

THE COMOVEMENTS BETWEEN FUTURES MARKETS FOR CRUDE OIL: EVIDENCE FROM A STRUCTURAL GARCH MODEL*

Fabrizio Spargoli and Paolo Zagaglia

This version: August 8, 2007

ABSTRACT

This paper studies the linkages between the prices of oil futures traded on the New York Mercantile Exchange and the Intercontinental Exchange of London. We estimate a structural BEKK-GARCH model that allows for non-zero correlation between the structural innovations. We identify the structural parameters through restrictions on the reduced-form GARCH model.

We find that the oil futures traded on the NYMEX and ICE can be used for mutual hedging purposes only when the structural conditional variances of both innovations are modest and, as such, no turbulent events have taken place. Periods with positive structural correlations are instead associated with peaks in the structural conditional variance of both innovations. During times of market turmoil, the structural variance of the returns on NYMEX futures becomes larger than that of ICE futures. This means that, when there are common shocks to both markets, the NYMEX reacts more strongly than the ICE. Our empirical evidence explains the negative reduced-form correlation between the two returns which is observed in turbulent periods.

JEL Classification: C22, G19

Keywords: oil prices, futures markets, GARCH, structural VAR

**Spargoli* (corresponding author): Università Politecnica delle Marche, fspargol@mta01.univpm.it.
Zagaglia: Stockholm University and European Central Bank, pzaga@ne.su.se. We are grateful to the Norwegian School of Management (BI) for providing the data. Zagaglia acknowledges the kind support of the Monetary Policy Stance Division of the European Central Bank, where this paper was finalized. The opinions expressed herein are those of the authors and do not necessarily reflect the views of the ECB.

Introduction

Commodity markets have by far lost their original function of trade and physical delivery of goods, and have become suitable for speculative and hedging purposes. In fact, the most part of trading in commodity markets is conducted through futures contracts, which are generally cash-settled rather than physically delivered. This trend has gone hand in hand with the suggestion by market commentators that speculative activity is the major contributor behind the steady surge in crude oil prices. For instance, the IMF World Economic Outlook of September 2006 reports that “*the share of non-commercial contracts* - reported by the U.S. Commodity Futures Trading Commission n.d.r. - *has steadily increased since 1995 from 9 percent to 16 percent of the total*”.

In this paper we study the linkages between the prices on crude oil products traded in the world’s largest commodity markets, the New York Mercantile Exchange – NYMEX - and the Intercontinental Exchange – ICE – of London. We use daily time series of the returns on the NYMEX light sweet crude oil futures and the ICE Brent crude futures in order to estimate a structural GARCH model in the spirit of Rigobon and Sachs (2003a). This amounts to identifying a structural VAR model for the two returns through restrictions on the reduced-form GARCH. The restrictions arise from a set of hypotheses about the conditional variance of innovations in structural form. Differently from Rigobon and Sachs (2003a), we assume that the joint evolution of the variances follows from a multivariate BEKK-GARCH model, which allows for non-zero correlation between the structural innovations.

Our results show that the oil futures traded on the NYMEX and ICE can be used for mutual hedging purposes only when the structural conditional variances of both innovations are modest and, as such, no turbulent events have taken place. Periods with positive structural correlations are instead associated with peaks in the structural conditional variance of both innovations. During times of market turmoil, the structural variance of the returns on NYMEX futures becomes larger than that of ICE futures. This means that, when there are common shocks to both markets, the former reacts more strongly than the latter. Our empirical evidence explains the negative reduced-form correlation between the two returns observed in turbulent periods.

This paper is organized as follows. In the second section we present a selective overview of the institutional characteristics of oil futures trading on NYMEX and ICE. The third section outlines the estimation methodology. The fourth section presents the results, and section five draws the main conclusions.

The institutional features of NYMEX and ICE: An overview

The New York Mercantile Exchange is the world's largest physical commodity futures exchange and the most important trading forum for energy and precious metals. It originated from the merger between New York's two largest exchanges, the New York Mercantile Exchange and the Commodity Exchange, in 1994. It operates through two divisions: the NYMEX division, where energy, platinum and palladium are traded, and the COMEX division, which is entitled for all other metals. The most part of trading is conducted through futures contracts, which were introduced by the Exchange in 1981 and rapidly overcame traditional trading as a mean of exchange. The overwhelming majority of exchange trading activity is executed by open outcry on the trading floor during the day. However, energy and metals futures contracts are also available for trading on the CME Globex electronic trading platform when the trading floor is closed, making the markets available for a more than 22 hours a day. Besides standard futures contracts, the Exchange also lists NYMEX “miNY” energy futures, fractional light, sweet crude oil and natural gas futures contracts which are suited to small investors and traders. In fact they are reduced-size contracts traded through an electronic trading system.

Established in May 2000, in June 2001, the ICE expanded its business into futures trading by acquiring the International Petroleum Exchange (IPE), now ICE Futures. Since 2003, ICE has partnered with the Chicago Climate Exchange to host its electronic marketplace. In April of 2005, the entire ICE portfolio of energy futures became fully electronic. In January of 2007, ICE acquired the New York Board of Trade.

