THE IMPACT OF ENERGY DERIVATIVES ON THE CRUDE OIL MARKET

Jeff Fleming
Jones Graduate School of Management, Rice University

Barbara Ostdiek
Jones Graduate School of Management, Rice University

Abstract

We examine the effects of energy derivatives trading on the crude oil market. There is a common public and regulatory perception that derivative securities increase volatility and can have a destabilizing effect on the underlying market. Consistent with this view, we find an abnormal increase in volatility for three consecutive weeks following the introduction of NYMEX crude oil futures. While there is also evidence of a longer-term volatility increase, this is likely due to exogenous factors such as the continuing deregulation of the energy markets. Subsequent introductions of crude oil options and derivatives on other energy commodities have no effect on crude oil volatility. We also examine the effects of derivatives trading on the depth and liquidity of the crude oil market. This analysis reveals a strong inverse relation between the open interest in crude oil futures and spot market volatility. Specifically, when open interest is greater, the volatility shock associated with a given unexpected increase in volume is much smaller.

April 7, 1998

Jones Graduate School – MS 531, Rice University, P.O. Box 1892, Houston, TX 77251-1892

THE IMPACT OF ENERGY DERIVATIVES ON THE CRUDE OIL MARKET

Abstract

We examine the effects of energy derivatives trading on the crude oil market. There is a common public and regulatory perception that derivative securities increase volatility and can have a destabilizing effect on the underlying market. Consistent with this view, we find an abnormal increase in volatility for three consecutive weeks following the introduction of NYMEX crude oil futures. While there is also evidence of a longer-term volatility increase, this is likely due to exogenous factors such as the continuing deregulation of the energy markets. Subsequent introductions of crude oil options and derivatives on other energy commodities have no effect on crude oil volatility. We also examine the effects of derivatives trading on the depth and liquidity of the crude oil market. This analysis reveals a strong inverse relation between the open interest in crude oil futures and spot market volatility. Specifically, when open interest is greater, the volatility shock associated with a given unexpected increase in volume is much smaller.

I. Introduction

Beginning in the 1970s, deregulation dramatically increased the degree of price uncertainty in the energy markets, prompting the development of the first exchange-traded energy derivative securities. The success and growth of these contracts attracted a broader range of participants to the energy markets and stimulated trading in an even wider variety of energy derivatives. Today, many exchanges and over-the-counter markets worldwide offer futures, futures options, swap contracts, and exotic options on a broad range of energy products including crude oil, fuel oil, gasoil, heating oil, unleaded gasoline, and natural gas.

It is well-known that derivative securities provide economic benefits.[1] The key attribute of these securities is their leverage, i.e., for a fraction of the cost of buying the underlying asset, they create a price exposure similar to that of physical ownership. As a result, they provide an efficient means of offsetting exposures among hedgers or transferring risk from hedgers to speculators. In addition, derivatives promote information dissemination and price discovery. The leverage and low trading costs in these markets attract speculators, and as their presence increases, so does the amount of information impounded into the market price. These effects ultimately influence the underlying commodity price through arbitrage activity, leading to a more broadly based market in which the current price corresponds more closely to its true value. Because this price influences production, storage, and consumption decisions, derivatives markets contribute to the efficient allocation of resources in the economy.

Nonetheless, the tightened cross-market linkages that result from derivatives trading also fuel a common public and regulatory perception that derivatives generate or exacerbate volatility in the underlying asset market. These concerns are often voiced in the context of their "destabilizing" effects around major declines in the market. Following the 1987 stock market crash, for example, John Shad, former chairman of the Securities and Exchange Commission argued, "Futures and options are the tail wagging the dog. They have escalated the leverage and volatility of the markets to precipitous, unacceptable levels."[2] This concern has led to studies commissioned by the Securities and Exchange Commission, the Commodity Futures Trading Commission, and a Presidential Task Force, and has been a driving force behind the

---

[1] Peck (1985) and Stoll and Whaley (1985) examine the economic benefits of futures and options, respectively.
[2] *The Wall Street Journal*, January 15, 1988.

adoption of program trading curbs, circuit breakers and daily price limits in the futures markets, and the staggering of stock index futures and options expirations.

There exists little theoretical or empirical evidence, however, to justify these actions. In perfect markets, derivatives should have no effect on the underlying asset market because they are redundant securities, i.e., they can be synthetically created by some combination of the asset and riskless bonds. With market imperfections, derivatives make the market more complete [Ross (1976) and Hakansson (1982)] by allowing investment choices that were previously cost inefficient or impossible due to regulatory or institutional constraints. Since investors benefit from an expanded opportunity set, the required returns and risks in existing asset markets should fall. In addition, Danthine (1978) argues that derivatives, by promoting information-based trading, increase the depth and liquidity of the market and reduce volatility. Grossman (1988) shows that option trading allows diverse opinions about volatility to be revealed which can reduce volatility. Detemple and Selden (1991) show that option trading can allow more efficient risk sharing which increases the demand for the asset and reduces volatility. Stein (1987) is the only theoretical study that implies volatility could increase, arguing that poorly informed speculators can have a destabilizing effect on the market.

