

THE ASYMMETRIC EFFECTS OF OIL PRICE SHOCKS

SAJJADUR RAHMAN

University of Saskatchewan

APOSTOLOS SERLETIS

University of Calgary

In this paper we investigate the effects of oil price uncertainty and its asymmetry on real economic activity in the United States, in the context of a bivariate vector autoregression with GARCH-in-mean errors. The model allows for the possibilities of spillovers and asymmetries in the variance–covariance structure for real output growth and the change in the real price of oil. Our measure of oil price uncertainty is the conditional variance of the oil price–change forecast error. We isolate the effects of volatility in the change in the price of oil and its asymmetry on output growth and employ simulation methods to calculate generalized impulse response functions and volatility impulse response functions to trace the effects of independent shocks on the conditional means and the conditional variances, respectively, of the variables. We find that oil price uncertainty has a negative effect on output, and that shocks to the price of oil and its uncertainty have asymmetric effects on output.

Keywords: Crude Oil, Volatility, Vector Autoregression, Multivariate GARCH-in-Mean VAR

1. INTRODUCTION

Questions regarding the relationship between the price of oil and economic activity are fundamental empirical issues in macroeconomics. Hamilton (1983) showed that oil prices had significant predictive content for real economic activity in the United States prior to 1972, whereas Hooker (1996) argued that the estimated linear relations between oil prices and economic activity appear much weaker after 1973. In the debate that followed, it has been suggested that the apparent weakening of the relationship between oil prices and economic activity is illusory and that the true relationship between oil prices and real economic activity is asymmetric, with the correlation between oil price decreases and output significantly different

This paper builds on material from Chapter 3 of Sajjadur Rahman's Ph.D. dissertation at the University of Calgary. We would like to thank Gary Koop, John Elder, Christian Hafner, Ólan Henry, Helmut Herwartz, Kalvinder Shields, Chris Auld, Zuzana Janko, and seminar participants at the University of Calgary, the University of Saskatchewan, and the University of New Brunswick. We also thank the members of Sajjadur's dissertation committee: Philip Chang, Daniel Gordon, John Keating, Gordon Sick, and David Walls. Address correspondence to: Apostolos Serletis, Department of Economics, University of Calgary, Calgary, Alberta T2N 1N4, Canada; e-mail: serletis@ucalgary.ca; URL: <http://econ.ucalgary.ca/serletis.htm>.

from the correlation between oil price increases and output—see, for example, Mork (1989) and Hamilton (2003). More recently, however, Edelstein and Kilian (2007, 2008) evaluated alternative hypotheses and argued that the evidence of asymmetry cited in the literature is driven by a combination of ignoring the effects of the 1986 Tax Reform Act on fixed investment and the aggregation of energy and non-energy related investment.

There is a vast literature that investigates the effects of oil prices on the real economy. However, relatively few studies consider the effects of uncertainty about oil prices on real economic activity. Lee et al. (1995) were the first to employ recent advances in financial econometrics and model oil price uncertainty using a univariate GARCH (1,1) model. They calculated an oil price shock variable, reflecting the unanticipated component as well as the time-varying conditional variance of oil price changes, introduced it into various vector autoregression (VAR) systems, and found that oil price volatility is highly significant in explaining economic growth. They also found evidence of asymmetry, in the sense that positive shocks have a strong effect on growth whereas negative shocks do not. The Lee et al. (1995) approach, however, is subject to the generated regressor problem, described by Pagan (1984).

More recently, Elder and Serletis (2010) examined the direct effects of oil price uncertainty on real economic activity in the United States, over the modern OPEC period, in the context of a structural VAR that is modified to accommodate GARCH-in-mean errors, as detailed in Engle and Kroner (1995) and Elder (2004). As a measure of uncertainty about the impending oil price, they use the conditional standard deviation of the forecast error for the change in the price of oil. Their main result is that uncertainty about the price of oil has had a negative and significant effect on real economic activity over the post-1975 period, even controlling for lagged oil prices and lagged real output. Their estimated effect is robust to a number of different specifications, including alternative measures of the price of oil and of economic activity, as well as alternative sample periods. They also find that accounting for oil price uncertainty tends to reinforce the decline in real GDP in response to higher oil prices, although moderating the short-run response of real GDP to lower oil prices.

In this paper we move the empirical literature forward by investigating the asymmetric effects of uncertainty on output growth and oil price changes, as well as the response of uncertainty about output growth and oil price changes to shocks. In doing so, we use an extremely general bivariate framework in which a vector autoregression is modified to accommodate GARCH-in-mean errors, as detailed in Engle and Kroner (1995), Grier et al. (2004), and Shields et al. (2005). The model allows for the possibilities of spillovers and asymmetries in the variance-covariance structure for real activity and the real price of oil. As in Elder and Serletis (2010), our measure of oil price change volatility is the conditional variance of the oil price change forecast error. We isolate the effects of oil price change volatility and its asymmetry on output growth and, following Koop et al. (1996), Grier et al. (2004), and Hafner and Herwartz (2006), we employ simulation

methods to calculate generalized impulse response functions (GIRFs) and volatility impulse response functions (VIRFs) to trace the effects of independent shocks on the conditional means and the conditional variances, respectively, of the variables.

We find that our bivariate, GARCH-in-mean, asymmetric VAR-BEKK model embodies a reasonable description of the monthly U.S. data, over the period from 1980:1 to 2010:7. We present evidence that increased uncertainty about the change in the real price of oil is associated with a lower average rate of growth in real economic activity. Generalized impulse response experiments highlight the asymmetric effects of positive and negative shocks in the change in the real price of oil on output growth. Also, volatility impulse response experiments reveal that the effect of good news (negative shocks in the change in the real price of oil) on the conditional variance of the change in the real price of oil differs in magnitude and persistence from that of bad news of similar magnitude. This result suggests that, given the relationship between oil price volatility and output, the asymmetric response of oil price volatility to oil price shocks might be a contributing factor in explaining the asymmetric relationship between oil prices and economic activity.

The paper is organized as follows. Section 2 presents the data and Section 3 provides a brief description of the bivariate, GARCH-in-mean, asymmetric VAR-BEKK model. Sections 4, 5, and 6 assess the appropriateness of the econometric methodology by various information criteria and present and discuss the empirical results. The final section concludes the paper.

2. THE DATA

We use monthly data for the United States from the Federal Reserve Economic Database (FRED) maintained by the Federal Reserve Bank of St. Louis, over the period from 1980:1 to 2010:7, on two variables—the industrial production index (y_t) and the real price of oil (oil_t). In particular, we use the spot price on West Texas Intermediate (WTI) crude oil as the nominal price of oil and divide it by the consumer price index (CPI) to obtain the real price of oil. Following Bernanke et al. (1997), Lee and Ni (2002), Hamilton and Herrera (2004), and Edelstein and Kilian (2008), we use the industrial production index as a proxy variable for real output. It is to be noted that industrial output reflects only manufacturing, mining, and utilities, and represents only about 20% of total output. The industrial production index, however, captures economic activity that is likely to be directly affected by oil prices and uncertainty about oil prices.

Table 1 presents summary statistics for the annualized logarithmic first differences of y_t and oil_t , denoted as $\Delta \ln y_t$ and $\Delta \ln oil_t$, and Figures 1 and 2 plot the $\ln y_t$ and $\Delta \ln y_t$ and $\ln oil_t$ and $\Delta \ln oil_t$ series, respectively, with shaded area indicating NBER recessions. Both $\Delta \ln y_t$ and $\Delta \ln oil_t$ are skewed and there is a significant amount of excess kurtosis present in the data. Moreover, a Jarque–Bera (1980) test for normality, distributed as a $\chi^2(2)$ under the null hypothesis

TABLE 1. Summary statistics

A. Summary statistics					
Variable	Mean	Variance	Skewness	Excess kurtosis	J-B normality
$\Delta \ln y_t$	1.951	70.174	-0.956	4.375	348.789 (0.000)
$\Delta \ln \text{oil}_t$	-0.607	9583.058	-0.344	3.374	181.375 (0.000)
B. Unit root and stationary tests					
Variable	Unit root tests			KPSS stationarity tests	
	ADF(τ)	ADF(μ)	ADF	$\hat{\eta}_\mu$	$\hat{\eta}_\tau$
$\Delta \ln y_t$	-6.508	-6.898	-6.898	0.181	0.155
$\Delta \ln \text{oil}_t$	-8.274	-8.265	-8.459	0.200	0.019
5% cv	-1.941	-2.871	-3.425	0.463	0.146
C. Tests for serial correlation and ARCH					
Variable	Q(4)	Q(12)	ARCH(4)		
$\Delta \ln y_t$	100.734 (0.000)	127.531 (0.000)	28.192 (0.000)		
$\Delta \ln \text{oil}_t$	33.913 (0.000)	50.807 (0.000)	27.187 (0.000)		
D. Engle and Ng (1993) Tests for sign and size bias in variance					
Variable	Sign	Negative size	Positive size	Joint	
$\Delta \ln y_t$	24.440 (0.013)	-7.846 (0.000)	-1.622 (0.171)	37.102 (0.000)	
$\Delta \ln \text{oil}_t$	3543.713 (0.062)	-87.300 (0.000)	-0.786 (0.963)	32.087 (0.000)	

Note: Numbers in parentheses are tail areas of tests. Annualized logarithmic first differences are used.

of normality, suggests that each of $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ fails to satisfy the null hypothesis of the test.