As regards the future contracts, in this paper we consider the light sweet crude oil futures, traded in the NYMEX, and the Brent crude futures contract, traded in the ICE. The first one is the world's most liquid and largest-volume futures contract on a physical commodity. Because of its excellent liquidity and price transparency, the contract is used as a principal international pricing benchmark. The Brent Crude Futures Contract, together with West Texas Intermediate Crude futures, accounts for nearly half of the world's global crude futures by volume of commodity traded.

The structural multivariate GARCH model

Let us assume that the evolution of the variables can be summarized by a structural VAR model

$$Ax_t = \psi + \Phi(L)x_t + \eta_t$$

where η_t is the vector of structural shocks, and A is the structural parameter matrix

$$A = \begin{bmatrix} 1 & a_{12} \\ a_{21} & 1 \end{bmatrix}$$

Direct estimation of the matrix A through OLS leads to asymptotically-biased estimates, owing to the endogeneity of some of the variables. Therefore, the structural parameters should be derived from the reduced form of the model through an identification procedure, as usual when dealing with structural VARs.

One of the solutions to the identification problem relies on the existence of heteroskedasticity. This idea has been originally introduced by Wright (1928) and recently developed by Rigobon (2003). The heteroskedasticity approach to identification amounts to using the information from time-varying volatility as a source of information on the relation between endogenous variables. This would allow us to identify the structural parameters of the model without need for additional assumptions.

In Rigobon (2003) and Rigobon and Sachs (2003b, 2004), identification is obtained through regimes of volatility. In other words, these authors consider subsamples across which there are shifts in the volatility pattern. A natural extension of this methodological framework involves the modelling of heteroskedasticity through GARCH processes so that regimes changes are continuous.

Rigobon and Sachs (2003a) use this formulation to study the relation between yields on Treasury bills with short (3 months) and long (10 years) maturity, and the Standard & Poor's 500. Assuming that the structural shocks have a zero mean, are independent but not i.i.d., the authors postulate that their variances follow the GARCH process

$$\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix}_t = \begin{bmatrix} \sqrt{h_1} & 0 & 0 \\ 0 & \sqrt{h_2} & 0 \\ 0 & 0 & \sqrt{h_3} \end{bmatrix} \begin{bmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{bmatrix}_t, \quad \begin{bmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{bmatrix}_t \sim N \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right)$$

with $h_t = \sqrt{h_t} \sqrt{h_t} = V(\eta_t)$ defined as

$$\underline{h}_t = \underline{c} + \Gamma \underline{h}_{t-1} + \Lambda \underline{\eta}_{t-1}^2$$

with:

$$\underline{h}_t = \begin{bmatrix} h_{11} \\ h_{22} \\ h_{33} \end{bmatrix}_t, \quad \underline{c} = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}, \quad \underline{\eta}_{t-1}^2 = \begin{bmatrix} \eta_1^2 \\ \eta_2^2 \\ \eta_3^2 \end{bmatrix}_{t-1}$$

The matrices Γ and Λ are square with dimension 3. Their elements are restricted to be positive. Since the shocks of the reduced form are a linear combination of the structural shocks, they also have a conditional variance that follows a GARCH process. In particular,

$$\begin{bmatrix} H_{11} \\ H_{12} \\ H_{22} \\ H_{13} \\ H_{23} \\ H_{33} \end{bmatrix}_t = B_t \underline{c} + B_t \Gamma (B^2)^{-1} \begin{bmatrix} H_{11} \\ H_{22} \\ H_{33} \end{bmatrix}_{t-1} + B_t \Lambda (B^2)^{-1} \begin{bmatrix} v_1^2 \\ v_2^2 \\ v_3^2 \end{bmatrix}_{t-1}$$

with

$$B_t = \begin{bmatrix} b_{11}^2 & b_{12}^2 & b_{13}^2 \\ b_{11}b_{21} & b_{12}b_{22} & b_{13}b_{23} \\ b_{21}^2 & b_{22}^2 & b_{23}^2 \\ b_{11}b_{31} & b_{12}b_{32} & b_{13}b_{33} \\ b_{21}b_{31} & b_{22}b_{32} & b_{23}b_{33} \\ b_{31}^2 & b_{32}^2 & b_{33}^2 \end{bmatrix} \quad \text{and} \quad B = A^{-1}$$

In this model, the restrictions that yield identification are imposed on the covariance matrix of the reduced form. This, in turn, depends on the heteroskedasticity of the structural shocks.