The empirical evidence is generally consistent with these theoretical implications.[3] The evidence tends to focus on stock option introductions due to the large quantity of listing events, and most of these studies [e.g., Skinner (1989) and Conrad (1989)] find a reduction in volatility following introduction. In addition, Damodaran and Lim (1991) and Skinner (1990), respectively, find that the speed with which information is incorporated into price and the accuracy of this information increases after options are introduced. Kumar, Sarin, and Shastri (1998) find a decrease in the adverse selection component of the bid/ask spread and a reduction in the pricing error variance after option introduction, signaling an improvement in pricing efficiency and market quality. In other markets, Edwards (1988) finds reductions in volatility following the introductions of stock index futures and treasury bill futures, while Harris (1989) shows that the volatility of S&P 500 stocks increased after the introduction of S&P 500 futures.

There is also evidence that volatility decreases when the trading activity in existing derivatives markets increases. Bessembinder and Seguin (1992), for example, find that stock

---

[3] Damodaran and Subrahmanyam (1992) provide a thorough survey of the evidence regarding the effects of derivative securities on the underlying asset market.

market volatility is inversely related to both the open interest and trading volume of S&P 500 futures after controlling for spot market volume. Bessembinder and Seguin (1993) find that spot volatility is positively related to unexpected volume and negatively related to expected open interest for eight currency, interest rate, and commodity futures contracts. For the currency and agricultural contracts, spot volatility decreases when unexpected open interest increases. These findings indicate that futures trading increases the depth and liquidity of the underlying asset market, mitigating the impact of volume shocks on volatility.

In general, there is little research regarding physical commodity derivatives, and this research is primarily focused on agricultural futures contracts.[4] For our analysis of the energy markets, there are at least two reasons we might expect results that differ from past research. First, in these markets, it is difficult to trade on "bad news" that would negatively affect the market price without using derivative securities. Therefore, if derivatives provide benefits of increased informational efficiency, their effects may be more pronounced in the energy markets. Second, there tend to be strong informational linkages across energy markets. Information that affects crude oil prices can also affect, say, natural gas or heating oil prices. Given these linkages, the introduction of natural gas or heating oil derivatives could influence the crude oil market through its effect on the transfer of information across markets.

To examine the effect of derivative introductions in these markets, we must address two complications. First, in a typical event study, we average the abnormal effects around an event across many observations to control for factors other than the event. This is not possible here. The introduction of a given energy derivative contract only happens once, and we have only one price history from which to draw our inference. In essence, our event study has a sample size of one. Second, the timing of the oil futures introduction closely corresponds with that of the degregulation of the U.S. oil market. Therefore, our sample of "free-floating" spot prices extends just a year prior to the introduction.

We address these complications by fitting a stochastic volatility model to the sample of post-introduction prices. The model controls for the time-series structure of volatility, capturing the nature of volatility persistence, mean reversion in volatility, and the volatility of volatility in the crude oil market. We then examine whether, given the structure imposed by

---

[4] These studies include Working (1960) [onion futures], Powers (1970) [pork bellies and cattle futures], Taylor and Leuthold (1974) [cattle futures], and Cox (1976). All of these studies find a reduction in spot market volatility following the introduction of exchanged-traded futures.

the model, the volatility shocks around the futures contract introduction date seem abnormal. By using just the post-introduction sample for estimation, the fitted model is not influenced by the volatility process that prevailed at the time of introduction. However, if this process is consistent with the post-introduction process and the introduction had no effect on volatility, then the innovations around the introduction date should not appear unusual.

Our results indicate that volatility increased after the introduction of crude oil futures. Positive abnormal volatility shocks are observed for three consecutive weeks following the introduction. We also find evidence of a much longer-term (more than a year) volatility increase, but it is inappropriate to simply attribute this effect to derivatives. The increase coincides with the growth of the energy derivatives markets which was spurred by volatility induced by continuing deregulation of the energy markets. Given this linkage, it is difficult to disentangle the cause from the effect. After the introduction of crude oil futures, there is little evidence that subsequent introductions had any effect on oil market volatility. In particular, we find no volatility effects around the introduction of crude oil options and no pattern in the effects across the time-series of introductions on other energy products. This evidence contradicts the idea that subsequent introductions should gradually complete the market.

To more fully assess the impact of derivatives on the crude oil market, we also examine the ongoing dynamics between futures trading activity and spot market volatility. This analysis reveals a strong positive relation between unexpected futures volume and unexpected volatility. This relation is weaker, but still positive, for the long-term trend and expected volume components. We also find evidence of asymmetry in the volume-volatility relation. Specifically, an increase in unexpected volume is associated with an increase in spot market volatility that is 80% larger than the decrease in volatility associated with an equivalent decrease in unexpected volume.

In contrast to the volume-volatility relation, we find that the overall size of the crude oil futures market (measured by open interest) is negatively related to spot market volatility. The relation is strongest for the unexpected component of open interest, but is also present for the long-term trend and expected open interest. This finding indicates that the futures market provides depth and liquidity to the crude oil market. Moreover, when combined with the positive volume-volatility relation, it implies that the unexpected change in open interest for a given shock to futures volume either mitigates or amplifies the effect on spot volatility. For

example, the volatility increase associated with unexpected volume is approximately 40% less when it is accompanied by an unexpected increase in open interest than when open interest remains unchanged. This result may reflect not only changes in market depth but also the nature of the trades that accompany the increased volume.