A battery of unit root and stationarity tests are conducted in Table 1 in $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$. In particular, we report the augmented Dickey–Fuller (ADF) test [see Dickey and Fuller (1981)] and, given that unit root tests have low power against relevant trend stationary alternatives, we also present Kwiatkowski et al. (1992) tests, known as KPSS tests, for level and trend stationarity. As can be seen, the null hypothesis of a unit root can be rejected at conventional significance levels. Moreover, the t -statistics $\hat{\eta}_\mu$ and $\hat{\eta}_\tau$ that test the null hypotheses of level and trend stationarity are small relative to their 5% critical values of 0.463 and 0.146 (respectively), given in Kwiatkowski et al. (1992). We thus conclude that $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ are stationary [integrated of order zero, or $I(0)$, in the terminology of Engle and Granger (1987)].




FIGURE 1. Logged real output and the real output growth rate.




FIGURE 2. Logged real oil price and the rate of change in the real price of oil.

In panel C of Table 1, we conduct Ljung–Box (1979) tests for serial correlation in $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$. The Q -statistics, $Q(4)$ and $Q(12)$, are asymptotically distributed as $\chi^2(36)$ on the null hypothesis of no autocorrelation. Clearly, there is significant serial dependence in the data. We also present (in the last column of panel C) Engle’s (1982) ARCH χ^2 test statistic, distributed as a $\chi^2(1)$ on the null of no ARCH. The test indicates that there is strong evidence of conditional heteroskedasticity in each of the $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ series.

Finally, as we are interested in the asymmetry of the volatility response to news, in panel D of Table 1 we present Engle and Ng (1993) tests for “sign bias,” “negative size bias,” and “positive size bias,” based on the respective regression equations

$$\widehat{\varepsilon}_t^2 = \phi_0 + \phi_1 D_{t-1}^- + \xi_t, \quad (1)$$

$$\widehat{\varepsilon}_t^2 = \phi_0 + \phi_1 D_{t-1}^- \widehat{\varepsilon}_{t-1} + \xi_t, \quad (2)$$

$$\widehat{\varepsilon}_t^2 = \phi_0 + \phi_1 D_{t-1}^+ \widehat{\varepsilon}_{t-1} + \xi_t, \quad (3)$$

where $\widehat{\varepsilon}_t$ is the residual from a fourth-order autoregression of the raw data ($\Delta \ln y_t$ or $\Delta \ln \text{oil}_t$), treated as a collective measure of news at time t , D_{t-1}^- is a dummy variable that takes a value of one when $\widehat{\varepsilon}_{t-1}$ is negative (bad news) and zero otherwise, $D_{t-1}^+ = 1 - D_{t-1}^-$, picking up the observations with positive innovations (good news), and ϕ_0 and ϕ_1 are parameters. The t -ratio of the ϕ_1 coefficient in each of the regression equations (1)–(3) is defined as the test statistic.

The sign bias test in equation (1) examines the impact that positive and negative shocks have on volatility which is not predicted by the volatility model under consideration. In particular, if the response of volatility to shocks is asymmetric (that is, positive and negative shocks to $\widehat{\varepsilon}_{t-1}$ impact differently upon the conditional variance, $\widehat{\varepsilon}_t^2$), then ϕ_1 will be statistically significant. Irrespective of whether the response of volatility to shocks is symmetric or asymmetric, the size (or magnitude) of the shock could also affect volatility. The negative size bias test in equation (2) focuses on the asymmetric effects of negative shocks (that is, whether small and large negative shocks to $\widehat{\varepsilon}_{t-1}$ impact differently upon the conditional variance, $\widehat{\varepsilon}_t^2$). In this case, D_{t-1}^- is used as a slope dummy variable in equation (2) and negative size bias is present if ϕ_1 is statistically significant. The positive size bias test in equation (3) focuses on the different effects that large and small positive shocks have on volatility, and positive size bias is present if ϕ_1 is statistically significant in (3). We also conduct a joint test for both sign and size bias using the regression equation

$$\widehat{\varepsilon}_t^2 = \phi_0 + \phi_1 D_{t-1}^- + \phi_2 D_{t-1}^- \widehat{\varepsilon}_{t-1} + \phi_3 D_{t-1}^+ \widehat{\varepsilon}_{t-1} + \xi_t. \quad (4)$$

In the joint test in equation (4), the test statistic is equal to $T \times R^2$ (where R^2 is from the regression) and follows a χ^2 distribution with three degrees of freedom under the null hypothesis of no asymmetric effects.

As can be seen in panel D of Table 1, the conditional volatility of output growth is sensitive to the sign and size of the innovation. In particular, there is strong

evidence of sign and negative size bias in the output growth volatility, and the joint test for both sign and size bias is highly significant. Also, the conditional volatility of the change in the price of oil displays negative size bias and the joint test for both sign and size bias is significant at conventional levels.

3. ECONOMETRIC METHODOLOGY

Given the evidence of conditional heteroskedasticity in the $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ series, we characterize the joint data-generating process underlying $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ as a bivariate GARCH-in-mean model, as follows:

$$y_t = a + \sum_{i=1}^p \Gamma_i y_{t-i} + \sum_{j=0}^q \Psi_j \sqrt{h_{t-j}} + e_t \tag{5}$$

$$e_t | \Omega_{t-1} \sim (\mathbf{0}, H_t), \quad H_t = \begin{bmatrix} h_{\Delta \ln y \Delta \ln y, t} & h_{\Delta \ln y \Delta \ln \text{oil}, t} \\ h_{\Delta \ln \text{oil} \Delta \ln y, t} & h_{\Delta \ln \text{oil} \Delta \ln \text{oil}, t} \end{bmatrix},$$

where $\mathbf{0}$ is the null vector, Ω_{t-1} denotes the available information set in period $t - 1$, and

$$y_t = \begin{bmatrix} \Delta \ln y_t \\ \Delta \ln \text{oil}_t \end{bmatrix}; e_t = \begin{bmatrix} e_{\Delta \ln y, t} \\ e_{\Delta \ln \text{oil}, t} \end{bmatrix}; h_t = \begin{bmatrix} h_{\Delta \ln y \Delta \ln y, t} \\ h_{\Delta \ln \text{oil} \Delta \ln \text{oil}, t} \end{bmatrix};$$

$$a = \begin{bmatrix} a_{\Delta \ln y} \\ a_{\Delta \ln \text{oil}} \end{bmatrix}; \Gamma_i = \begin{bmatrix} \gamma_{11}^{(i)} & \gamma_{12}^{(i)} \\ \gamma_{21}^{(i)} & \gamma_{22}^{(i)} \end{bmatrix}; \Psi_j = \begin{bmatrix} \psi_{11}^{(j)} & \psi_{12}^{(j)} \\ \psi_{21}^{(j)} & \psi_{22}^{(j)} \end{bmatrix}.$$

Notice that we have not added any error correction term to the model, as the null hypothesis of no cointegration between output ($\ln y_t$) and the real price of oil ($\ln \text{oil}_t$) cannot be rejected.

Multivariate GARCH models require that we specify volatilities of $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$, measured by conditional variances. Several different specifications have been proposed in the literature, including the VECM model of Bollerslev et al. (1988), the CCORR model of Bollerslev (1990), the FARCH specification of Engle et al. (1990), the BEKK model proposed by Engle and Kroner (1995), and the DCC model of Engle (2002). However, none of these specifications is capable of capturing the asymmetry of the volatility response to news.

