsymbol{G}_i \operatorname{diag}(\boldsymbol{H}_{t-i}), \qquad (2)$$

where diag is the operator that extracts the diagonal from a square matrix. We also have the baseline assumption that the conditional variance of $z_{i,t}$ depends only on its own past squared errors and its own past conditional variances, so that the parameter matrices $\boldsymbol{F}_j$ and $\boldsymbol{G}_i$ are also diagonal. The MGARCH-M VAR model can be estimated by full information maximum likelihood (FIML).

Kilian and Vigfusson (2011) note concern that the underlying GARCH model utilized by Elder and Serletis (2010), and consequently in this paper, may not be the most appropriate measure of uncertainty about oil prices. There are, of course, many possible measures of uncertainty about impending oil prices, but one might think a very reasonable alternative would be a statistical measure of the dispersion in the conditional forecast of next period's oil price, generated by a model that provides a good statistical fit to the data. That, of course, describes precisely the metric we use.

A related issue noted by Kilian and Vigfusson (2011) is that this measure of uncertainty about the price of oil may not capture the long-run uncertainty in oil prices that might be most relevant for large investment decisions. This may or may not be true, because the quantity of interest is not directly observable, but it might well be the case that if markets are highly uncertain about next month's realization of oil prices, then they may also be highly uncertain about oil prices two years hence.

Kilian and Vigfusson also commit a reductive fallacy by asserting that irreversible investment decisions apply only to large-scale fixed investment, such as automobile plants. Dixit and Pindyck (1994) are quite clear that many irreversible investment decisions occur with much shorter time horizons, such as the hiring and training of workers, as summarized previously. The study by Kellog (2010), mentioned previously, highlights the varied nature and short horizon of many such investments.

3. SPECIFICATION

The model we estimate is a very general four-variable simultaneous equations system that is generally consistent with, for example, Bernanke et al. (1997), Elder (2004), Hamilton and Herrera (2004), and many others. Our four-variable VAR uses monthly observations of inflation, the growth rate of industrial production, the growth rate of nominal oil prices, and a short-term nominal interest rate, and is very similar to that of Bredin et al. (2010). Our price series is the personal consumption expenditures (PCE) price index, excluding food and energy. Excluding food and energy ensures that the role of energy prices will be attributed to our oil price variable.

Production is measured by the index of industrial production compiled by the Federal Reserve Board, which captures the output in energy-intensive sectors such as mining, manufacturing, and gas and electric utilities. Manufacturing output, in particular, includes the industries specified by the North American Industry Classification System, plus the logging and publishing industries. Kilian and Vigfusson (2011) express concern that the measures of output used by Elder and Serletis (2010) are not weighted by the share of energy in value added. Although we have not estimated our model with such measures of output, we do not believe this issue to be material. As noted by Kilian (2008, p. 875), "while not trivial, the observed fluctuations in the energy share in value added are largely immaterial for estimates of energy price shocks, because the share does not fluctuate enough on a quarter-to-quarter basis. Weighted and unweighted quarterly energy price changes have a correlation of 99 percent."

The price of oil is measured by the spot West Texas Intermediate (WTI) crude oil price at Chicago. We use the nominal price of oil in our model, rather than the real price of oil, to better segregate uncertainty about oil prices from the aggregate price level. We understand that there is not universal agreement on whether the nominal price of oil or the real price of oil is more appropriate. Although this issue is not important for our primary empirical results, we use the nominal price of oil rather than the real price of oil, because the real price of oil introduces an additional component of measurement error into our estimated measure of oil price uncertainty, which would then capture uncertainty about both nominal oil prices and whichever measure of aggregate inflation we use. See also Hamilton (2008) regarding this issue.

The short-term interest rate is the federal funds rate. Inclusion of the federal funds rate facilitates identifying the effects of oil shocks, as it controls for the endogenous response of monetary policy, which Bernanke et al. (1997) have argued to be empirically important. We should also note that the inclusion of the federal funds rate is not central to our primary results, which also hold in a bivariate version of the model.