The remainder of this study is organized as follows. Section II describes the data used in our analysis and some preliminary evidence regarding the structure of crude oil volatility. Section III develops our stochastic volatility model for the oil market, our estimation strategy, and the estimation results. Section IV examines the effects of energy derivative introductions on crude oil market volatility, and Section V examines the depth and liquidity effects of derivatives trading on the crude oil market. Section VI provides a summary and conclusions.

II. Data and Preliminary Analysis

Table 1 lists the primary energy futures and futures option contracts, along with their respective introduction dates. Each of these contracts is traded at either the New York Mercantile Exchange (NYMEX) or the International Petroleum Exchange (IPE). Our study focuses on the West Texas Intermediate (WTI) crude oil market, the commodity underlying the NYMEX crude oil futures contract. The contract is denominated in 1,000 U.S. barrels (42,000 gallons) of light, sweet crude oil for delivery in Cushing, Oklahoma. Futures contracts are currently traded for thirty consecutive months plus five long-dated maturities extending out seven years.

To examine the effect of derivative introductions on the oil market, we need a sample of spot oil prices that begins prior to the introduction of crude oil futures. Reliable data for this period are scarce because the introduction closely coincides with the deregulation of the U.S. oil market. Although the *Wall Street Journal* and several industry publications reported "posted prices" prior to deregulation, these prices do not necessarily represent actual spot market prices. The data we use for this analysis are from DataStream International.[5] Prices for WTI near (oil for prompt month delivery) are available on a weekly basis beginning February 2, 1982, and on a daily basis beginning September 1, 1983. Daily spot prices for sweet Cushing crude begin April 5, 1983. For the oil futures introduction analysis, we use the weekly WTI prices and, to maintain consistency, we use the daily WTI prices to examine subsequent introductions. For our analysis of the relation between futures trading activity and spot market volatility, we use

---

[5] DataStream International obtains the spot oil data from Independent Chemical in the United Kingdom.

the daily sweet Cushing prices and the total daily futures volume and open interest across all available NYMEX crude oil contracts. These futures data also are obtained from DataStream International. All of our data series extends through the end of 1997. In addition, we obtained annual world oil production data from the American Petroleum Institute's *Basic Petroleum Data Book*.

A. Summary Statistics

Table 2 summarizes the price series used for our analysis. Over the course of our sample period, crude oil prices fell from nearly $34 per barrel in 1982 to under $18 by the end of 1997, an average annual return of about –4%. Oil prices ranged from a low of $10.80 in July 1986 to a high of $40.85 in October 1990. The high variability of oil prices relative to most financial assets is apparent from the annual returns reported in Table 2. Prices increased more than 25% during three different years of the sample, and they fell by 35–40% in three others.

Our first objective is to examine the volatility of oil returns. There is considerable evidence that volatility changes over time, but conditional volatility is not observable and we must rely on estimates to examine the nature of time-variation. The most common approach [e.g., Poterba and Summers (1986) and French, Schwert, and Stambaugh (1987)] is to consider the standard deviation of returns over a fixed window of observations. Table 2 reports these standard deviations for each year of our sample. No real patterns are apparent, except perhaps that the estimates appear to be relatively low in the first couple years. It is difficult, however, to attribute the subsequent increase in standard deviation to the introduction of oil futures in 1983. The estimates are quite noisy, and the standard deviations based on weekly observations actually indicate a reduction in volatility in 1983, and again in 1984. After this, the estimates range from 50–60% in 1986, 1990, and 1991, down to about 20% in 1992 and 1995.

B. Rolling Volatility Estimates

Relying on the standard deviation to detect variation in conditional volatility is problematic because it assumes volatility is constant within each estimation window (i.e., a year). We can reduce this problem by shortening the window length, but a reasonable number of datapoints are required within each window to obtain precise estimates. We address these issues by adopting a "rolling" estimation approach. We use a window of observations around time t to estimate the conditional volatility, s_t, and we move the window forward one period to estimate

s_{t+1}. Because volatility time-varies within each window, observations nearer to t should convey more information about s_t. We accommodate this by giving more weight to these observations in forming our estimate of s_t. Foster and Nelson (1996) show that, under reasonable smoothness restrictions, this approach yields consistent and asymptotically normal estimators.

To apply the rolling estimation approach, we define the estimator,

$$\hat{s}_t^2 = \sum_{l=-t+1}^{T-t} w_{t+l}(r_{t+l} - m_{t+l})^2, \tag{1}$$

where r_{t+l} and m_{t+l}, respectively, are the conditional return and mean return, w_{t+l} is the weight placed on the innovation at time $t + l$, and T is the number of observations in the sample.[6] Foster and Nelson (1996) show that, if volatility is stochastic, the optimal weighting function for a two-sided rolling estimator is

$$w_{t+l} = (a_t/2)e^{-a_t|l|}, \tag{2}$$

where a_t is the decay rate. This estimator is two-sided because it uses both leads and lags of r_t to estimate s_t^2. To construct a one-sided estimator (i.e., based only on past information), we set $w_{t+l} = 0$ for $l > 0$, and double each of the weights for $l \leq 0$.