In this regard, given the asymmetric effects of news on volatility in the $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ series, we use an asymmetric version of the BEKK model, introduced by Grier et al. (2004),

$$H_t = C' C + \sum_{j=1}^f B'_j H_{t-j} B_j + \sum_{k=1}^g A'_k e_{t-k} e'_{t-k} A_k + D' u_{t-1} u'_{t-1} D, \tag{6}$$

where C , B_j , A_k , and D are 2×2 matrices (for all values of j and k), with C being a triangular matrix to ensure positive definiteness of H . In equation (6), $u_t = (u_{\Delta \ln y,t}, u_{\Delta \ln \text{oil},t})'$ and captures potential asymmetric responses. In particular, if the change in the price of oil, $\Delta \ln \text{oil}_t$, is higher than expected, we take that to be bad news. We therefore capture bad news about oil price changes by a positive oil price change residual, by defining $u_{\Delta \ln \text{oil},t} = \max\{e_{\Delta \ln \text{oil},t}, 0\}$. We also capture bad news about output growth by defining $u_{\Delta \ln y,t} = \min\{e_{\Delta \ln y,t}, 0\}$. Hence, $u_t = (u_{\Delta \ln y,t}, u_{\Delta \ln \text{oil},t})' = (\min\{e_{\Delta \ln y,t}, 0\}, \max\{e_{\Delta \ln \text{oil},t}, 0\})'$.

The specification in equation (6) allows past volatilities, H_{t-j} , as well as lagged values of ee' and uu' to show up in estimating current volatilities of $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$. Moreover, the introduction of the uu' term in (6) extends the BEKK model by relaxing the assumption of symmetry, thereby allowing for different relative responses to positive and negative shocks in the conditional variance-covariance matrix, H .

There are $n + n^2(p + q) + n(n + 1)/2 + n^2(f + g + 1)$ parameters in (5) and (6) and to deal with estimation problems in the large parameter space, we assume that $f = g = 1$ in equation (6), consistent with recent empirical evidence regarding the superiority of GARCH(1,1) models—see, for example, Hansen and Lunde (2005). It is also to be noted that we have not included an interest rate variable in the model (in the y_t equation), although it would seem to be important as oil prices affect output through an indirect effect on the rate of interest. We have kept the dimension of the model low because of computational and degree of freedom problems in the large parameter space. For example, with $n = 2$, $p = q = 2$ in equation (5) and $f = g = 1$ in equation (6), the model has 33 parameters to be estimated. If we introduce one more variable into the model (such as the interest rate), then we would have to estimate 81 parameters. Moreover, the tests that we conduct in Section 4 indicate that the exclusion of such a variable is not expected to result in significant misspecification error.

In order to estimate our bivariate GARCH-in-mean asymmetric BEKK model, we construct the likelihood function, ignoring the constant term and assuming that the statistical innovations are conditionally Gaussian,

$$l_t = -\frac{1}{2} \sum_{t=t+1}^T \log |H_t| - \frac{1}{2} \sum_{t=t+1}^T (e_t' H_t^{-1} e_t),$$

where e_t and H_t are evaluated at their estimates. The log-likelihood is maximized with respect to the parameters Γ_i ($i = 1, \dots, p$), Ψ_j ($j = 1, \dots, q$), C , B , A , and D . As we are using the BEKK model, we do not need to impose any restrictions on the variance parameters to make H_t positive definite. Moreover, we are estimating all the parameters simultaneously rather than estimating mean and variance parameters separately, thus avoiding the Lee et al. (1995) problem of generated regressors.

4. EMPIRICAL EVIDENCE

Initially we used the AIC and SIC criteria to select the optimal values of p and q in (5). However, because of computational difficulties and remaining serial correlation and ARCH effects in the standardized residuals, we set $p = 3$ and $q = 1$ in equation (5). Hence, with $p = 3$ and $q = 1$ in equation (5), and $f = g = 1$ in equation (6), we estimate a total of 37 parameters. Maximum likelihood (ML) estimates of the parameters and diagnostic test statistics are presented in Tables 2 and 3.

We conducted a battery of misspecification tests, using robustified versions of the standard test statistics based on the standardized residuals

$$z_i = \frac{e_{i,t}}{\sqrt{\hat{h}_{ij,t}}}, \quad \text{for } i, j = \Delta \ln y, \Delta \ln \text{ oil}.$$

As shown in panel A of Table 2, the Ljung–Box Q -statistics for testing serial correlation cannot reject the null hypothesis of no autocorrelation (at conventional significance levels) for the values and the squared values of the standardized residuals, suggesting that there is no evidence of conditional heteroskedasticity. Moreover, the failure of the data to reject the null hypotheses of $E(z) = 0$ and $E(z^2) = 1$ implicitly indicates that our bivariate asymmetric GARCH-in-mean model does not bear significant misspecification error—see, for example, Kroner and Ng (1998).

In Table 3, we also present diagnostic tests suggested by Engle and Ng (1993) and Kroner and Ng (1998), based on the “generalized residuals,” defined as $e_{i,t}e_{j,t} - h_{ij,t}$ for $i, j = \Delta \ln y, \Delta \ln \text{ oil}$. For all symmetric GARCH models, the news impact curve—see Engle and Ng (1993)—is symmetric and centered at $e_{i,t-1} = 0$. A generalized residual can be thought of as the distance between a point on the scatterplot of $e_{i,t}e_{j,t}$ and a corresponding point on the news impact curve. If the conditional heteroskedasticity part of the model is correct, $E_{t-1}(e_{i,t}e_{j,t} - h_{ij,t}) = 0$ for all values of i and j , generalized residuals should be uncorrelated with all information known at time $t - 1$. In other words, the unconditional expectation of $e_{i,t}e_{j,t}$ should be equal to its conditional one, $h_{ij,t}$.

The Engle and Ng (1993) and Kroner and Ng (1998) misspecification indicators test whether we can predict the generalized residuals by some variables observed in the past, but that are not included in the model—this is exactly the intuition behind $E_{t-1}(e_{i,t}e_{j,t} - h_{ij,t}) = 0$. In this regard, we follow Kroner and Ng (1998) and Shields et al. (2005) and define two sets of misspecification indicators. In a two-dimensional space, we first partition $(e_{\Delta \ln y,t-1}, e_{\Delta \ln \text{ oil},t-1})$ into four quadrants in terms of the possible signs of the two residuals. Then, to shed light on any possible sign bias of the model, we define the first set of indicator functions as $I(e_{\Delta \ln y,t-1} < 0)$, $I(e_{\Delta \ln \text{ oil},t-1} < 0)$, $I(e_{\Delta \ln y,t-1} < 0; e_{\Delta \ln \text{ oil},t-1} < 0)$, $I(e_{\Delta \ln y,t-1} > 0; e_{\Delta \ln \text{ oil},t-1} < 0)$, $I(e_{\Delta \ln y,t-1} < 0; e_{\Delta \ln \text{ oil},t-1} > 0)$, and $I(e_{\Delta \ln y,t-1} > 0; e_{\Delta \ln \text{ oil},t-1} > 0)$, where $I(\cdot)$ equals one if the argument is true and zero otherwise. Significance of any of these indicator functions indicates

TABLE 2. The bivariate GARCH-in-mean asymmetric BEKK model

A. Conditional mean equation						
$\mathbf{a} = \begin{bmatrix} 0.725 \\ (0.895) \\ -47.680 \\ (0.005) \end{bmatrix}; \Gamma_1 = \begin{bmatrix} 0.213 & 0.002 \\ (0.000) & (0.614) \\ 0.528 & 0.174 \\ (0.303) & (0.000) \end{bmatrix}; \Gamma_2 = \begin{bmatrix} 0.263 & -0.006 \\ (0.000) & (0.027) \\ 1.633 & -0.100 \\ (0.000) & (0.112) \end{bmatrix}; \Gamma_3 = \begin{bmatrix} 0.179 & -0.002 \\ (0.163) & (0.824) \\ -0.276 & 0.043 \\ (0.512) & (0.380) \end{bmatrix};$						
$\Psi_1 = \begin{bmatrix} 0.144 & -0.071 \\ (0.000) & (0.005) \\ -1.735 & 0.276 \\ (0.309) & (0.726) \end{bmatrix}; \Psi_2 = \begin{bmatrix} 0.046 & 0.056 \\ (0.941) & (0.000) \\ 3.210 & 0.123 \\ (0.205) & (0.855) \end{bmatrix}.$						
Residual diagnostics						
	Mean	Variance	$Q(4)$	$Q^2(4)$	$Q(12)$	$Q^2(12)$
z_{y_t}	-0.007 (0.892)	0.994 (0.997)	1.139 (0.887)	2.701 (0.608)	15.358 (0.222)	12.029 (0.443)
z_{oil_t}	-0.020 (0.698)	1.003 (1.001)	7.016 (0.135)	2.650 (0.617)	13.517 (0.332)	8.868 (0.714)