TABLE 1. Data description

Series	Transformation	Description
Infl	$100 \times \ln(PCE_t/PCE_{t-1})$	Personal conumption expenditures (PCE) price index, less food and energy, seasonally adjusted
Production	$100 \times \ln(IP_t/IP_{t-1})$	Industrial production index, seasonally adjusted
Durables	$100 \times \ln(IPDUR_t/IPDUR_{t-1})$	Industrial production — durables manufacturing, seasonally adjusted
Oil	$100 \times \ln(P_t/P_{t-1})$	West Texas Intermediate crude oil price
Funds	None	Daily average of fed funds rate, not seasonally adjusted
NOPI_12	$\max\left[0, \ln\left(\frac{P_t}{\max\,(P_{t-1}, \cdots, P_{t-12})}\right)\right]$	Net oil price increase over one year, as defined by Hamilton (2003)
NOPI_36	$\max\left[0, \ln\left(\frac{P_t}{\max\,(P_{t-1}, \cdots, P_{t-36})}\right)\right]$	Net oil price increase over three year, as defined by Hamilton (2003)
Oilnorm	$Oilnorm = \frac{Oil(t) - Infl(t)}{\sqrt{\hat{H}_{3,3}(t)}}$	Volatility adjusted real oil price, as defined by Hamilton (2003)
Oilnorm$^+$	$\max\,(Oilnorm, 0)$	Volatility adjusted real oil price, positive values only, as defined by Hamilton (2003)

Table 1 provides a description of the data. Our sample begins in 1979:1, when oil prices began to exhibit increased volatility. Prior to 1979, oil prices remained unchanged for considerable periods and then adjusted rapidly. Our starting date is roughly consistent with the dates of structural change utilized by Hamilton (2008). We terminate our sample in 2009:12, which includes the most recent recession and the episode of greatest volatility in oil prices.

As noted by Elder (2004), the proper approach to estimating the multivariate GARCH-in-Mean VAR is to difference the variables that are integrated. Therefore, in our model, we use the first difference of the log of the PCE price index, denoted by *Infl*, the first difference of the logged oil price, denoted by *Oil*, and the first difference of the log of the industrial production index, denoted by *Production*. The mnemonic for the federal funds rate is *Funds*.

Identification of the underlying structural parameters in equations (1) and (2) is achieved with exclusion restrictions on $\boldsymbol{B}$ and the assumption that the structural disturbances are conditionally uncorrelated. The exclusion restrictions are generally consistent with economic priors. In particular, we assume that *Oil* shocks and *Funds* shocks affect output and inflation with a lag. *Oil* and *Funds* are derived from prices in unregulated markets, so it is natural to put them at the bottom of the ordering. In particular, *Oil* is an unregulated commodity price with a deep and active futures market that adjusts rapidly to new information. We therefore allow *Oil* to respond to contemporaneous innovations in both *Infl* and *Production*. We allow *Funds* to respond to contemporaneous innovations in *Infl*, *Production*, and *Oil*. With one remaining free parameter, we allow *Production* to respond to contemporaneous innovations in *Infl*. These identifying restrictions imply that $\boldsymbol{B}$ is lower triangular with the ordering *Infl*, *Production*, *Oil*, and *Funds*, although triangularity of $\boldsymbol{B}$ is not required by our estimation scheme or computer programs. We should also note that our results tend not to be affected by the choice of ordering, such as putting *Production* prior to *Infl*; the motivation for putting *Infl* prior to *Production* is based on the proposition that aggregate prices are "sticky," or somehow slow to respond to current macroeconomic conditions.

Uncertainty about oil prices is measured by the conditional standard deviation of *Oil*, which is $\sqrt{H_{3,3}(t)}$ in this model. As noted by Elder (2004), this is the standard deviation of the one–period ahead forecast error, in this case for *Oil*, conditional on the contemporaneous information set. That is, $\text{var}(\boldsymbol{e}_t \mid \Omega_{t-1}) = H_{3,3}(t)$. It is therefore a statistical measure of how uncertain markets are about the impending realization of *Oil*. Because we are interested in measuring the effect of *Oil* uncertainty on *Production*, we allow *Oil* uncertainty to enter the *Production* equation, and estimate the appropriate element of $\boldsymbol{\Lambda}$ (that is, $\Lambda_{2,3}$) in equation (1).

Relative to the bivariate model in Elder and Serletis (2010), this four-variable model has several important advantages. First, the four-variable model has fewer parametric restrictions, and is therefore less likely to suffer from omitted variable bias. Second, the four-variable model conditions on more information, and so our measure of oil price uncertainty is likely to be more realistic. Third, by including separate variables for the nominal price of oil and the aggregate price level, we isolate the effects of oil prices from those of aggregate prices, diminishing measurement and contamination issues that might be associated with using the real price of oil.

Consistent with Elder and Serletis (2010), we include a full year of lags, given the arguments advanced by Hamilton and Herrera (2004) and Edelstein and Kilian (2007). These authors stress that the primary effect of oil prices on real output occurs at or before one year, and so emphasize the importance of including at least one year of lags.