Foster and Nelson (1996) show that the optimal choice of a_t is f_t/q_t, where f_t^2 is the conditional variance of volatility innovations and $(q_t^2 + s_t^4)/s_t^4$ is the conditional coefficient of returns kurtosis. We eliminate the time-dependency in a_t by assuming the volatility innovations are proportional to volatility ($f_t = f s_t^2$) and the coefficient of kurtosis is constant ($q_t = q s_t^2$). If we also assume that the conditional distribution of returns is normal ($q = 2$), then setting $a_t = f/2$ is optimal. Using the estimation procedure developed in Fleming, Kirby, and Ostdiek (1998b) yields $a = 0.1155$ for daily returns, and $a = 0.1443$ for weekly returns.

Figure 1 plots the time-series of rolling, exponentially weighted volatility estimates obtained from equation (1). The trends in the daily estimates (Panel A) and the weekly estimates (Panel B) are similar. (Note the difference in x-axis due to the earlier start of the weekly sample.) The largest volatility shocks occur in 1986 when oil prices fell by nearly \$10/bbl and in 1990 following Iraq's invasion of Kuwait. Aside from these, there is a general upward trend from 1984 through 1988, with a sharp swing from 1989 to 1991, and relatively steady, lower

---

[6] We model the conditional mean using the procedure developed in Fleming, Kirby, and Ostdiek (1998b). We first remove day-of-the-week effects from returns by regressing raw returns on a set of six variables: a dummy variable for each weekday and a variable equal to the number of nontrading days between return observations. We then use the residuals from this regression to remove day-of-the-week and nontrading day seasonalities in squared returns.

volatility thereafter. The most significant difference in the daily and weekly estimates occurs in 1996 when several large 2-3 day price swings are not detected with weekly observations.

The patterns shown in Figure 1 are generally consistent with the standard deviations reported in Table 2, but two additional features of volatility are now observable. First, the time-series estimates in Figure 1 allow us to detect finer variations in volatility. We can see, for example, that weekly volatility is locally high at the beginning of 1983 (prior to the introduction of crude oil futures), and then falls steadily over the remainder of the year. Second, we can observe stylized facts regarding the time-series structure of volatility. In particular, like most financial time-series, crude oil volatility is persistent and tends to mean-revert over time. These observations motivate our strategy for evaluating the effect of derivative introductions on volatility. Specifically, we must model the time-series structure of volatility in order to evaluate whether any variation around the introduction date is unusual.

III. A Stochastic Volatility Model

In this section, we develop and estimate a stochastic volatility model for the crude oil market. The model captures the structure of mean-reversion, persistence, and volatility of volatility apparent in the data, and allows us to assess whether the volatility realizations following the introduction of energy derivatives are inconsistent with this structure. We begin by outlining the specification and the intuition behind the model. Then, we describe our estimation strategy and results. Finally, we generate the volatility residuals under the model and examine whether the model adequately captures the time-series structure of volatility in the oil market.

A. The Stochastic Volatility Specification

Our analysis is based on the volatility model developed in Fleming, Kirby, and Ostdiek (1998a). The setup is similar to Clark (1973) and Tauchen and Pitts (1983) where we have an economy that consists of a large number of active speculators with heterogeneous expectations about asset value. As new information arrives in the market, traders revise their expectations and initiate a round of trading. Over the course of a day, these information arrivals generate a large number of unpredictable price changes. If we let $\boldsymbol{e}_{it}$ represent the incremental return generated by event i, then the return on day t can be modeled as

$$r_t = \boldsymbol{m}_t + \sum_{i=1}^{I_t} \boldsymbol{e}_{it}, \tag{3}$$

where I_t is the number of information events that occur. We assume e_t is *iid* normal with mean zero and variance s_e^2, but note that because we can rewrite the summation in equation (3) as $s_e z_t \sqrt{I_t}$, where $z_t \equiv 1/\sqrt{I_t} \sum_i^{I_t} (e_{it}/s_e)$, the central limit theorem implies $z_t \xrightarrow{d} N(0,1)$ as $I_t \longrightarrow \infty$. Therefore, even if e_t is non-normal and exhibits weak forms of serial dependence, the conditional distribution of r_t should be approximately normal with mean mean m_t and variance $s_e^2 I_t$.

We impose more time-series structure by exploiting the relation between information flow and the volatility of returns ($s_t = s_e \sqrt{I_t}$).[7] As noted above, volatility is persistent and empirical research indicates that increases in volatility are more likely than decreases of the same magnitude (i.e., asymmetry). We capture these features by focusing on the representation,

$$r_t = m_t + \exp(\tfrac{1}{2} h_t) z_t, \tag{4}$$

where $h_t \equiv \ln s_t^2$, and modeling h_t as an AR(1) process,[8]

$$h_t = g + f_h h_{t-1} + u_t, \tag{5}$$

where u_t is *iid* with mean zero and independent of z_t.