TABLE 2. Continued

B. Conditional variance–covariance structure

$$C = \begin{bmatrix} 5.400 & -3.338 \\ (0.000) & (0.435) \\ & 22.662 \\ & (0.001) \end{bmatrix}; B = \begin{bmatrix} 0.002 & 1.213 \\ (0.994) & (0.354) \\ & 0.001 & 0.807 \\ & (0.919) & (0.000) \end{bmatrix};$$

$$A = \begin{bmatrix} 0.393 & 0.298 \\ (0.277) & (0.842) \\ & 0.018 & 0.546 \\ & (0.133) & (0.000) \end{bmatrix}; D = \begin{bmatrix} -0.898 & 1.491 \\ (0.000) & (0.093) \\ & -0.016 & 0.037 \\ & (0.335) & (0.752) \end{bmatrix}.$$

Hypothesis testing

Diagonal VAR	$H_0 : \gamma_{12}^{(i)} = \gamma_{21}^{(i)} = 0, \text{ for } i = 1, 2, 3$	(0.000)
No GARCH	$H_0 : \alpha_{ij} = \beta_{ij} = \delta_{ij} = 0, \text{ for all } i, j$	(0.000)
No GARCH-M	$H_0 : \psi_{ij}^k = 0, \text{ for all } i, j, k$	(0.000)
No asymmetry	$H_0 : \delta_{ij} = 0, \text{ for } i, j = 1, 2$	(0.000)
Diagonal GARCH	$H_0 : \alpha_{12} = \alpha_{21} = \beta_{12} = \beta_{21} = \delta_{12} = \delta_{21} = 0$	(0.000)

Note: Model: Equations (2) and (3) with $p = 3, q = 1,$ and $f = g = 1.$ Numbers in parentheses are tail areas of tests.

TABLE 3. Diagnostic tests based on the news impact curve

	$e_{\Delta \ln y, t}^2 - h_{\Delta \ln y \Delta \ln y, t}$	$e_{\Delta \ln y, t} e_{\Delta \ln oil, t} - h_{\Delta \ln y \Delta \ln oil, t}$	$e_{\Delta \ln oil, t}^2 - h_{\Delta \ln oil \Delta \ln oil, t}$
$I(e_{\Delta \ln y, t-1} < 0)$	1.107 (0.292)	0.337 (0.561)	2.441 (0.118)
$I(e_{\Delta \ln oil, t-1} < 0)$	0.027 (0.867)	0.620 (0.430)	1.639 (0.200)
$I(e_{\Delta \ln y, t-1} < 0, e_{\Delta \ln oil, t-1} < 0)$	31.178 (0.000)	5.890 (0.015)	0.772 (0.379)
$I(e_{\Delta \ln y, t-1} > 0, e_{\Delta \ln oil, t-1} < 0)$	18.607 (0.000)	12.271 (0.000)	3.187 (0.074)
$I(e_{\Delta \ln y, t-1} < 0, e_{\Delta \ln oil, t-1} > 0)$	1.414 (0.234)	20.550 (0.000)	8.859 (0.003)
$I(e_{\Delta \ln y, t-1} > 0, e_{\Delta \ln oil, t-1} > 0)$	1.322 (0.250)	3.452 (0.063)	1.205 (0.272)
$e_{\Delta \ln y, t-1}^2 I(e_{\Delta \ln y, t-1} < 0)$	0.392 (0.530)	1.954 (0.162)	2.953 (0.085)
$e_{\Delta \ln y, t-1}^2 I(e_{\Delta \ln oil, t-1} < 0)$	0.688 (0.407)	5.476 (0.019)	10.516 (0.001)
$e_{\Delta \ln oil, t-1}^2 I(e_{\Delta \ln y, t-1} < 0)$	0.350 (0.553)	0.514 (0.473)	0.009 (0.923)
$e_{\Delta \ln oil, t-1}^2 I(e_{\Delta \ln oil, t-1} < 0)$	0.153 (0.695)	2.536 (0.111)	1.826 (0.176)

Note: Numbers in parentheses are tail areas of tests.

that the model (5)–(6) is incapable of predicting the effects of some shocks to either $\Delta \ln y_t$ or $\Delta \ln oil_t$. Moreover, because the possible effect of a shock could be a function of both the size and the sign of the shock, we define a second set of indicator functions, $e_{\Delta \ln y, t-1}^2 I(e_{\Delta \ln y, t-1} < 0)$, $e_{\Delta \ln y, t-1}^2 I(e_{\Delta \ln y, t-1} > 0)$, $e_{\Delta \ln oil, t-1}^2 I(e_{\Delta \ln y, t-1} < 0)$, and $e_{\Delta \ln oil, t-1}^2 I(e_{\Delta \ln oil, t-1} < 0)$. These indicators are technically scaled versions of the former ones, with the magnitude of the shocks as a scale measure.

We conducted indicator tests and we report the results in Table 3. As can be seen in Table 3, most of the indicators fail to reject the null hypothesis of no misspecification—all test statistics in Table 3 are distributed as $\chi^2(1)$. Hence, our model (5)–(6) captures the effects of all sign bias and sign size scale-dependent shocks in predicting volatility, and there is no significant misspecification error. This means that the exclusion of the interest rate variable (in y_t), mentioned earlier, is not expected to lead to significant misspecification problems.

Turning now back to panel B of Table 2, the diagonality restriction, $\gamma_{12}^{(i)} = \gamma_{21}^{(i)} = 0$ for $i = 1, 2, 3$, is rejected, meaning that the data provide strong evidence of the existence of dynamic interactions between $\Delta \ln y_t$ and $\Delta \ln oil_t$. The null hypothesis of homoskedastic disturbances requires the **A**, **B**, and **D** matrices to be jointly nonsignificant (that is, $\alpha_{ij} = \beta_{ij} = \delta_{ij} = 0$ for all i, j) and is rejected at the 1% level or better, suggesting that there is significant conditional heteroskedasticity in the data. The null hypothesis of symmetric conditional variance–covariances, which requires all elements of the **D** matrix to be jointly insignificant (that is, $\delta_{ij} = 0$ for all i, j), is rejected at the 1% level or better, implying the existence of some asymmetries in the data that the model is capable of capturing. Also, the null

hypothesis of a diagonal covariance process requires the off-diagonal elements of the A , B , and D matrices to be jointly insignificant (that is, $\alpha_{12} = \alpha_{21} = \beta_{12} = \beta_{21} = \delta_{12} = \delta_{21} = 0$), but these estimated coefficients are jointly significant at the 1% level of significance.

Thus the $\Delta \ln y_t - \Delta \ln \text{oil}_t$ process is strongly conditionally heteroskedastic, with innovations in oil price changes significantly influencing the conditional variance of output growth in an asymmetric way. Moreover, the sign as well as the size of oil price change innovations are important. To establish the relationship between volatility in the change in the price of oil and output growth, in Table 2 we test the null hypothesis that the volatility of $\Delta \ln \text{oil}_t$ does not (Granger) cause output growth, $H_0 : \psi_{12} = 0$. We strongly reject the null hypothesis, finding strong evidence in support of the hypothesis that $\Delta \ln \text{oil}_t$ volatility Granger-causes output growth.

In Figures 3–5 we plot the conditional standard deviations of output growth and the change in the price of oil, as well as the conditional covariance implied by our estimates of the asymmetric VAR-BEKK model in Table 2. In Figure 3, the biggest episode of output growth volatility coincides with the 2009 NBER recession—the largest recession in the sample. Regarding the change in the real price of oil, $\Delta \ln \text{oil}_t$, Figure 4 shows that the largest episodes of oil price change volatility took place in 1986, 1990, and 2009. All of these volatility jumps in $\Delta \ln \text{oil}_t$ do not coincide with NBER recessions, but the relatively higher volatility jump in the oil price change in 2009 coincides with the largest recession in the sample. The volatility jumps in 1986 and 1999 are the results of steep oil price declines rather than oil price increases. Finally, the conditional covariance between $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$, shown in Figure 5, is highest in 1986, 1991, and 2009.

5. GENERALIZED IMPULSE RESPONSE FUNCTIONS

As van Dijk et al. (2007) recently put it, “it generally is difficult, if not impossible, to fully understand and interpret nonlinear time series models by considering the estimated values of the model parameters only.” Thus, in order to quantify the dynamic response of output growth and oil price changes to shocks and to investigate the statistical significance of the asymmetry in the variance–covariance structure, we calculate GIRFs, introduced by Koop et al. (1996) and recently used by Grier et al. (2004), based on our bivariate, GARCH-in-mean, asymmetric VAR-BEKK model (5)–(6).