The formulation of Rigobon and Sachs (2003a), however, does not guarantee that variance-covariance matrices are positive-definite, which is a problem typical of every vector – vech – GARCH. In this paper, we rely on a BEKK-GARCH (Engle and Kroner, 1995) in order to cope with this problem. In particular, we assume that structural form innovations η_t are distributed according to

$$\eta_t \sim N(0, h_t), \quad h_t = CC' + Gh_{t-1}G' + T\eta_{t-1}\eta_{t-1}'T'$$

where C is a triangular matrix whose elements are all positive, G and T are two parameters matrix such that G_{11} and T_{11} are constrained to be positive. Given that the degree of generality, the order of the autoregressive component and the order of the moving average component are all equal to 1, Proposition 2.1 in Engle and Kroner (1995) guarantee that these restrictions are sufficient for the identification of the parameters of the GARCH model.

We should stress that this model implies that structural form innovations are correlated, contrary to what Rigobon and Sachs assume. Using time series of the returns on oil futures traded in two different markets, in fact, it is very likely that their evolution depends on common factors that make the structural form innovations of the two series linked each other to some extent.

Identification of the structural parameters is achieved like in Rigobon and Sachs (2003a) through restrictions on the conditional variance-covariance matrix of the reduced form innovations,

which are represented by the BEKK-GARCH model we put forward. For the purpose of estimation, we begin with the OLS estimate of the VAR model

$$x_t = c + F(L)x_t + v_t$$

where $c = A^{-1}\psi$, $F(L) = A^{-1}\Phi(L)$ and $v_t = A^{-1}\eta_t$ are the reduced form innovations, whose variance-covariance matrix is a combination of the variance-covariance matrix of the structural form innovations, that is

$$H_t = A^{-1}h_t A^{-1} \rightarrow H_t = B h_t B'$$

$$H_t = BCC'B' + BGh_{t-1}G'B' + BT\eta_{t-1}\eta_{t-1}'T'B'$$

In this formulation the variance-covariance matrix of the reduced form innovations is a function of the structural innovations, which we do not have. However, we can use the equality $\eta_t = Av_t$ to show that

$$\eta_t\eta_t' = Av_t v_t' A' \quad h_t = AH_t A'$$

and to represent H_t in terms of the reduced form innovations as

$$H_t = BCC'B' + BGAH_{t-1}A'G'B' + BTA v_{t-1} v_{t-1}' AT'B'$$

It should be stressed that H_t is positive-definite by construction because it is given by the sum of positive-definite components. Furthermore, the model can be seen as an augmented BEKK-GARCH model, given that the reduced form depends also on the structural parameters in matrix A. It is from that dependence that we are able to identify the structural parameters.

After obtaining the residuals from the model in reduced form, we can estimate the parameters by maximum likelihood on the function

$$\ell_t(\theta) = -\frac{n}{2} \ln 2\pi - \frac{1}{2} \left[\ln |H_t| + v_t' H_t^{-1} v_t \right]$$

with $\theta' \equiv [k', \text{vec}(A)', \text{vec}(G)', \text{vec}(T)']$ and H_t is the covariance matrix defined earlier. In the practical implementation of the estimation algorithm, special care must be used to address the presence of kinks and local maxima in the likelihood function. We have chosen to run a number of initial steps through *simulated annealing* in order to obtain robust estimates of the initial points for the maximization step. In the second round, we have used gradient-based optimization methods conditional on the initial point from simulated annealing.

After estimating the model, we compute impulse-response functions. In structural GARCH models, these functions show the impact that shocks produce on the conditional second

moments of the variables in the system. However, differently from the impulse response function for a standard VAR, the impulse responses of a structural GARCH depend both on the magnitude of the shock and on the period during which the shock itself takes place. This is due to the fact that the residuals enter the model in quadratic form. Hence, differently from the case of linear models, the magnitude of the effects of a shock is not proportional to the size of the shock itself. This allows us to compute a distribution of impulse responses following each shock. To that end, we use the concept of “Volatility Impulse Response Functions” – VIRF - proposed by Hafner e Herwartz (2006). The impulse-response function for a vech-GARCH model can be written as

$$V_t(\xi_0) = E[\text{vech}(H_t) | \xi_0, I_{-1}] - E[\text{vech}(H_t) | I_{-1}]$$

The response at time t of the variances and covariances following a shock η in t=0 - denoted as $V_t(\eta_0)$ - is equal to the difference, conditioned on the information set at time -1 (I_{-1}) and on the shock η_0 , of the variance (or covariance) at t from its expected value conditional on the information set of period -1. In Appendix I, we show how to obtain the analytical formulas used in the computation of the impulse responses.

Results

We estimate the model using daily data from the 26th of April 1993 to the 26th of April 2007. The data series on the prices of the futures are downloaded from Bloomberg. As regards the expiration date, we decided to consider front month future contracts in both cases, because they are the most actively traded and volatile. We calculate the returns in percentage points from the two series. The sample includes a total of 3653 observations. The time series of the returns on oil futures traded on NYMEX and ICE are plotted in figure 1.