The AR(1) structure in equation (5) yields a volatility specification that is similar in many respects to an EGARCH model [Nelson (1991)]. Volatility is constrained to be non-negative, it follows an exponential autoregressive process, and is asymmetric in levels. An important difference, however, is that under our model volatility is stochastic rather than known conditional on past prices. This feature is attractive because the information flow to financial markets is unpredictable and it is information that generates volatility. As a result, our specification may better capture salient features of the return generating process.

B. Model Estimation

We estimate and test our volatility specification by forming a set of moment restrictions from equations (4) and (5) and applying Hansen's (1982) generalized method of moments (GMM). We assume $|f_h| < 1$ in equation (5), so h_t is stationary with mean $m_h = g(1 - f_h)$ and variance $s_h^2 = s_u^2 (1 - f_h^2)$. The autocorrelation of return innovations is zero at all lags, but there can be a substantial degree of higher-order dependence apparent in the logarithm of squared returns,

$$\ln r_t^2 = h_t + \ln z_t^2. \tag{6}$$

---

[7] Ross (1989) develops a similar relation between the information flow and volatility.
[8] This specification is similar to those examined by Taylor (1994) and Harvey, Ruiz, and Shephard (1994).

Because z_t is standard normal, the mean and variance of $\ln z_t^2$ are -1.27 and 4.93 [Abramowitz and Stegun (1972)]. Defining $y_t \equiv \ln r_t^2 - E[\ln z_t^2]$, we obtain the transformed system,

$$
\begin{aligned}
y_t &= h_t + x_t \\
h_t &= g + f_h h_{t-1} + u_t,
\end{aligned}
\tag{7}
$$

where $x_t \equiv \ln z_t^2 - E[\ln z_t^2]$, is mean zero with variance 4.93 and independent of h_t. Under our stated assumptions, we can obtain the following moment restrictions for y_t,[9]

$$
\begin{aligned}
E[y_t] &= E[h_t] \\
\text{var}(y_t) &= \text{var}(h_t) + \text{var}(x_t) \\
\text{cov}(y_t, y_{t+k}) &= (f)^k \, \text{var}(h_t)
\end{aligned}
\tag{8}
$$

for all integers $k > 0$.

To impose these moment restrictions and estimate the parameters of the model, we define the GMM disturbance vector,

$$
e_t(q) = \begin{bmatrix} y_t - m_h \\ (y_t - m_h)^2 - s_h^2 - s_x^2 \\ (y_t - m_h)(y_{t+k} - m_h) - (f_h)^k [(y_t - m_h)^2 - s_x^2] \end{bmatrix}
\tag{9}
$$

where $q = [m_h, s_h^2, f_h]'$ is the vector of unknown parameters, $k = 1, 2, \ldots, l$ counts the number of autocorrelation restrictions used in the estimation, and $s_x^2 = 4.93$. The first two restrictions identify the mean and variance of the log volatility, h_t, and the l remaining restrictions identify the AR(1) parameter of the h_t process, f_h.

We construct the y_t series used in the estimation by removing from the raw data any seasonal patterns in returns and volatility. First, we remove returns seasonality by using the residuals from a regression of raw returns on six variables: a dummy variable for each weekday and a variable that counts the number of nontrading days between observations. Second, we remove volatility seasonality by regressing these residuals on the Monday dummy and nontrading day variables. Adding 1.27 to the intercept and residuals from this regression yields the seasonally adjusted series that we use in the estimation.[10]

We estimate the system by minimizing $g_T(q)' \hat{S}^{-1} g_T(q)'$ where $g_T(q)' \equiv \frac{1}{T-l} \sum_{t=1}^{T-l} e_t(q)$ and $\hat{S}$ is a consistent estimate of the GMM covariance matrix. For the asymptotic distribution theory

---

[9] Note that y_t is the sum of an AR(1) component and a noise process. Therefore, its autocovariance function should be identical to that of an ARMA(1,1) process.

of GMM to hold, we assume that the series is stationary and ergodic and that the regularity conditions in Hansen (1982) are satisfied. Our choice of $\hat{S}$ adjusts for conditional heteroskedasticity and autocorrelation using Parzens weights and Andrews' (1991) method of bandwidth selection. The system in equation (9) has $l + 2$ moment conditions and three unknown parameters, leaving $l - 1$ overidentifying restrictions. As a result, the GMM procedure yields a direct test for specification error in the form of an overidentifying test statistic [Hansen (1982)]. Since there is no theoretical guidance for choosing the optimal l, we estimate the system using $l = 10, 20, 30$, and 40 for daily observations and $l = 12, 16, 20$, and 24 for weekly observations.

Table 3 reports the estimation results. In general, the parameter estimates are fairly insensitive to the lag length. The mean of h_t is stable, and although f_h increases slightly for longer lags at the daily level, no such pattern is apparent at the weekly level. All of the estimates of f_h indicate a slow decay in the autocorrelation function of h_t, suggesting a long lag length is necessary to capture the persistence of volatility. Therefore, for the remainder of the study, we rely on the estimation results using $l = 40$ for daily returns and $l = 24$ for weekly returns. These lag lengths encompass periods of about two months and six months, respectively.