Traditional impulse-response functions, which are more usefully applied to linear models than to nonlinear ones, measure the effect of a shock (say of size δ) hitting the system at time t on the state of the system at time $t + n$, given that no other shocks hit the system. As Koop et al. (1996, p. 121) put it, “the idea is very similar to Keynesian multiplier analysis, with the difference that the analysis is carried out with respect to shocks or ‘innovations’ of macroeconomic time series, rather than the series themselves (such as investment or government expenditure).” In the case of multivariate nonlinear models, however, traditional




FIGURE 3. Conditional standard deviation of output growth.

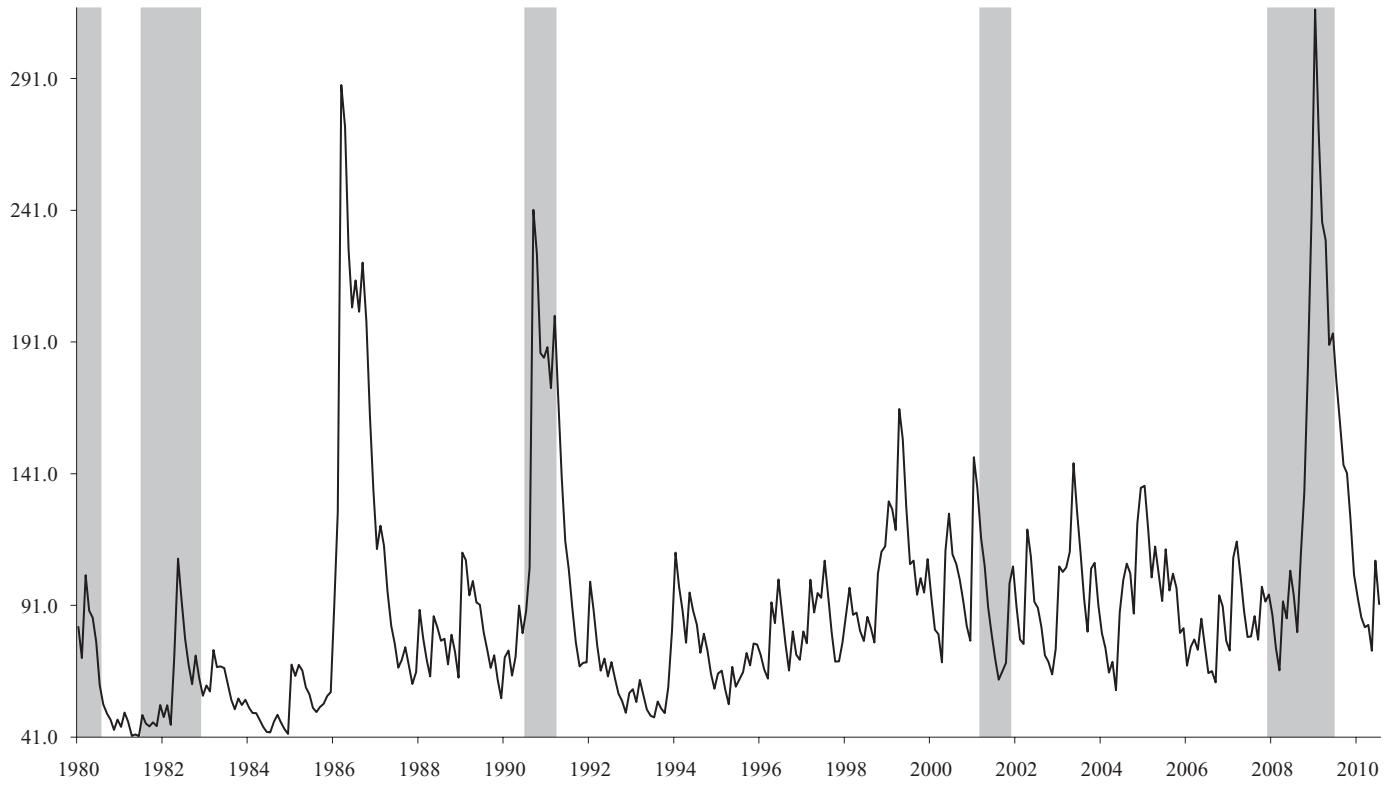


FIGURE 4. Conditional standard deviation of the change in the price of oil.




FIGURE 5. Covariance between output growth and the change in the price of oil.

impulse-response functions depend on the sign and size of the shock as well as the history of the system (i.e., expansionary or contractionary) before the shock hits—see, for example, Potter (2000).

In our asymmetric bivariate, GARCH-in-mean, VAR-BEKK model, shocks impact on output growth and the change in the price of oil through the conditional mean as described in equation (5) and with lags through the conditional variance as described in equation (6). Moreover, the impulse responses of $\Delta \ln y_t$ and $\Delta \ln \text{Oil}_t$ depend on the composition of the $e_{\Delta \ln y_t}$ and $e_{\Delta \ln \text{Oil}_t}$ shocks—that is, the effect of a shock to $\Delta \ln \text{Oil}_t$ is not isolated from having a contemporaneous effect on $\Delta \ln y_t$ and vice versa. The GIRFs that we use in this paper provide a method of dealing with the problems of shock, history, and composition dependence of impulse responses in multivariate (linear and) nonlinear models.

In particular, assuming that \mathbf{y}_t is a random vector, the GIRF for an arbitrary current shock, \mathbf{v}_t , and history, ω_{t-1} , is defined as

$$\text{GIRF}_{\mathbf{y}}(n, \mathbf{v}_t, \omega_{t-1}) = E[\mathbf{y}_{t+n} | \mathbf{v}_t, \omega_{t-1}] - E[\mathbf{y}_{t+n} | \omega_{t-1}] \quad (7)$$

for $n = 0, 1, 2, \dots$. If \mathbf{v}_t and ω_{t-1} are realizations of the random variables \mathbf{V}_t and Ω_{t-1} (where Ω_{t-1} is the set containing information used to forecast \mathbf{y}_t) that generate realizations of $\{\mathbf{y}_t\}$, then according to Koop et al. (1996), the GIRF in (7) can be considered to be a realization of a random variable defined by

$$\text{GIRF}_{\mathbf{y}}(n, \mathbf{V}_t, \Omega_{t-1}) = E[\mathbf{y}_{t+n} | \mathbf{V}_t, \Omega_{t-1}] - E[\mathbf{y}_{t+n} | \Omega_{t-1}]. \quad (8)$$

Equation (8) is the difference between two conditional expectations, $E[\mathbf{y}_{t+n} | \mathbf{V}_t, \Omega_{t-1}]$ and $E[\mathbf{y}_{t+n} | \Omega_{t-1}]$. It is also to be noted that generalized impulse-responses are expected values (which are not random values and will not have error bands).

The computation of GIRFs in the case of multivariate nonlinear models is made difficult by the inability to construct analytical expressions for the conditional expectations $E[\mathbf{y}_{t+n} | \mathbf{V}_t, \Omega_{t-1}]$ and $E[\mathbf{y}_{t+n} | \Omega_{t-1}]$ in equation (8). To deal with this problem, Monte Carlo methods of stochastic simulation are used to construct the GIRFs. Here, we allow for time-varying composition dependence and follow the algorithm described in Koop et al. (1996). In particular, using 310 data points as histories, we first transform the estimated residuals by using the variance–covariance structure and Jordan decomposition. Then at each history, 50 realizations are drawn randomly, thereby obtaining identical and independent distributions over time. Recovering the time-varying dependence among the residuals, 15,500 realizations of impulse responses are calculated for each horizon. Finally, the whole process is replicated 150 times to average out the effects of impulses.

The GIRFs to an average shock in $\Delta \ln y_t$ and $\Delta \ln \text{Oil}_t$ are shown in Figures 6 and 7—they show the effect on $\Delta \ln y_t$ and $\Delta \ln \text{Oil}_t$ of an average initial shock in $\Delta \ln y_t$ and $\Delta \ln \text{Oil}_t$. As can be seen, none of the shocks is very persistent, although the effect of the shock to $\Delta \ln \text{Oil}_t$ on $\Delta \ln y_t$ is more persistent than




FIGURE 6. GIRF of output growth to an oil price change shock.




FIGURE 7. GIRF of the oil price change to an output growth shock.

the effect of the shock to $\Delta \ln y_t$ on $\Delta \ln \text{oil}_t$, as it takes longer for $\Delta \ln y_t$ to return to its original value. Shocks to output growth and the change in the price of oil provide a large stimulus to $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ for the first few months. In particular, in response to an oil price change shock, output growth declines by more than 0.75% in the first quarter of the year and returns to its mean within one and one-half years. Also the change in the price of oil responds very strongly (almost 12%) to the innovation in output growth within the first few months of the year.