4. EMPIRICAL RESULTS

The system we estimate is therefore a 12-lag multivariate GARCH-in-Mean VAR with monthly observations on *Infl*, *Production*, *Oil*, and *Funds* over a usable sample of 1980:1–2009:12. Table 2 reports the point estimates for the multivariate GARCH variance function parameters. There are strong GARCH effects in both

TABLE 2. Estimates of variance function of the multivariate GARCH-In-Mean VAR, equations (1)–(2)

Equation	Conditional variance	Constant	$e_i^2(t-1)$	$H_{i,i}(t-1)$
Infl	$H_{1,1}(t)$	0.001*	0.484*	0.232
		(3.411)	(3.08)	(1.42)
Production	$H_{2,2}(t)$	0.002*	0.861*	0.000
		(14.68)	(6.60)	(0.00)
Oil	$H_{3,3}(t)$	0.037	0.369*	0.613*
		(1.89)	(6.65)	(11.06)
Funds	$H_{4,4}(t)$	0.004*	0.685*	0.310*
		(4.07)	(8.17)	(3.76)

Notes: These are the parameter estimates for the free elements in $\boldsymbol{F}$ and $\boldsymbol{G}$ from the model given by equations (1) and (2) with $e_t \sim N(\mathbf{0}, \boldsymbol{H}_t)$. Each row in the table represents an equation from the associated multivariate GARCH-in-Mean VAR. Asymptotic t-statistics are in parentheses. A coefficient of 0.000 indicates that the nonnegativity constraint is binding.
*Significance at the 5% level.

Oil and *Funds*, with the coefficients on the lagged squared errors and the lagged conditional variances both being highly significant, and the volatility processes being very persistent. There is also evidence of ARCH in *Production* and *Infl* over this sample period.

The effect of the conditional volatility in *Oil* on *Production* is reported in Table 3 as model (1). In model (1), the point estimate on this coefficient is -0.015, with an asymptotic t-statistic of 3.48. Our model therefore indicates that after controlling for the effects of lagged output, the price level, oil prices, and the short-term interest rate, the conditional volatility of oil prices has tended to cause industry production to decline. This result is stronger than the comparable result in a similar model with quarterly GDP data reported in Elder and Serletis (2010). This should not be surprising, given that the real GDP data include sectors of the

TABLE 3. Coefficient estimates for oil uncertainty in the output equation of model (1)–(2)

Model	Variables	Sample	Coefficient on $\sqrt{H_{3,3}(t)}$
(1) MGARCH-M VAR	{*Infl*, *Production*, *Oil*, *Funds*}	1980:1–2009:12	-0.015^{**}
			(3.48)
(2) MGARCH-M VAR	{*Infl*, *Durables*, *Oil*, *Funds*}	1980:1–2009:12	-0.026^{*}
			(2.52)

Notes: These are estimates for the free element in from the model given by equations (1) and (2) with $\boldsymbol{e}_t \mid \Omega_{t-1} \sim$ iid $N(\mathbf{0}, \boldsymbol{H}_t)$. $\sqrt{H_{3,3}(t)}$ denotes the conditional standard deviation of *Oil*. Absolute asymptotic t-statistics are in parentheses.
*Significance at the 5% level.
**Significance at the 1% level.

economy, such as service sectors, that are not likely to be very sensitive to oil prices, whereas our production data derive primarily from energy-intensive areas.

The first three panels of Figure 1 plot industrial production (the year-over-year growth rate is shown), oil prices, and the conditional standard deviation of *Oil* over the 1980:1–2009:12 sample, with NBER recessions shaded. These plots are similar to those of Elder and Serletis (2010), but the additional granularity provided by the higher-frequency data is informative, as is the extended sample. For example, in monthly data, the increase in oil price volatility in 1980 and 1982 is more clear, as well as the general pattern of oil price volatility being high when oil prices make large moves either up or down. As in Elder and Serletis (2010), the spike in oil price volatility during the mid-1980s is evident, as is the elevated volatility in oil prices, but without spikes, from 2002 to 2005, as oil prices rose steadily. Because increases in oil price volatility tend to cause production to decline, these results provide at least a partial explanation for why manufacturing activity did not surge during the mid-1980s as oil prices collapsed, and why the slow but steady rises in oil prices since 2002 were less disruptive than the rapid price increase during, for example, the 1990 recession. For other, more detailed, interpretations of this period, see Hamilton (2009).