The final two lines in each panel of Table 3 report the overidentifying test statistics for our stochastic volatility model. None of these statistics indicate rejection. The statistics become less significant with longer lag lengths, but this is consistent with our argument that longer lags are necessary to capture the strong volatility persistence. Therefore, we conclude that the GMM estimation reveals little evidence of model misspecification.

C. Fitted Volatilities

We now want to use our fitted volatility model to evaluate whether the residuals under the model seem abnormal following the introduction of derivatives. Although our GMM approach yields parameter estimates for the model, it does not produce a fitted time-series of volatility estimates (or residuals).[11] We generate these estimates using the Kalman filter.[12]

---

[10] In the empirical work, we estimate the GMM system using both daily and weekly data. For the weekly deseasonalization, we omit the day-of-the-week dummy variables and use only the nontrading day variable.

[11] As an alternative, we could simply include pre- (or post-) introduction dummy variables in our GMM system and directly estimate the structural change following the introduction. This method, however, effectively partitions the available data into subsamples, making it more difficult to obtain precise estimates using GMM. As a result, this approach yields no meaningful evidence of structural change for any of the introduction dates reported in Table 1.

[12] The Kalman filter is a standard technique for constructing time-series estimates of random variables. A number of econometric texts provide detailed descriptions of the procedure. See, for example, Hamilton (1994).

To fit the filter to our stochastic volatility specification, we express equation (7) as

$$y_t = h_t + \mathbf{x}_t$$
$$h_t = \mathbf{m}_h(1-\mathbf{f}_h) + \mathbf{f}_h h_{t-1} + u_t, \tag{10}$$

where $\mathbf{m}_h(1 - \mathbf{f}_h) = \mathbf{g}_h$ is the constant in the AR(1) specification of volatility. We parameterize equation (10) using the consistent estimates obtained from our GMM analysis. The filtering algorithm takes the observed y_t series and, for each day in the sample, delivers two estimates of h_t. The first estimate is the best linear forecast of h_t given all of the data available through time $t-1$ (i.e., a one-sided estimate). The second, commonly called the smoothed estimate, is the best linear estimate based on the entire sample (i.e., a two-sided estimate).

Figure 2 plots the fitted volatilities. Comparing these estimates to the rolling volatility estimates in Figure 1 (note that the y-scales for the two figures are slightly different) reveals that the fitted volatilities exhibit less time-series variation. In other words, we observe fewer extreme volatilities in Figure 2. This should not be surprising, however, since the Kalman filter procedure generates a best linear fit of the unobservable volatility at each point, and, therefore, unusual price changes influence this estimate less than they influence the rolling estimate. Aside from this difference, the patterns shown in Figures 1 and 2 are generally comparable.

D. Diagnostics

As a final robustness check before using the Kalman filter estimates to evaluate the effect of derivative introductions, we conduct a series of specification tests similar to those used to evaluate GARCH models. Our model implies that the time t return is drawn from a normal distribution with mean $\mathbf{m}_t$ and variance I_t. Therefore, if the model is well-specified, the standardized, seasonally-adjusted returns (z_t) should be *iid* normal with mean zero and variance one. We construct the z_t series from our Kalman filter estimates of h_t,

$$z_t = \frac{R_t - \mathbf{m}_1 MON_t - \cdots - \mathbf{m}_5 FRI_t - \mathbf{m}_6 NTDYS_t}{\exp[0.5h_t]\exp[0.5(\mathbf{a}_1 MON_t + \mathbf{a}_6 NTDYS_t)]} \ . \tag{11}$$

The second term in the denominator accounts for volatility seasonalities, and is the same adjustment we used to compute the y_t series for the GMM estimation.[13] If our model is well specified, the moments of z_t should match those of a standard normal random variable.

---

[13] To see this, note that the log of the square of equation (9) is $y_t - h_t - 1.27$.

Table 4 reports the specification results for both daily (Panel A) and weekly (Panel B) datasets. The first four columns report the mean, variance, skewness, and excess kurtosis of z_t (and the smoothed estimates, z_t^*), and the final three columns report the autocorrelations of the series, its absolute values, and its squared values. As a benchmark for comparison, we also report these statistics for the non-standardized, seasonally adjusted returns. Focusing on the standardized returns, both the one-sided and smoothed series exhibit substantial departures from normality. In particular, for each series, the variance is greater than one and both the skewness and excess kurtosis are positive.

We evaluate the significance of these results using simulations.[14] We use our GMM estimates to parameterize the return generating process in equations (4) and (5), and we simulate realizations of z_t and u_t to generate the h_t and y_t series. We then apply the Kalman filter to this y_t series to estimate h_t, we construct the standardized returns (z_t), and we compute each of the statistics reported in Table 4. We repeat this simulation 5000 times. In Table 4, under each statistic, we report the probability of realizing in the simulations a value lower than that observed in the data. These probabilities indicate that the variance, skewness, and kurtosis of both daily and weekly returns are significantly greater than we would expect under the model, as are the autocorrelations of absolute and squared daily returns.