In Figures 8 and 9, we differentiate between positive and negative shocks, in order to address issues regarding the asymmetry of shocks. As can be seen in Figure 8, output growth declines because of a positive $\Delta \ln \text{oil}_t$ shock and increases in response to a negative $\Delta \ln \text{oil}_t$ shock. The responses are not mirror images of each other, suggesting that output growth, $\Delta \ln y_t$, responds asymmetrically to shocks in the change in the price of oil, $\Delta \ln \text{oil}_t$. Also, the response of output growth to a negative $\Delta \ln \text{oil}_t$ shock returns to zero faster than to a positive shock of equal magnitude, suggesting that positive shocks in the change in the price of oil have more persistent effects on output growth than negative ones. Figure 9 shows the GIRFs of $\Delta \ln \text{oil}_t$ to positive and negative output growth shocks. A positive $\Delta \ln y_t$ shock raises $\Delta \ln \text{oil}_t$ and a negative $\Delta \ln y_t$ shock lowers $\Delta \ln \text{oil}_t$. The response of $\Delta \ln \text{oil}_t$ to output growth shocks is also asymmetric—a positive output growth shock has a larger effect on the change in the price of oil compared to a negative $\Delta \ln y_t$ shock.

Given the asymmetric nature of the specification of our bivariate asymmetric VAR-BEKK model, we follow Van Dijk et al. (2007) and use the GIRFs to positive and negative shocks to compute a random asymmetry measure, defined as follows,

$$\text{ASY}_y(n, \mathbf{V}_t^+, \Omega_{t-1}) = \text{GIRF}_y(n, \mathbf{V}_t^+, \Omega_{t-1}) + \text{GIRF}_y(n, -\mathbf{V}_t^+, \Omega_{t-1}), \quad (9)$$

where $\text{GIRF}_y(n, \mathbf{V}_t^+, \Omega_{t-1})$ denotes the GIRF derived from conditioning on the set of all possible positive shocks, $\text{GIRF}_y(n, -\mathbf{V}_t^+, \Omega_{t-1})$ denotes the GIRF derived from conditioning on the set of all possible negative shocks, and $\mathbf{V}_t^+ = \{v_t | v_t > 0\}$. The distribution of $\text{ASY}(n, \mathbf{V}_t^+, \Omega_{t-1})$ can provide an indication of the asymmetric effects of positive and negative shocks. In particular, if $\text{ASY}(n, \mathbf{V}_t^+, \Omega_{t-1})$ has a symmetric distribution with a mean of zero, then positive and negative shocks have exactly the same effect (with opposite sign).

We have computed the asymmetry measures for $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ and show the distributions of the respective $\text{ASY}_y(n, \mathbf{V}_t^+, \Omega_{t-1})$ measures in Figures 10 and 11 at horizons $n = 6$, $n = 9$, and $n = 12$. As can be seen in Figure 10, on average output growth exhibits more persistence to a positive $\Delta \ln \text{oil}_t$ shock than to a negative one. In particular, the loss of output growth due to a positive $\Delta \ln \text{oil}_t$ shock (bad news) at horizon $n = 9$ is 0.083% in excess of the gain in output growth from a negative $\Delta \ln \text{oil}_t$ shock (good news) of equal magnitude. Figure 11 shows the asymmetry measure for an output growth shock on $\Delta \ln \text{oil}_t$. We find a stronger effect of a positive output growth shock on the change in the

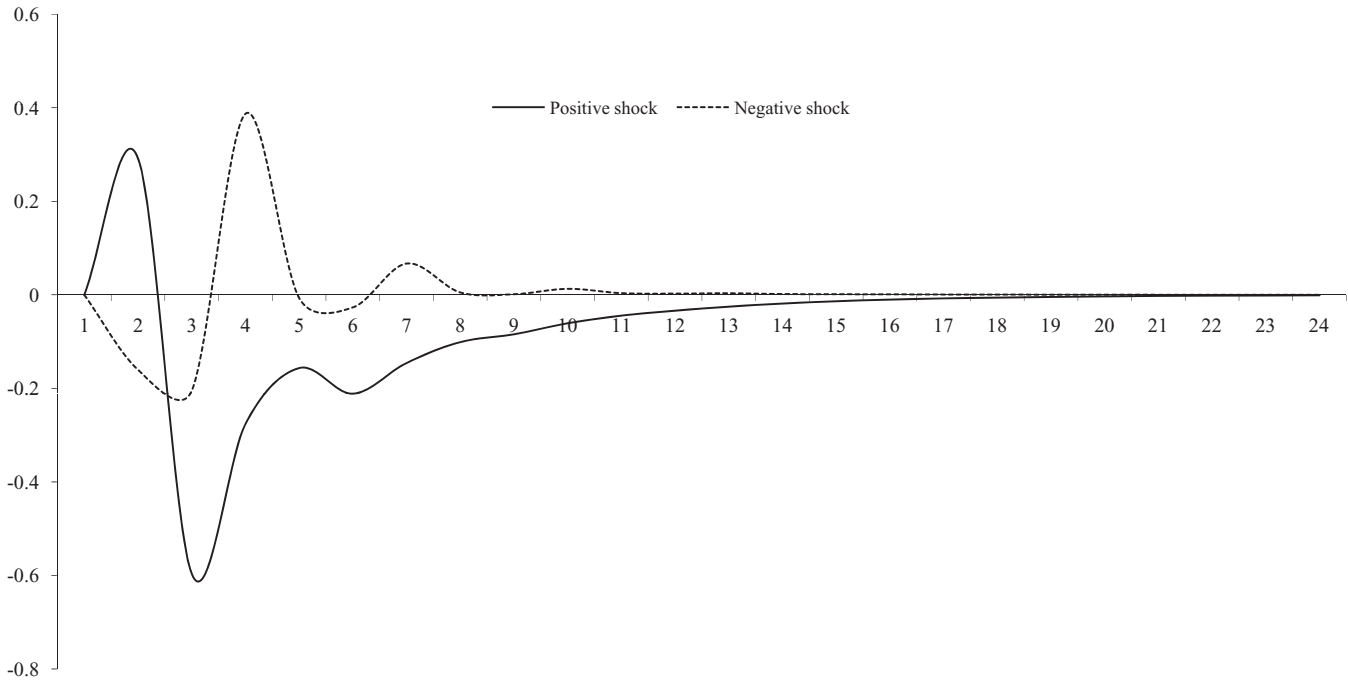


FIGURE 8. GIRFs of output growth to positive and negative shocks to the change in the price of oil.




FIGURE 9. GIRFs of the change in the price of oil to positive and negative output growth shocks.




FIGURE 10. The distribution of the asymmetry measure based on the GIRFs of output growth to shocks in the change in the price of oil.




FIGURE 11. The distribution of the asymmetry measure based on the GIRFs of the change in the price of oil to shocks in output growth.

price of oil than of a negative shock of equal magnitude. On average at horizon 9, the increase in $\Delta \ln \text{oil}_t$ due to a positive output growth shock is 0.256% in excess of the decrease in $\Delta \ln \text{oil}_t$ due to a negative $\Delta \ln y_t$ shock.

6. VOLATILITY IMPULSE RESPONSE FUNCTIONS

The GIRFs, introduced by Koop et al. (1996), trace the effects of independent shocks (or news) on the conditional mean. Recently, Hafner and Herwartz (2006) introduced a new concept of impulse response functions, known as “volatility impulse response functions” (VIRFs), tracing the effects of independent shocks on the conditional variance—see also Shields et al. (2005) for an early application of the Hafner and Herwartz (2006) VIRFs concept.

We start with the conditional variance–covariance matrix of e_t , H_t , and define $Q_t = \text{vech}(H_t)$ to be a 3×1 random vector with the following elements (in that order): $h_{\Delta \ln y,t}$, $h_{\Delta \ln y \Delta \ln \text{oil},t}$, $h_{\Delta \ln \text{oil},t}$. Then the VIRFs of Q_t , for $n = 0, 1, \dots$, are given by

$$\text{VIRF}_Q(n, v_t, \omega_{t-1}) = E[Q_{t+n}|v_t, \omega_{t-1}] - E[Q_{t+n}|\omega_{t-1}]. \tag{10}$$

Hence, the VIRF is conditional on the initial shock and history, v_t and ω_{t-1} , and constructs the response by averaging out future innovations given the past and present. Following Koop et al. (1996) and assuming that v_t and ω_{t-1} are realizations of the random variables V_t and Ω_{t-1} that generate realizations of $\{Q_t\}$, the VIRF in (10) can be considered to be a realization of a random variable given by

$$\text{VIRF}_Q(n, V_t, \Omega_{t-1}) = E[Q_{t+n}|V_t, \Omega_{t-1}] - E[Q_{t+n}|\Omega_{t-1}].$$

As already noted, the first element of $\text{VIRF}_Q(n, V_t, \Omega_{t-1})$ gives the impulse response of the conditional variance of $\Delta \ln y_t$, $h_{\Delta \ln y,t}$, the second that of the conditional covariance, $h_{\Delta \ln y \Delta \ln \text{oil},t}$, and the third that of the conditional variance of $\Delta \ln \text{oil}_t$, $h_{\Delta \ln \text{oil},t}$. It should also be noted that in contrast to the GIRFs where positive and negative shocks produce opposite effects, VIRFs consist of the same (positive sign) effect irrespective of the sign of the shocks. Also, shock linearity does not hold in the case of VIRFs. Finally, unlike traditional impulse responses that do not depend on history, VIRFs depend on history through the conditional variance–covariance matrix at time $t = 0$ when the innovation occurs.