In order to obtain reduced-form residuals, we estimate a VAR model with a constant and a set of dummy variables to account for outlier observations. Outliers are found through E-views as observations that lie outside the intervals given by the third quartile plus 3 times the interquartile range, and the first quartile less 3 times the interquartile range. We detect 27 extreme observations. After including the dummy variables in the VAR, we also perform a set of Wald exclusion tests to select the best fitting model for the conditional mean, which turned out to be a VAR with 7 lags and 26 dummy variables.

Given the reduced form innovations, the maximization method outlined in the previous section yields the results reported in Table 2. All the coefficients are statistically different from zero, except for g_{21} for which, however, the t statistics is very close to the 5% threshold. From the point

estimates of the matrix A , the links between the returns on oil futures traded on NYMEX and ICE, ignoring lags and exogenous variables, are

$$\text{NYMEX}_t = 4.257 \text{ ICE}_t$$

$$\text{ICE}_t = -2.131 \text{ NYMEX}_t$$

These estimates imply that a 1 basis point increase in the return on ICE futures causes a 4.257 basis points increase in the return on NYMEX futures, while a 1 basis point increment of the latter leads to a 2.131 basis point decrease in the former. We can interpret these figures as showing a hedging motive for trading between the NYMEX and ICE. A positive return shock in the NYMEX causes a response of opposite sign in the ICE. In other words, when a shock happens to the price of Light Sweet crude oil futures, the evidence suggests that traders readjust their portfolios away from Brent Crude futures. The relation of substitutability between the two markets does not hold when a shock hits the price of Brent Crude futures. A positive price shock to the ICE drives up the prices in the NYMEX.

Figure 2 shows the structural conditional variances, while the structural conditional correlation is depicted in Figure 3. The conditional structural variances of innovations to NYMEX oil futures are greater than those of ICE oil futures on absolute value. Peaks occur at the same time for both conditional variances. The structural correlations have a floor of about -0.8 over the full sample. However, it is evident from Figure 3 that there are some peaks which make the structural correlation go up to almost 0.5. It should be noted that these peaks occur at the same time of the peaks in structural conditional variances. This means that when volatility is low the shocks in the two markets are negatively correlated, while they are positively correlated in periods of high volatility. This evidence can be due to the fact that the commodities underlying the two types of futures contracts are substitutes, so that a shock to the price of one of them implies an opposite shock to the price of the other. However, in case of case of turbulence due to exceptional events, such as international conflicts, the price of oil follows a common behaviour and, therefore, both markets are subject to a common shock.

Figure 4 shows the reduced-form conditional correlation between the returns on oil futures in the NYMEX and ICE, which takes into account the links between the two markets. Differently from the structural correlation, the reduced-form correlation has a ceiling about 0.8, and displays frequent peaks that make it negative. A possible interpretation of this fact is that investors, given the negative structural correlation of the two shocks, buy both types of futures for hedging purposes, making the returns positively correlated in reduced form. This happens only in non-turbulent periods, when the structural volatility of the two innovations is relatively regular. On the

other hand, in turbulent times, it is no longer possible to hedge the returns of NYMEX futures against those of ICE futures because their structural correlation becomes positive. The reduced-form conditional correlation, however, looks more irregular than the structural one. In particular, it shows bigger oscillations in the central part of the sample and smaller ones in the extreme parts, even if negative peaks occur all over the sample.

The evidence from the estimated matrix A and from the evolution of the conditional structural variances helps to understand the dynamics of the reduced-form correlation in periods when this becomes negative or only mildly positive, that is when investors do not hedge NYMEX futures against ICE futures. On those days, the structural conditional variance of both futures is higher in absolute value than over the rest of the sample. However, the structural conditional variance of the returns on NYMEX futures is larger than that of ICE futures returns. This is the case, for instance, for the central part of the sample, and for the observations between March 1995 and February 1997. The pattern of the reduced-form conditional correlation can be explained by the negative relationship in conditional mean between the returns on NYMEX futures and those on ICE futures. The fact that the former are more volatile than the latter implies that structural shocks to the NYMEX futures are larger and more important in terms of propagation, so that a negative correlation arises from the structural link between the two markets.

In order to analyze the persistence of the effects of the shocks, we present some evidence from the volatility impulse responses. As explained before, given that GARCH are non-linear in the innovations, the effect of a shock depends both on the size and on the timing. Therefore, the use of VIRFs that we can make is twofold. On the one hand, we can show traditional impulse responses from a given shock occurred at a specific point in time. On the other hand, we can compute the distribution of VIRFs, that is we can calculate impulse responses for each shock at each time horizon of the VIRFs, and then we can determine their frequency.

The first panel of figure 5 shows the impulse responses for the shock occurred on the 11th of September 2001, whereas the second panel plots the responses to the second Gulf war shock, which began on the 20th of March 2003. The first shock produces the largest impact on the covariance between the two returns, which decreases by 0.6 basis points. The reason for this can be tracked in the reaction of the variance of the returns on NYMEX futures, which increased by 0.3 basis points, and in the negative structural link between NYMEX and ICE futures. As regards the second Gulf war shock, we can make the same considerations, even if the effect on the conditional covariance is less severe. Overall, figure 5 shows that the effects of shocks tend to be absorbed as the time horizon increases. The same observation emerges from looking at the distribution of the VIRFs along the time horizon.