Despite these findings, there is also evidence that the model captures many features of observed returns. The deseasonalized returns (r) reported in Table 4 evidence large degrees of skewness and excess kurtosis at both the daily and weekly levels. The model explains much of this behavior, for example, reducing the skewness in daily returns by a factor of 17 and the excess kurtosis by a factor of 6. The model also reduces the intertemporal dependence apparent in squared daily returns, absolute and squared weekly returns, and the mean reversion apparent in weekly returns. These findings indicate that, although there is evidence of misspecification, the model performs rather well given its simple AR(1) structure.

IV. The Effects of Derivative Introductions on Crude Oil Volatility

A. Introduction of Crude Oil Futures

We now use our stochastic volatility model to evaluate the effect of energy derivative introductions on the structure of crude oil volatility. We focus first on the introduction of crude oil

---

[14] Note that a simple GMM test for normality would be inappropriate because the volatilities used to construct the standardized returns are measured with error.

Table 4 reports the specification results for both daily (Panel A) and weekly (Panel B) datasets. The first four columns report the mean, variance, skewness, and excess kurtosis of z_t (and the smoothed estimates, z_t^*), and the final three columns report the autocorrelations of the series, its absolute values, and its squared values. As a benchmark for comparison, we also report these statistics for the non-standardized, seasonally adjusted returns. Focusing on the standardized returns, both the one-sided and smoothed series exhibit substantial departures from normality. In particular, for each series, the variance is greater than one and both the skewness and excess kurtosis are positive.

We evaluate the significance of these results using simulations.[14] We use our GMM estimates to parameterize the return generating process in equations (4) and (5), and we simulate realizations of z_t and u_t to generate the h_t and y_t series. We then apply the Kalman filter to this y_t series to estimate h_t, we construct the standardized returns (z_t), and we compute each of the statistics reported in Table 4. We repeat this simulation 5000 times. In Table 4, under each statistic, we report the probability of realizing in the simulations a value lower than that observed in the data. These probabilities indicate that the variance, skewness, and kurtosis of both daily and weekly returns are significantly greater than we would expect under the model, as are the autocorrelations of absolute and squared daily returns.

Despite these findings, there is also evidence that the model captures many features of observed returns. The deseasonalized returns (r) reported in Table 4 evidence large degrees of skewness and excess kurtosis at both the daily and weekly levels. The model explains much of this behavior, for example, reducing the skewness in daily returns by a factor of 17 and the excess kurtosis by a factor of 6. The model also reduces the intertemporal dependence apparent in squared daily returns, absolute and squared weekly returns, and the mean reversion apparent in weekly returns. These findings indicate that, although there is evidence of misspecification, the model performs rather well given its simple AR(1) structure.

IV. The Effects of Derivative Introductions on Crude Oil Volatility

A. Introduction of Crude Oil Futures

We now use our stochastic volatility model to evaluate the effect of energy derivative introductions on the structure of crude oil volatility. We focus first on the introduction of crude oil

---

[14] Note that a simple GMM test for normality would be inappropriate because the volatilities used to construct the standardized returns are measured with error.

futures on March 30, 1983. Our strategy is as follows. We first fit our stochastic volatility model using the post-introduction sample of weekly data. We then use the resulting parameter estimates to calibrate the Kalman filter and estimate the weekly series of h_t for the entire sample (both pre- and post-introduction). Finally, we evaluate the significance of the h_t realizations subsequent to the introduction date. If the structure of volatility changed following the introduction, then these realizations will be inconsistent with our fitted model.

The GMM estimation results using the post-introduction sample (770 observations) are similar to those reported in Table 3 using the entire sample (829 observations). For a lag length of $l=24$, the estimates of $\boldsymbol{m}_h$, $\boldsymbol{s}_h^2$, and $\boldsymbol{f}_h$, respectively, are –6.8652, 1.4169, and 0.9562. The largest change from the overall sample is for the $\boldsymbol{s}_h^2$ estimate, but with a standard error over 0.26 this change is not statistically significant. The J-statistic for the post-introduction period is 16.66 (p-value = 0.8256). These findings suggest that excluding the pre-introduction sample does not meaningfully alter our fitted stochastic volatility model.

We now use these fitted parameter estimates in our Kalman filter procedure to estimate the h_t series for the entire sample (both pre- and post-introduction). For this analysis, we use the one-sided (rather than the smoothed) estimates from the filter so the current volatility estimates are not influenced by future innovations. On the last Friday before the introduction of crude oil futures, March 25, 1983, our estimate of h_t is –9.3070, which implies an annualized volatility rate of $\exp(\frac{1}{2}h_t)\sqrt{52} = 6.87\%$. Now, we need to determine whether the next k volatility realizations, conditioned on $\boldsymbol{s}_t$, are consistent with our fitted volatility model.

Fleming, Kirby, and Ostdiek (1998b) demonstrate that, under the model,

$$h_{t+k}|h_t \sim N\left(\frac{\boldsymbol{g}_h(1-\boldsymbol{f}_h^k)}{1-\boldsymbol{f}_h} + \boldsymbol{f}_h^k h_t, \; \boldsymbol{s}_h^2(1-\boldsymbol{f}_h^{2k})\right). \tag{12}$$

Given the volatility level on March 25, the $E[h_{t+k}|h_t]$ in equation (12) implies a volatility for the following week of 7.25%. The realized volatility was greater than expected, 8.70%. Using the distribution in equation (12), the probability of realizing a volatility less than 8.70% is 0.8534. This indicates that the increase in volatility during this week was not statistically significant.