During the recession of 2008, manufacturing production and oil prices both collapsed. Without doubt, the collapse in oil prices during 2008 was partially due to a decline in the demand for oil. Our model should capture this effect to the extent that it is captured by the relationship between oil prices and current and lagged manufacturing activity. Our estimates suggest that, after accounting for the decline in the demand for oil, the surge in the conditional volatility of oil prices further contributed to the decline in manufacturing activity. This at least lends credence to the popular view among policymakers at that time that oil price volatility had significant and substantial effects on the real economy. See, for example, "We Must Address Oil-Market Volatility," by Gordon Brown and Nicolas Sarkozy (*Wall Street Journal*, July 8, 2009).

Next, we investigate the robustness of our results by using an alternative measure of the level of economic activity. In particular, the theory suggests that uncertainty about oil prices should affect production, investment, and consumption decisions that are either irreversible or costly to reverse. We should therefore expect uncertainty about oil prices to have a more pronounced effect on the production of relatively illiquid durable goods than on a broader index of manufacturing output. In Table 3 under model (2), we report the results from reestimating our model with an index of the production of durable goods as the measure of output. Durables manufacturing is a component of the industrial production index, representing more than 40% of total manufacturing, mining, and utility output (as of 2003). This index of durables manufacturing includes transportation equipment (e.g., automobiles, aircraft, and related parts, representing 11% of the total index), computers and electronics (representing about 8%), machinery (e.g., construction equipment representing about 6%), and fabricated metal products (e.g., tools, representing about 6%).

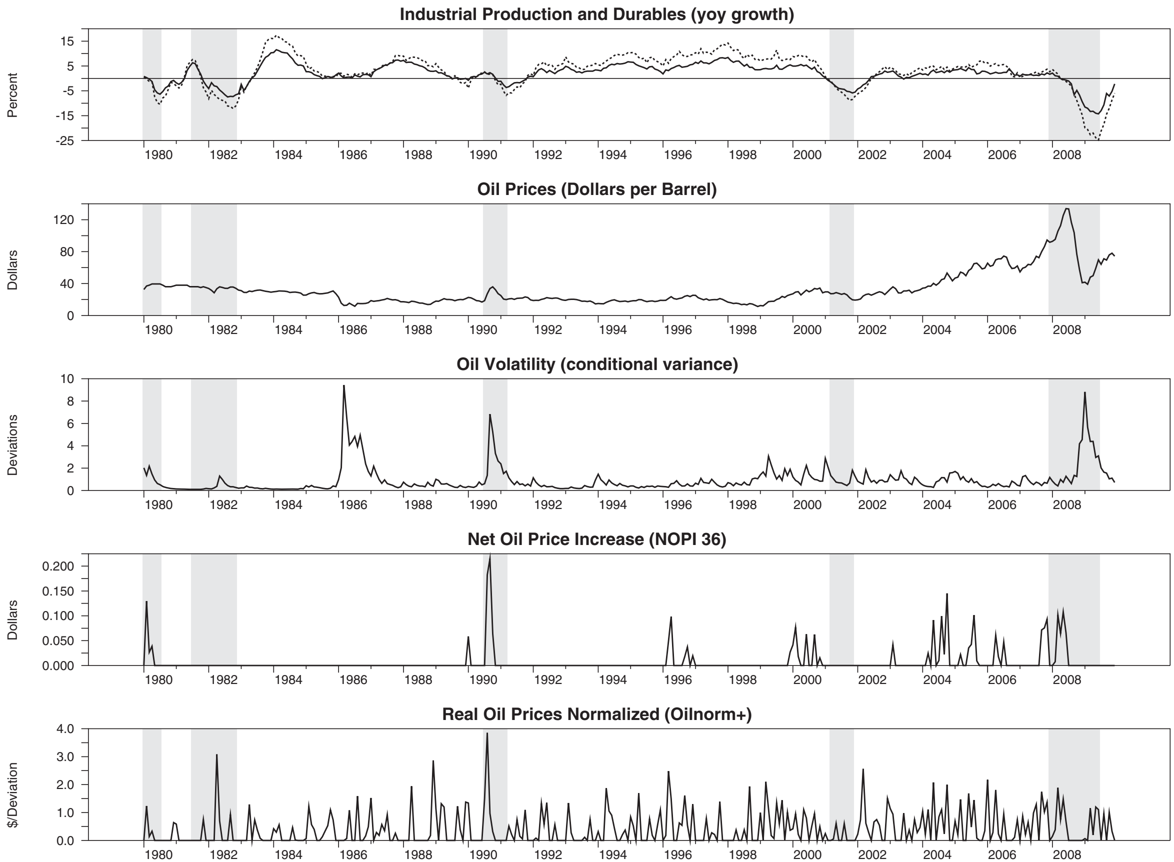

FIGURE 1. Output and oil price volatility. NBER recessions shaded.