Figure 6 shows the 1st, 10th and 25th percentiles, and figure 7 displays the 50th, the 75th, the 90th and the 99th percentiles. At a first glance, what emerges is that the percentiles get closer and closer to zero day by day, which means that the effects of the shocks tend to be absorbed. Furthermore, it should be noted that the distribution of the VIRF for the reduced-form conditional variance of the returns on NYMEX futures is positive on the entire time span, meaning that all the shocks that occur in the sample have a positive effect. On the contrary, the distribution of the VIRFs for the reduced-form conditional covariance is characterized by negative values. Furthermore, the effects of the shocks on impact are larger for both the conditional variance of the returns on NYMEX futures and the conditional correlation, given that their extreme – 1st and 99th – percentiles are far apart from each other.

Conclusions

In this paper we analyzed the inter-relations between NYMEX and ICE using daily time series of, respectively, the returns on the light sweet crude oil futures and the brent crude futures. To this end, we estimated a structural BEKK-GARCH model in the spirit of Rigobon and Sachs (2003a), i.e. we identify structural parameters through restrictions on the reduced-form GARCH model. Contrary to Rigobon and Sachs (2003a), however, our model guarantees that variance-covariance matrix is positive-definite and allows for a non-zero correlation between the structural innovations. Furthermore, we use the volatility impulse responses functions of Hafner and Herwartz (2006) in order to estimate the size and persistence of the effects of structural innovations.

The main conclusion that can be drawn from our analysis is the following. In normal periods, namely when the structural conditional variances of both innovations are regular, NYMEX and ICE futures are used by investors for hedging purposes, given that the structural correlation of their innovations is negative. However, in turbulent periods when there are peaks in the structural conditional variance of both innovations, the structural correlation between them is positive and hedging is no more feasible. Furthermore, in those periods we observe that the structural variance of the returns on NYMEX futures becomes larger than that of ICE futures, meaning that when there are common shocks to both markets the former reacts more strongly than the latter. This is evidence, together with the estimated negative structural link between NYMEX and ICE returns, is able to explain the negative or less positive reduced-form correlation between the two returns which is observed in turbulent periods.

References

- Engle, Robert F., and Kenneth F. Kroner, “Multivariate Simultaneous Generalized ARCH”, *Econometric Theory*, 11, 1995
- Hafner, Christianam M., and Helmut Herwartz, “Volatility Impulse Responses for Multivariate GARCH Models: An Exchange Rate Illustration”, *Journal of International Money and Finance*, 25(5), 2006
- Rigobon, Roberto, “Identification through Heteroskedasticity”, *Review of Economics and Statistics*, 85(4), 2003
- Rigobon, Roberto, and Brian Sack, “Spillovers across U.S. Financial Markets”, *unpublished manuscript*, MIT Sloan School of Management, 2003(a)
- Rigobon, Roberto, and Brian Sack, “Measuring the Reaction of Monetary Policy to the Stock Market”, *Quarterly Journal of Economics*, 118(2), 2003(b)
- Rigobon, Roberto, and Brian Sack, “The Impact of Monetary Policy on Asset Prices”, *forthcoming on the Journal of Monetary Economics*, 2004
- Wright, Philip, *The Tariff on animal and Vegetable Oils*, The institute of Economics, The Macmillan Company, New York, 1928

Appendix I

The VIRF for the first period takes the form

$$V_1(\xi_0) = E[\text{vech}(H_1) | \xi_0, I_{-1}] - E[\text{vech}(H_1) | I_{-1}]$$

with

$$\begin{aligned} E[\text{vec}(H_1) | \xi_0, I_{-1}] &= E\left[\text{vec}(BCC'B' + BGAH_0A'G'B' + BTABh_0^{\frac{1}{2}}\xi_0\xi_0^{\frac{1}{2}}h_0^{\frac{1}{2}}B'A'T'B') | \xi_0, I_{-1}\right] \\ &= E\left[\text{vec}(BCC'B' + BGAH_0A'G'B' + BTh_0^{\frac{1}{2}}\xi_0\xi_0^{\frac{1}{2}}h_0^{\frac{1}{2}}T'B') | \xi_0, I_{-1}\right] \end{aligned}$$

given that $v_t = Bh_t^{\frac{1}{2}}\eta_t$ and $B = A^{-1}$, and

$$E[\text{vec}(H_1) | I_{-1}] = E[\text{vec}(BCC'B' + BGAH_0A'G'B' + BTAv_0v_0'A'T'B') | I_{-1}]$$

where $h_t^{\frac{1}{2}}$ can be obtained from a Jordan decomposition of h_t . In particular, labelling λ_i ($i=1, \dots, N$, where N is the number of variables) the eigenvalues of h_t and γ_i the corresponding eigenvectors, the symmetric matrix $h_t^{\frac{1}{2}}$ is defined as:

$$h_t^{\frac{1}{2}} = \Gamma_t \Lambda_t^{\frac{1}{2}} \Gamma_t'$$

This implies that

$$V_1(\xi_0) = B \otimes B' \text{vec} \left(Th_0^{\frac{1}{2}} \xi_0 \xi_0' h_0^{\frac{1}{2}} T' - TAH_0 A'T' \right)$$

since the first expected value is conditional to the shock at $t=0$, while the second is conditioned only on the information set at time $t=-1$. Following the same logic, we can find proper expressions for the impulse responses in the subsequent periods. For two periods ahead, the responses are

$$V_2(\xi_0) = E[\text{vech}(H_2) | \xi_0, I_{-1}] - E[\text{vech}(H_2) | I_{-1}]$$

with

$$E[\text{vec}(H_2) | I_{-1}] = E[\text{vec}(BCC'B' + BGAH_1 A'G'B' + BTAv_1 v_1' A'T'B') | I_{-1}]$$

$$E[\text{vec}(H_2) | \xi_0, I_{-1}] = E[\text{vec}(BCC'B' + BGAH_1 A'G'B' + BTAv_1 v_1' A'T'B') | \xi_0, I_{-1}]$$

where

$$H_1 = BCC'B' + BGAH_0 A'G'B' + BTh_0^{\frac{1}{2}} \xi_0 \xi_0' h_0^{\frac{1}{2}} T'B'$$

Since the expected value is conditional to ξ_0 , this gives

$$V_2(\xi_0) = B \otimes B' \text{vec} \left[G \left(Th_0^{\frac{1}{2}} \xi_0 \xi_0' h_0^{\frac{1}{2}} T' - TAH_0 A'T' \right) G' \right]$$

Following the same logic we can write the response s periods ahead, which is given by

$$V_s(\xi_0) = B \otimes B' \text{vec} \left[\underbrace{G \dots G}_{s-1 \text{ times}} \left(Th_0^{\frac{1}{2}} \xi_0 \xi_0' h_0^{\frac{1}{2}} T' - TAH_0 A'T' \right) \underbrace{G' \dots G'}_{s-1 \text{ times}} \right]$$