Using an analytic version of the VIRF, as described in Hafner and Herwartz (2006), we show the VIRFs to shocks in $\Delta \ln y_t$ and $\Delta \ln \text{oil}_t$ in Figures 12 and 13. In Figure 12 (the responses of oil price change volatility are on the secondary y-axis), shocks to the change in the price of oil, $\Delta \ln \text{oil}_t$, produce much higher responses in the conditional variance of the change in the price of oil, $h_{\Delta \ln \text{oil},t}$, than




FIGURE 12. Volatility responses to oil price change shocks.

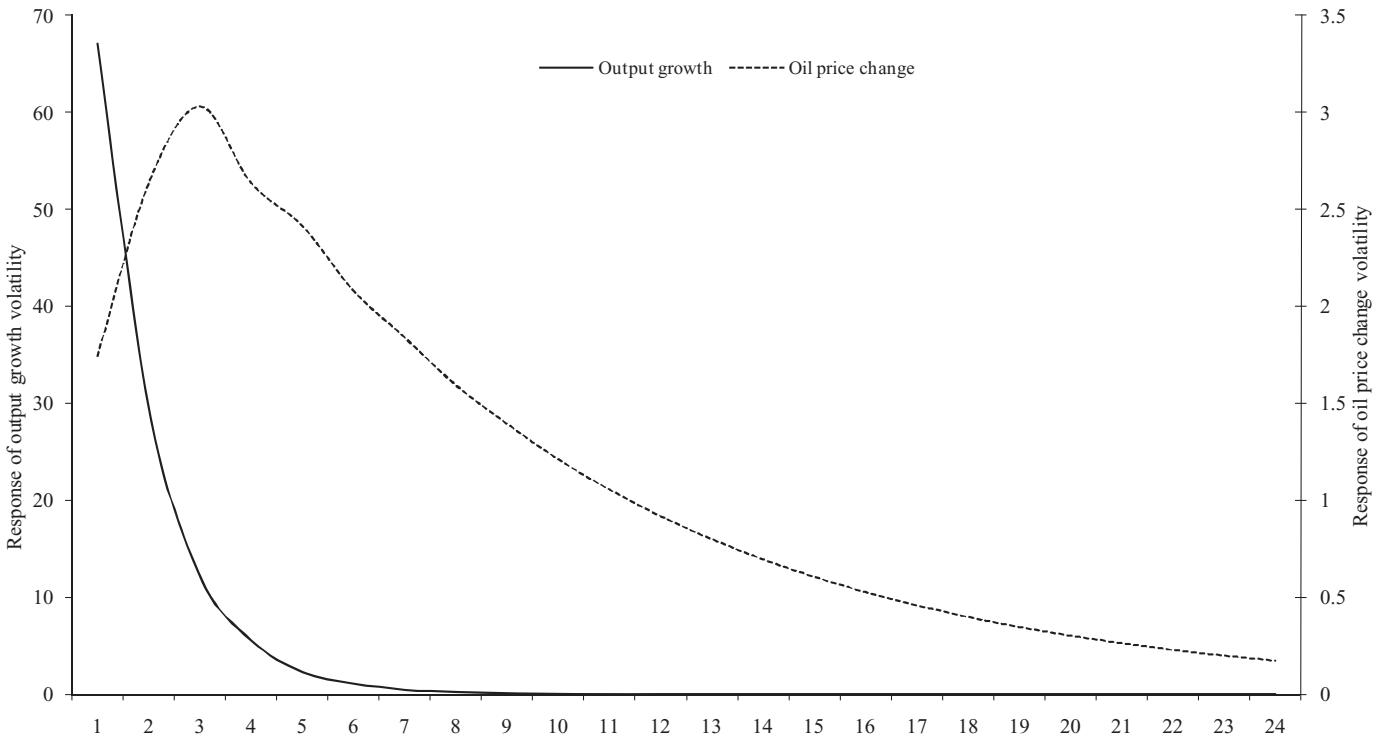


FIGURE 13. Volatility responses to output growth shocks.

in the conditional variance of output growth, $h_{\Delta \ln y,t}$. Moreover, the responses of both oil price change and output growth volatility are persistent—they take more than two years to return to zero. In Figure 13 (again the responses of oil price change volatility are on the secondary y-axis), shocks to output growth also have a very remarkable impact on the conditional variance of the change in the price of oil, $h_{\Delta \ln oil,t}$, and the conditional variance of output growth, $h_{\Delta \ln y,t}$. The peak response of the change in the price of oil is much lower than that of output growth volatility and $h_{\Delta \ln oil,t}$ takes longer to return to its original position compared to $h_{\Delta \ln y,t}$. In Figure 14, we distinguish again between positive and negative shocks in order to investigate the effect of the asymmetry of shocks on the conditional variance of the change in the real price of oil. As is shown, negative oil price shocks have a larger impact on oil price change volatility than positive oil price shocks do.

As with the GIRFs, we use the VIRFs to positive and negative innovations to compute the random asymmetry measure

$$ASY_Q(n, V_t^+, \Omega_{t-1}) = VIRF_Q(n, V_t^+, \Omega_{t-1}) - VIRF_Q(n, -V_t^+, \Omega_{t-1}), \quad (11)$$

where $VIRF_Q(n, V_t^+, \Omega_{t-1})$ denotes the VIRF derived from conditioning on the set of all possible positive innovations, $VIRF_Q(n, -V_t^+, \Omega_{t-1})$ denotes the VIRF derived from conditioning on the set of all possible negative innovations, and $V_t^+ = \{v_t | v_t > 0\}$. The distribution of $ASY_Q(n, V_t^+, \Omega_{t-1})$ will be centered at zero if positive and negative shocks have exactly the same effect. The difference between the random asymmetry measures (9) and (11) is that the latter is the difference between $VIRF_Q(n, V_t^+, \Omega_{t-1})$ and $VIRF_Q(n, -V_t^+, \Omega_{t-1})$, whereas the former is the sum of $GIRF_y(n, V_t^+, \Omega_{t-1})$ and $GIRF_y(n, -V_t^+, \Omega_{t-1})$. This is because unlike the GIRFs, where positive and negative shocks cause the response functions to take opposite signs, the VIRFs are made up of the squares of the innovations and are thus of the same sign.

We have computed the asymmetry measures $ASY_Q(n, V_t^+, \Omega_{t-1})$ and show the distributions of these measures in Figures 15 and 16 at horizon $n = 3, 6, 9$. The distribution in Figure 15 indicates that on the average, negative shocks to the change in the price of oil have more persistent effects on oil price change volatility than positive shocks do. The asymmetry measure for an oil price change shock to oil price change volatility at horizon $n = 3$ is -4.13% . Given the negative relationship between oil price change volatility and output growth, this asymmetric response of oil price change volatility tends to further weaken the short run response of output growth to lower oil prices. This result further helps in explaining the asymmetric response of output growth to oil price shocks. Also, in Figure 16, positive shocks to the change in the price of oil have more persistent effects on the volatility of output growth than negative shocks do. The asymmetry measure for oil price change shocks to the volatility in the change in the price of oil at horizon $n = 3$ is 0.13% .