TABLE 4. Coefficient estimates for oil uncertainty in the output equation of model (3)

Model	Variables	Sample	Coefficient on $\sqrt{\hat{H}_{3,3}(t)}$	Exclusion test on *NOPI* and *Oilnorm* (p-value)
(1) VAR	{*Infl*, *Production*, *Oil*, *Funds*}	1980:1–2009:12	−0.043** (3.92)	*N/A*
(2) VAR	{*Infl*, *Production*, *Oil*, *NOPI_12*, *Funds*}	1980:1–2009:12	−0.046** (3.03)	0.05
(3) VAR	{*Infl*, *Production*, *Oil*, *NOPI_36*, *Funds*}	1980:1–2009:12	−0.031* (2.52)	0.01
(4) VAR	{*Infl*, *Production*, *Oil*, *Oilnorm*, *Funds*}	1980:1–2009:12	−0.036** (3.18)	0.33
(5) VAR	{*Infl*, *Production*, *Oil*, $Oilnorm^{+}$, *Funds*}	1980:1–2009:12	−0.034** (2.69)	0.15
(6) VAR	{*Infl*, *Production*, *Oil*, *Funds*}	1988:1–2009:12	−0.045** (2.85)	*N/A*
(7) VAR	{*Infl*, *Durables*, *Oil*, *NOPI_36*, *Funds*}	1980:1–2009:12	−0.390* (2.45)	0.28
(8) VAR	{*Infl*, *Durables*, *Oil*, *Oilnorm*, *Funds*}	1980:1–2009:12	−0.468** (3.30)	0.05
(9) VAR	{*Infl*, *Durables*, *Oil*, $Oilnorm^{+}$, *Funds*}	1980:1–2009:12	−0.419** (2.63)	0.18
(10) VAR	{*Infl*, *Production*, *Oil*, *Funds*}	1980:1–2009:12 (excluding outliers)	−0.023** (3.55)	*N/A*

Notes: These models are estimated as a standard VAR with contemporaneous *Oil* volatility, $\sqrt{\hat{H}_{3,3}(t)}$ generated from model (1) in Table 3, included in the conditional mean equation. The last column reports p-values for exclusion tests of the nonlinear transformation of the price of oil. Absolute asymptotic t-statistics are in parentheses.
*Significance at the 5% level.
**Significance at the 1% level.

With the manufacture of durables as the measure of output, the point estimate for the effect of *Oil* uncertainty on durables manufacturing is negative and significant at the 5% level [reported as model (2) in Table 3]. This coefficient is greater in magnitude, which may be because of the volatility characteristics of durables manufacturing, or may suggest a more pronounced effect of uncertainty about oil prices through the manufacture of durable goods.

We further investigate the robustness of our results by estimating different specifications over different sample periods, with more concentrated measures of production and various nonlinear transformations of the price of oil. The results from these specifications are reported in Table 4.

Hamilton (2003, 2011) suggests that realized oil prices, corresponding with their increased volatility since the late 1970s, are not the appropriate metric for

https://doi.org/10.1017/S1365100511000630 Published online by Cambridge University Press

capturing the effect of rising oil prices on real economic activity. In particular, he shows that the effect of oil prices is nonlinear, with sustained increases in oil prices having very different effects than transitory increases or sustained decreases in oil prices. Hamilton (2003) and Lee et al. (1995) suggest several alternative measures of oil prices that differentiate between transitory and sustained increases in oil prices. In particular, Hamilton (2003) considers measuring the effect of sustained increases in oil prices by the amount by which oil prices exceed their peak value over the previous 12 or 36 months, referring to this metric as the net oil price increase, which we denote by the mnemonics *NOPI_12* and *NOPI_36*. *NOPI_36* for our sample is plotted in the fourth panel of Figure 1.

Lee et al. (1995) suggest a different metric, based on the rate of change in real oil prices standardized by the conditional standard deviation of oil prices. They motivate this measure by arguing that a given change in oil prices should have a smaller effect when the price change is believed to be transitory. Hamilton (2003) finds empirical support for a similarly defined variable, which is based on a variable we denote *Oilnorm*. In our notation, *Oilnorm* can be represented as

$$Oilnorm = \frac{Oil(t) - Infl(t)}{\sqrt{\hat{H}_{3,3}(t)}}.$$

This measure is comparable to that defined by Hamilton (2003 p. 379), but is a modification of the metric calculated by Lee et al. (1995). These authors also use a transformation of this variable, replacing negative values by 0. We denote this transformation by $Oilnorm^{+}$, and plot it in the fifth panel of Figure 1.