Figure 1: Returns on oil futures traded on NYMEX and ICE

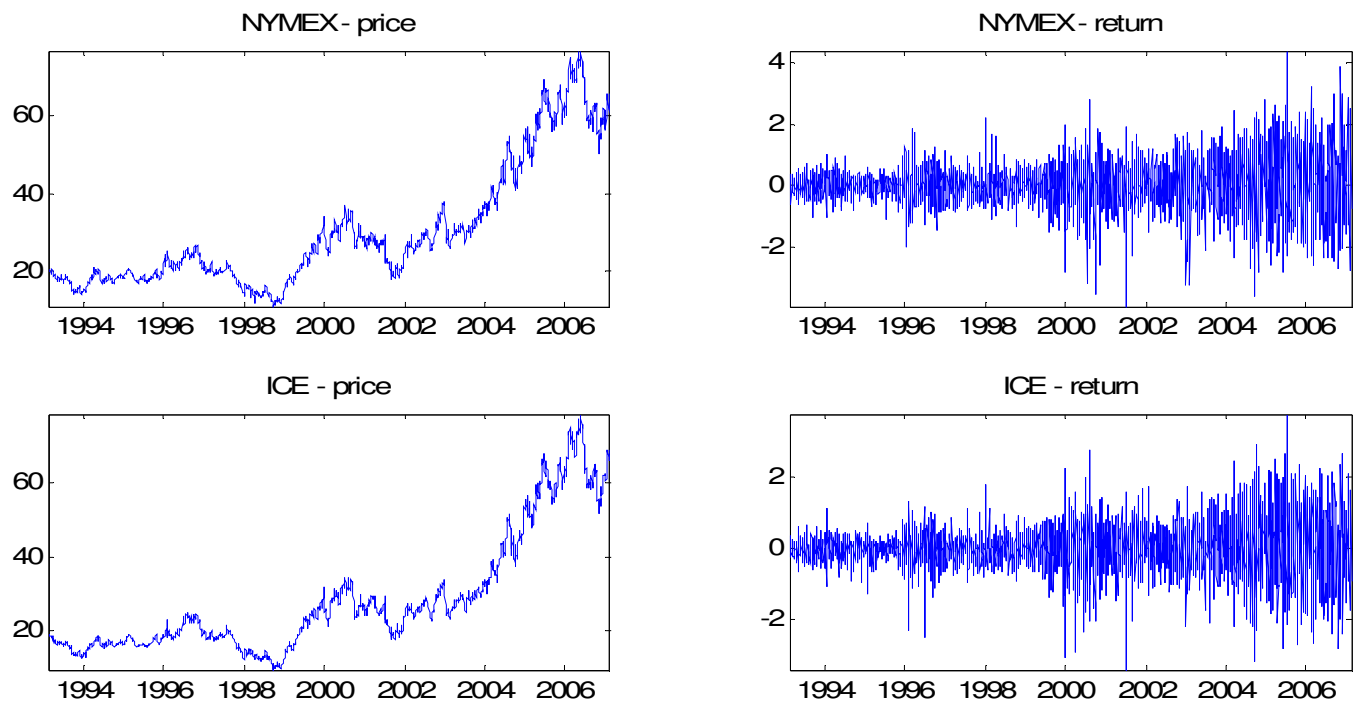


Figure 2: Structural Conditional Variances

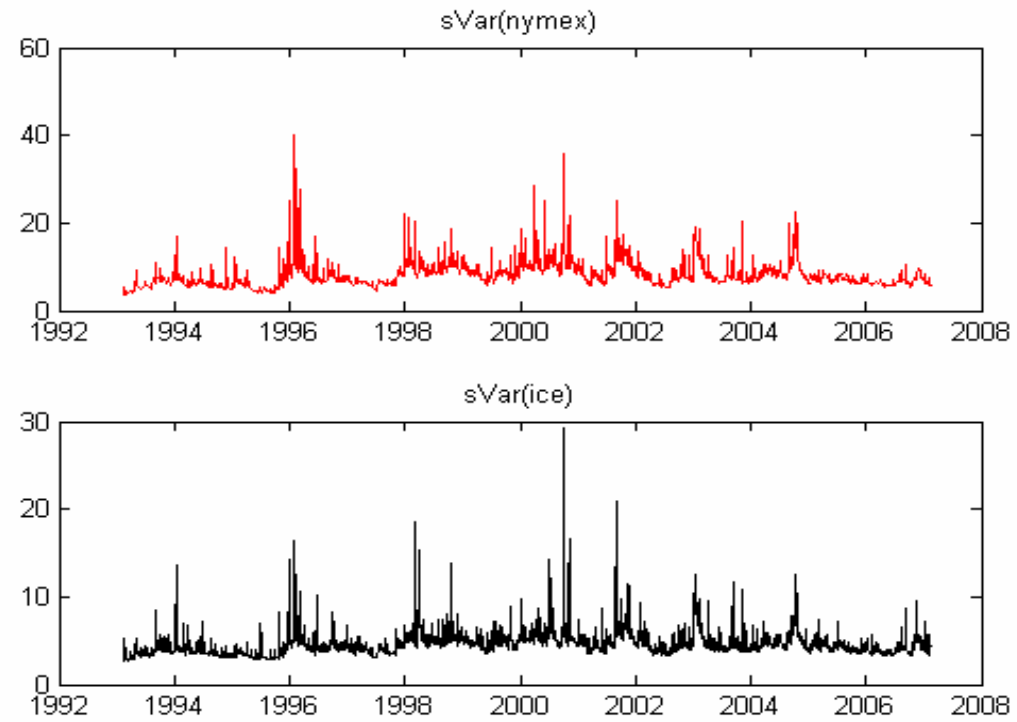


Figure 3: Structural Conditional Correlation




Figure 4: Reduced-form conditional correlation between the returns on NYMEX and ICE futures




Figure 5: VIRFs of reduced-form moments following the 11th of September and the second Gulf war shocks