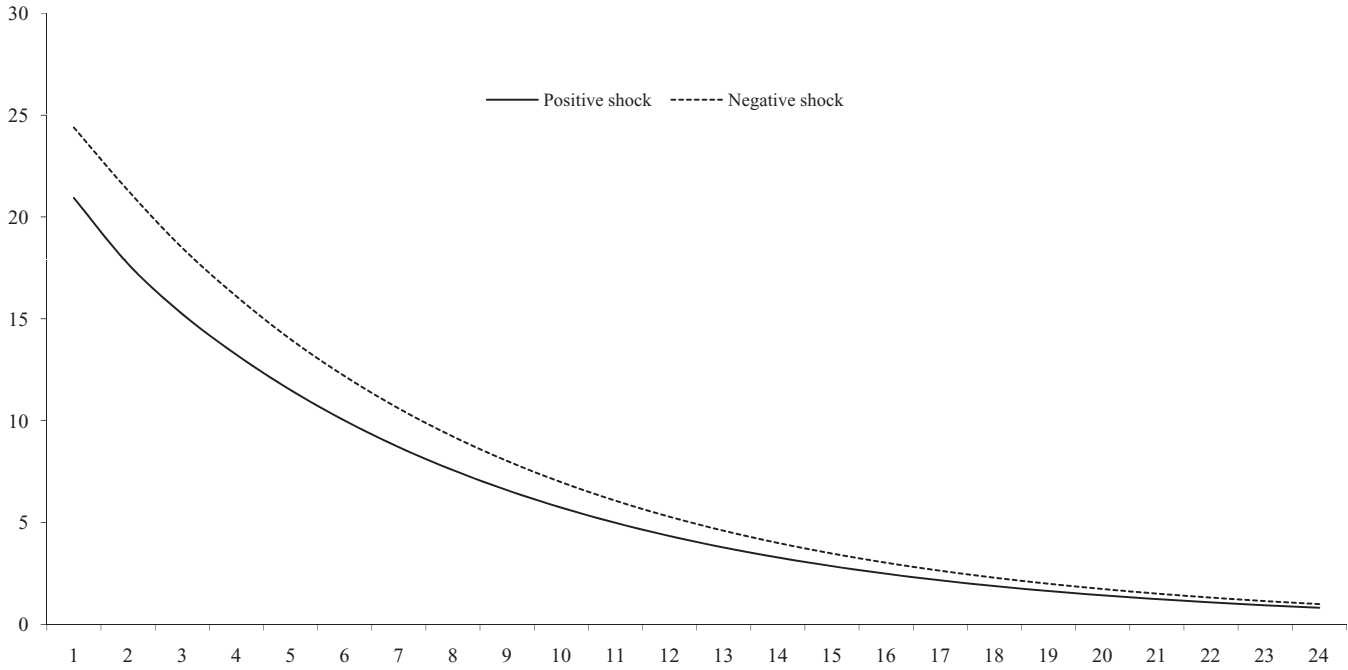


FIGURE 14. VIRFs of oil price change to positive and negative shocks to the change in the price of oil.




FIGURE 15. The distribution of the asymmetry measure based on the VIRFs of the change in price of oil to shocks in the change in the price of oil.




FIGURE 16. The distribution of the asymmetry measure based on the VIRFs of the change in the price of oil to shocks in output growth.

7. CONCLUSION

Recent empirical research regarding the relationship between the price of oil and real economic activity has focused on the role of uncertainty about oil prices—see, for example, Elder and Serletis (2010). In this paper, we examine the effects of oil price uncertainty and its asymmetry on real economic activity in the United States, in the context of a vector autoregression, in output growth and the change in the real price of oil, because the identification of higher-order VARs is usually highly questionable. In doing so, we modify our bivariate VAR to accommodate asymmetric GARCH-in-mean errors.

Our model is extremely general and allows for the possibilities of spillovers and asymmetries in the variance–covariance structure for real output growth and the change in the real price of oil. Our measure of oil price uncertainty is the conditional variance of the oil price change forecast error. We isolate the effects of volatility in the change in the price of oil and its asymmetry on output growth and, following Koop et al. (1996), Hafner and Herwartz (2006), and van Dijk et al. (2007), we employ simulation methods to calculate GIRFs and VIRFs to trace the effects of independent shocks on the conditional mean and the conditional variance, respectively, of output growth and the change in the real price of oil.

We find that our bivariate, GARCH-in-mean, asymmetric BEKK model embodies a reasonable description of the U.S. data on output growth and the change in the real price of oil. We show that the conditional variance–covariance process underlying output growth and the change in the real price of oil exhibits significant nondiagonality and asymmetry. We present evidence that increased uncertainty about the change in the real price of oil is associated with a lower average growth rate of real economic activity. Generalized impulse response experiments highlight the asymmetric effects of positive and negative shocks in the change in the real price of oil to output growth. Also, volatility impulse response experiments reveal that the effect of good news (negative shocks to the change in the real price of oil) on the conditional variance of the change in the real price of oil differs in magnitude and persistence from that of bad news of similar magnitude. This helps further in explaining the asymmetric response of output growth to oil price change shocks.

REFERENCES

- 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.
- Bollerslev, T. (1990) Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *Review of Economics and Statistics* 72, 498–505.
- Bollerslev, T., Robert F. Engle, and J. Wooldridge (1988) A capital asset pricing model with time varying covariances. *Journal of Political Economy* 96, 143–172.
- Dickey, David A. and Wayne A. Fuller (1981) Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* 49, 1057–1072.

- 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. *The B.E. Journal of Macroeconomics* (Contributions) 7, Article 35.
- Edelstein, Paul and Lutz Kilian (2008) Retail Energy Prices and Consumer Expenditures. University of Michigan and CEPR working paper.
- 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.
- Engle, Robert F. (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1008.
- Engle, Robert F. (2002) Dynamic conditional correlation: A simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R.F. and C.W. Granger (1987) Cointegration and error correction: Representation, estimation and testing. *Econometrica* 55, 251–276.
- Engle, Robert F. and Kenneth F. Kroner (1995) Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- Engle, Robert F. and Victor K. Ng (1993) Measuring and testing the impact of news on volatility. *Journal of Finance* 5, 1749–1778.
- Engle, Robert F., Victor K. Ng, and M. Rothschild (1990) Asset pricing with a factor-ARCH covariance structure: Empirical estimates for Treasury bills. *Journal of Econometrics* 45, 213–237.
- Grier, Kevin B., Ólan T. Henry, Nilss Olekalns, and Kalvinder Shields (2004) The asymmetric effects of uncertainty on inflation and output growth. *Journal of Applied Econometrics* 19, 551–565.
- Hafner, Cristian M. and Helmut Herwartz (2006) Volatility impulse response functions for multivariate GARCH models: An exchange rate illustration. *Journal of International Money and Finance* 25, 719–740.
- Hamilton, James D. (1983) Oil and the macroeconomy since World War II. *Journal of Political Economy* 91, 228–248.
- Hamilton, James D. (2003) What is an oil shock? *Journal of Econometrics* 113, 363–398.
- Hamilton, James D. and Anna Herrera (2004) Oil shocks and aggregate macroeconomic behavior: The role of monetary policy. *Journal of Money, Credit and Banking* 36, 265–286.
- Hansen, Peter R. and Asger Lunde (2005) A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics* 20, 873–889.
- Hooker, Mark A. (1996) What happened to the oil price–macroeconomy relationship? *Journal of Monetary Economics* 38, 195–213.
- Jarque, C.M. and A.K. Bera (1980) Efficient tests for normality, homoscedasticity, and serial independence of regression residuals. *Economics Letters* 6, 255–259.
- Koop, G., M.H. Pesaran, and S.M. Potter (1996) Impulse response analysis in non-linear multivariate models. *Journal of Econometrics* 74, 119–147.
- Kroner, Kenneth F. and Victor K. Ng (1998) Modeling asymmetric comovements of asset returns. *Review of Financial Studies* 11, 817–844.
- Kwiatkowski, D., P.C.B. Phillips, P. Schmidt, and Y. Shin (1992) Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics* 54, 159–178.
- Lee, K. and S. Ni (2002) On the dynamic effects of oil price shocks: A study using industry level data. *Journal of Monetary Economics* 49, 823–852.
- Lee, K., S. Ni, and R.A. Ratti (1995) Oil shocks and the macroeconomy: The role of price variability. *The Energy Journal* 16, 39–56.
- Ljung, T. and G. Box (1979) On a measure of lack of fit in time series models. *Biometrika* 66, 66–72.
- Mork, Knut A. (1989) Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy* 91, 740–744.
- Pagan, A. (1984) Econometric issues in the analysis of regressions with generated regressors. *International Economic Review* 25, 221–247.

- Potter, S.M. (2000) Nonlinear impulse response functions. *Journal of Economic Dynamics and Control* 24, 1425–1446.
- Shields, Kalvinder, Nills Olekalns, Ólan T. Henry, and Chris Brooks (2005) Measuring the response of macroeconomic uncertainty to shocks. *Review of Economics and Statistics* 87, 362–370.
- van Dijk, Dick, Philip Hans Franses, and H. Peter Boswijk (2007) Absorption of shocks in nonlinear autoregressive models. *Computational Statistics and Data Analysis* 51, 4206–4226.