In order to assess whether our measure of oil price uncertainty affects *Production* when controlling for these alternative measures of oil prices, we estimate VARs with the following variables: *Infl*, *Production*, *Oil*, the alternative measure of oil prices, and *Funds*. We then include our measure of *Oil* uncertainty, $\sqrt{\hat{H}_{3,3}(t)}$, generated from the multivariate GARCH-in-Mean VAR, equations (1) and (2), as a right hand–side variable. This VAR can be represented as

$$\boldsymbol{B z}_t = \boldsymbol{C} + \boldsymbol{\Gamma}_1 \boldsymbol{z}_{t-1} + \boldsymbol{\Gamma}_2 \boldsymbol{z}_{t-2} + \cdots + \boldsymbol{\Gamma}_p \boldsymbol{z}_{t-p} + \mathbf{D}(L)x_t + \boldsymbol{\Lambda}\sqrt{\hat{H}_{3,3}(t)} + \boldsymbol{e}_t, \quad (\mathbf{3})$$

where $\boldsymbol{e}_t \sim N(\mathbf{0}, \boldsymbol{\Phi})$ and x_t is the additional regressor, in turn (12 lags of) *NOPI_12*, *NOPI_36*, *Oilnorm,* and $Oilnorm^{+}$. We adopt this approach because attempting to estimate all the parameters simultaneously is either not possible or very complex. The method is comparable to that of Hamilton [2003, equations (3.2) and (3.3)] and Lee et al. (1995).

As a baseline, we first report the coefficient estimate and asymptotic t-statistic on oil price uncertainty when it is included as a generated regressor in our baseline VAR. The model is comparable to model (1) in Table 3, except that oil price uncertainty is included as a generated regressor. This is reported as model (1) in Table 4, and the coefficient is again negative and significant. We also estimated this model using the Chicago Fed National Activity Index (CFNAI) in order to shed some light on the transmission channel of oil prices, as the CFNAI includes, in addition to industrial production, other economic activity indicators such as consumption and housing that are likely to be affected by increased uncertainty. The resulting coefficient estimate is -0.262, with an absolute t-statistic of 2.98. Although we cannot directly compare the magnitude of this coefficient to that of the coefficient when *Production* is used, because of the construction of the index, the coefficient is still negative and statistically significant.

The effects of *Oil* uncertainty on *Production*, controlling for each of *NOPI_12*, *NOPI_36*, *Oilnorm,* and *Oilnorm*$^+$ are reported in Table 4 as models (2) through (5). In each case, the effect of *Oil* uncertainty is similar in magnitude to that in the original model and is statistically significant, with the smallest (absolute) t-statistic equal to 2.52 for the model with *NOPI_36*. These results reinforce our finding that the effect of oil price uncertainty is in addition to any effects associated solely with changes in the level of oil prices, as captured by these nonlinear transformations.

Elder and Serletis (2010) test whether oil price uncertainty affects output in more recent samples, post 1986, recognizing Edelstein and Kilian's (2007) suggestion that apparent asymmetries in the response of output to oil prices may be due to the Tax Reform Act of 1986. We also conduct that test here, by modifying the sample in our VAR to 1988:1–2009:12. The result is reported as model (6) in Table 4. Again, the effect of oil price volatility is negative and highly significant.

In models (7)–(9) in Table 4, we repeat these calculations with the production of durables as our measure of output. The results are comparable, with the coefficient on oil price volatility negative, larger in magnitude, and highly significant.

Finally, given the extreme variation in the volatility of oil prices, it is worthwhile to examine whether our results are sensitive to a small number of extreme observations. To investigate this issue, we estimate the VAR given by equation (3), except that we exclude observations for which the residual in the output equation is greater than one standard deviation. This drops the number of observations from 360 to 280. The parameter coefficient is reported in Table 4 as model (10). Again, we find support for our finding that oil price volatility has had a negative and significant effect on output. Similar results are obtained if residuals larger than two, three, or four standard deviations are dropped.

5. IMPULSE–RESPONSE ANALYSIS

We next conduct impulse–response analysis for our multivariate GARCH-in-Mean VAR. The method we use is described in Elder (2003, 2004). The procedure is relatively straightforward, even though the impulse–response functions are nonlinear. As noted by Elder (2003), the impulse–response function consists of two components—the effect of the shock on the conditional mean vector (analogous to that of a standard linear VAR) and the effect of the shock on the conditional volatility, which affects the conditional mean as detailed in equation (1). Gallant et al. (1993) show that the linear–quadratic structure of GARCH models causes the latter component to behave possibly less surprisingly than impulse responses of other nonlinear models, so that the persistence of this component is characterized in the multivariate case by the terms derived in Elder (2003). In our parameterization, the magnitude of the response of volatility also scales proportionally with the size of the shock.