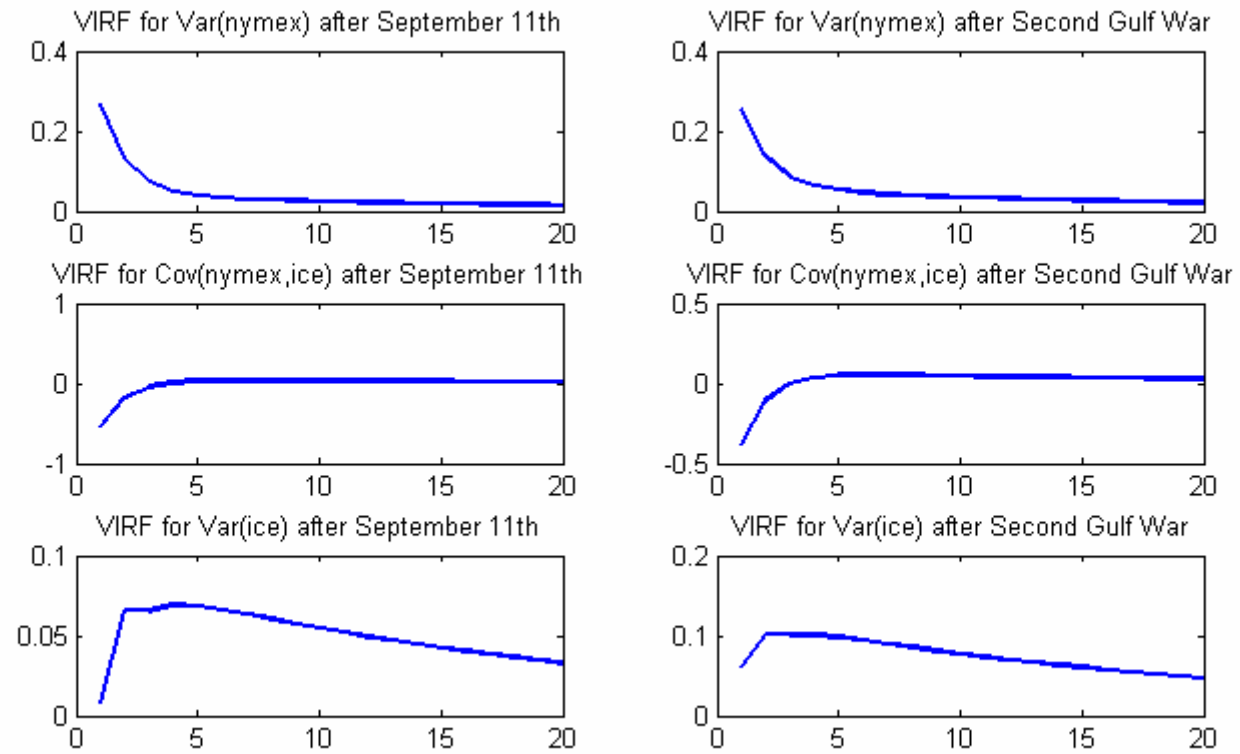


Figure 6: 1st, 10th and 25th percentiles of the VIRF distribution

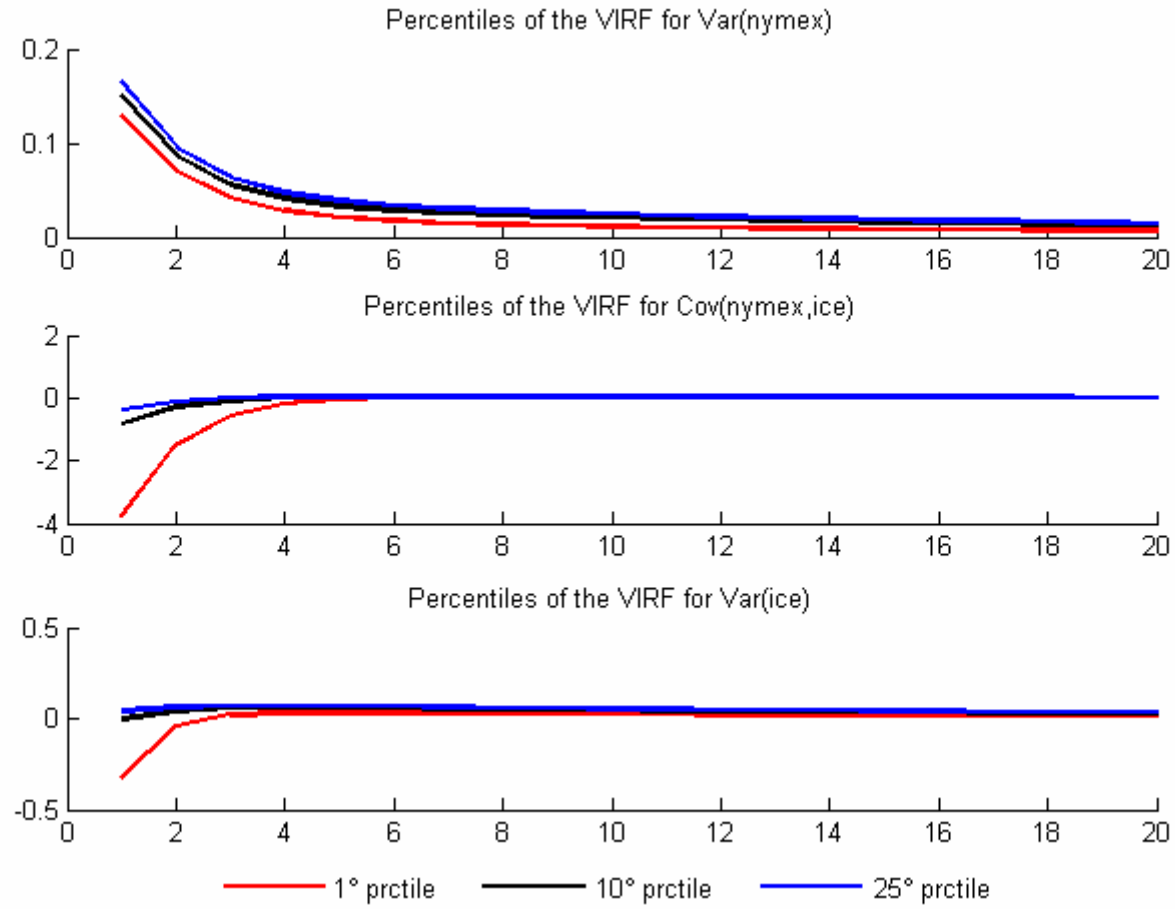


Figure 7: 50th, 75th, 90th and 99th percentiles of the VIRF distribution