We simulate two responses of the growth rate of industrial production to an oil shock, and report these in Figure 2. Kilian and Vigfusson (2011) express concern about the size of the shock that Elder and Serletis (2010) use to simulate their impulse responses. In particular, Elder and Serletis (2010) use the unconditional standard deviation of the change in the price of oil rather than the unconditional standard deviation of the residual. In practice, this is a distinction without a difference, as the ratio of the latter to the former is about 0.99.

The dashed line reports the response of industrial production to an oil shock, in which the uncertainty effect is ignored. This is comparable to the impulse–response function for a standard VAR and is calculated by constraining $\mathbf{\Lambda} = \mathbf{0}$ in equation (1). This impulse–response function indicates that a (unconditional) one–standard deviation increase in oil prices causes industrial production to fall, with the response most severe between four and seven months after the shock. Next, we simulate the effect of an increase in oil prices on industrial production, allowing for the channel through which oil price volatility affects industrial production. This is plotted as the solid line in Figure 2. Clearly, once the effect of oil price volatility is considered, the decline in industrial production in response to higher oil prices is decidedly more negative. In the "real options" view, the decline in production is more precipitous, because of the firm's increased uncertainty about the future path of oil prices.

6. CONCLUSION

Real options analysis predicts that uncertainty about the return to a project will induce firms to delay production and investment. Our empirical results are consistent with this prediction, with uncertainty about oil prices having a significantly negative effect on industrial production, after controlling for lagged oil prices, lagged output, lagged inflation, and lagged interest rates. The most pronounced effect of uncertainty about oil prices appears to be through the manufacture of durable goods, including automobiles and other transportation equipment. Our result is robust to alternative specifications, controlling for other nonlinear measures of oil prices, and various sample periods.

Our results suggest that the attention devoted by policy makers and the popular press to the extreme increase in the volatility of oil prices during the recession of 2008 was not unwarranted. The volatility in oil prices during this time period likely contributed to both the depth and severity of the decline in manufacturing activity, as well as the failure of manufacturing activity to rebound more immediately as oil prices fell.

Kilian and Vigfusson (2011) raise concerns that the models of Elder and Serletis (2010), and consequently the ones in this paper, do not fit all the features

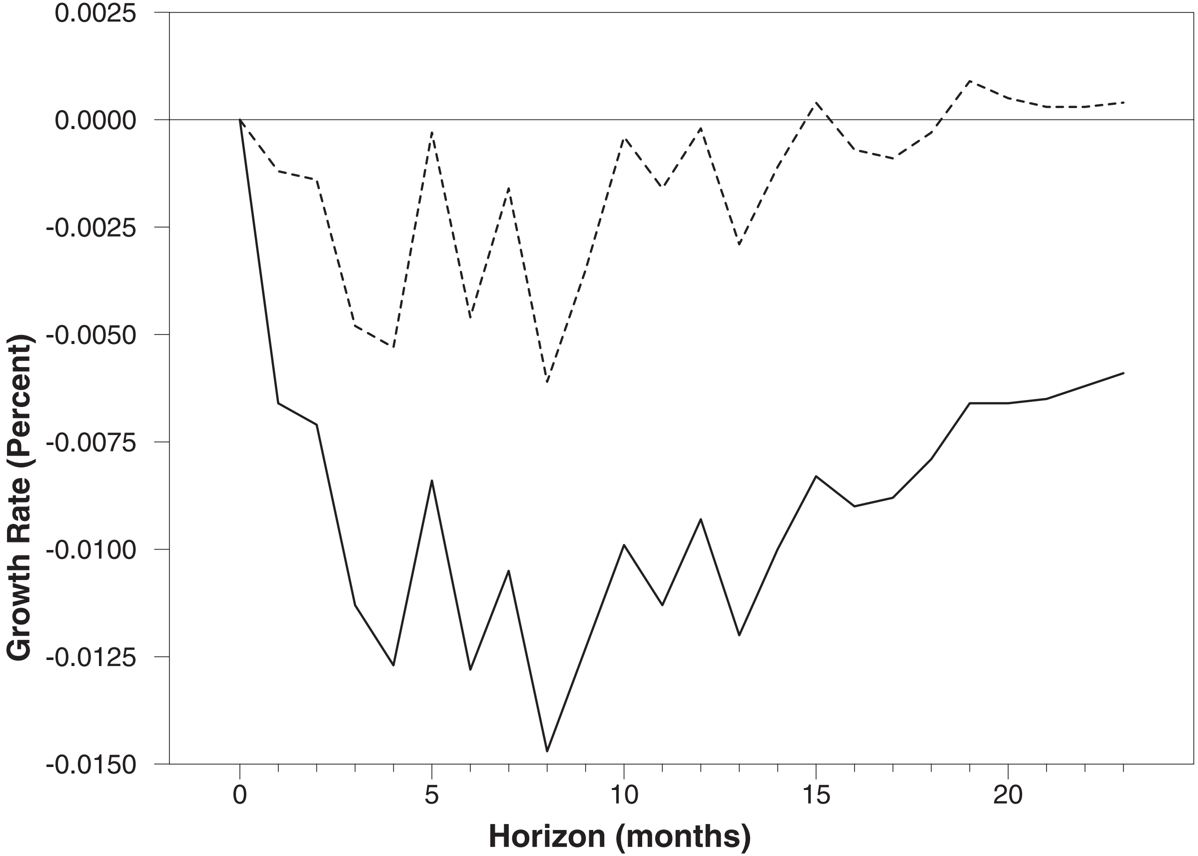

FIGURE 2. Response of industrial production to oil shock with (solid) and without (dashed) uncertainty effect.

of the data, and have some anomalous results, depending on the particular sample period and data analyzed. We acknowledge that we have not identified the single valid macroeconomic model that governs the relationship between oil prices and real economic activity. But nobody else has either, and we believe that macroeconomists will continue to develop, estimate, and refine stylistic models that fit salient features of the data. One test of whether such models are useful, however, is whether they describe important features of the data out of sample. This paper verifies this, by showing that the relationship identified by Elder and Serletis (2010) continues to capture a negative relation between oil price uncertainty and output growth after the most recent recession, which was preceded by an unprecedented spike in oil price uncertainty.

REFERENCES

Bernanke, Ben S. (1983) Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* 98, 85–106.

Bernanke, Ben S., Mark Gertler, and Mark Watson (1997) Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity* 1, 91–142.

Bredin, Don, John Elder, and Stilianos Fountas (2010) The Effects of Uncertainty about Oil Prices in G-7. Working paper, Colorado State University.

Brennan, Michael (1990) Latent assets. *Journal of Finance* 45, 709–730.

Brennan, Michael and Eduardo Schwartz (1985) Evaluating natural resource investment. *Journal of Business* 58, 1135–1157.

Dixit, Avinash and Robert S. Pindyck (1994) *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Edelstein, Paul and Lutz Kilian (2007) The response of business fixed investment to changes in energy prices: A test of some hypotheses about the transmission of energy price shocks. *B.E. Journal of Macroeconomics* (Contributions) 7, Article 35.

Elder, John (2003) An impulse-response-function for a vector autoregression with multivariate GARCH-in-Mean. *Economics Letters* 79, 21–26.

Elder, John (2004) Another perspective on the effects of inflation volatility. *Journal of Money, Credit and Banking* 36, 911–928.

Elder, John and Apostolos Serletis (2010) Oil price uncertainty. *Journal of Money, Credit and Banking* 42, 1138–1159.

Gallant, Ronald A., Peter E. Rossi, and George Tauchen (1993) Nonlinear dynamic structures. *Econometrica* 61(4), 871–907.

Gibson, Rajna and Eduardo S. Schwartz (1990) Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance* 45, 959–976.

Hamilton, James D. (2003) What is an oil shock? *Journal of Econometrics* 113, 363–398.

Hamilton, James D. (2008) Oil and the macroeconomy. In *New Palgrave Dictionary of Economics*, 2nd ed. London: Macmillan.

Hamilton, James D. (2009) Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity* 1, 215–261.

Hamilton, James D. (2011) Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics* [this issue].

Hamilton, James D. and Ana Maria Herrera (2004) Oil shocks and aggregate macroeconomic behavior: The role of monetary policy. *Journal of Money, Credit and Banking* 36, 265–286.

Kellogg, R. (2010) The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling. NBER working paper 16541 (http://www.nber.org/papers/w16541).

Kilian, Lutz (2008) The economic effects of energy price shocks. *Journal of Economic Literature* 46, 871–909.

Kilian, Lutz and Robert J. Vigfusson (2011) Nonlinearities in the oil price–output relationship. *Macroeconomic Dynamics* [this issue].

Lee, Kiseok, Shawn Ni, and Ronald A. Ratti (1995) Oil shocks and the macroeconomy: The role of price variability. *Energy Journal* 16, 39–56.

Majd, Saman and Robert S. Pindyck (1987) Time to build, option value, and investment decisions. *Journal of Financial Economics* 18, 7–27.