It may be misleading, however, to use the analytical distribution in equation (12) to measure abnormal volatility. Our fitted volatilities are measured with error because we first estimate the parameters of our volatility model, and then we use the Kalman filter to estimate the true h_t series. This yields a fitted h_t series that is "smoother" than the true (but unobserv-

able) one. To assess the impact of these issues, we compare the distribution of $h_{t+1} - \mathrm{E}[h_{t+1}|h_t]$ innovations under equation (12) to the empirical distribution. Across the entire post-introduction sample, less than 1% of the realizations fall in the upper 10% of the analytical distribution, and only 4% of the realizations fall in the lower 10%. This finding indicates that the distribution of the fitted h_t series is indeed quite different than the analytical distribution.

To control for this difference, we use the empirical distribution of the fitted volatility innovations to determine whether volatility around the introduction date is abnormal. We use our fitted model, and the fitted h_t series, to compute the realized u_t under equation (10). We then simulate the empirical distribution by drawing (with replacement) from the sample of u_t realizations beginning one year after the introduction date. Starting from h_t, we generate a sequence $u_{t+1}, \ldots, u_{t+52}$, and use equation (10) to compute the corresponding $h_{t+1}, \ldots, h_{t+52}$. Repeating this process 5000 times, we approximate the distribution of $h_{t+k}|h_t$ for $k = 1, \ldots, 52$.

The second, third, and fourth columns of Table 5, respectively, report the fitted volatilities and their simulated expected values and probabilities for the 52 weeks following the introduction of crude oil futures. The second line, for example, shows the increase in volatility from 6.87% to 8.70% during the first week. Based on the empirical distribution, this increase appears to be abnormally high (p-value = 0.9704). During the following three weeks, volatility continued to increase, up to 14.52%. This realization, given $s_t = 6.87\%$, is also significant (p-value = 0.9966). By the twelfth week, however, volatility fell to 6.27%, and the volatilities realized after this date perhaps seem unusually low rather than high.

The average volatility statistics, reported in the final three columns of Table 5, allow us to address whether the average realization during the subsequent k weeks (rather than just the endpoint) is abnormal. We approximate the distribution of the average volatility using the same simulations as before, except now for each h_{t+k} realization we compute the average of $s_{t+1}, \ldots, s_{t+k}$. Consistent with the individual realizations, the average volatility over the first four weeks is significantly greater than expected. This similarity is not surprising since volatility follows a fairly direct path in reaching s_{t+4}. The average volatility through $t+24$, on the other hand, is less abnormal than the s_{t+24} realization. This occurs because volatility increases sharply and then decreases sharply to reach s_{t+24}. After 52 weeks, the average volatility, like s_{t+52}, seems unusually low. The actual level of volatility increases over this

15

period, however, the initial volatility ($s_t = 6.87\%$) is below the long-term mean ($s = 23.29\%$), and the rate of mean-reversion is slower than expected under our stochastic volatility model.

Conditional on s_t, both the realized and average volatility through $t+k$ and $t+k+1$ are correlated. To focus purely on the innovations between any two dates, we also consider the expected step-ahead realizations, i.e., $E[s_{t+k+1}|s_{t+k}]$. These expected values, and the realized p-values, are reported in the fifth and sixth columns of Table 5. This evidence indicates that the most unlikely sequence of innovations occur in the first three weeks after the introduction (p-values of 0.9704, 0.9128, and 0.9692). If we assume these innovations are *iid* normal, then the sum of their squared, standardized values is distributed c_3^2. The realized value is 8.8981, and the probability of realizing a value this high is just 0.0307. Over the remaining 49 weeks, the sequence of step-ahead realizations exhibit no apparent pattern.

Based on this evidence, we conclude that volatility indeed increased following the introduction of crude oil futures. The increase is prominent over the first three to four weeks, although an isolated sharp volatility drop occurs in the twelfth week. As a result of this drop, the realized and average volatilities after a year seem lower than expected. This finding may, in part, actually be symptomatic of a longer-term volatility increase following the introduction. Trading activity was thin during the first year of the oil futures market, but both volume and open interest grew by 500% after the first year, and by over 2500% after five years. Therefore, any volatility effects in the spot market might develop over a period of time. Because our model estimation is based on the entire post-introduction sample, these effects would be present in the model, but not in the data, during the first year. Volatility may seem low during this year only because it fails to revert toward this higher long-term mean volatility.

B. Subsequent Energy Derivative Introductions

Next, we examine the effect of other energy derivatives introduced after the introduction of crude oil futures. If our previous results are due to increasing market completeness, we might expect similar results following the introduction of crude oil options. Options may further complete the market because they allow a one-sided payoff structure that may be difficult or costly to create when there are market imperfections. Moreover, crude oil prices are correlated with other energy prices, and introducing derivatives on these assets may affect crude oil volatility. Detemple and Jorion (1990) and Detemple and Selden (1991) model these direct and cross-market interactions, and show that the volatility effects should be greatest

following the first derivative introduction and they should decay with subsequent introductions as the market gradually becomes more complete.

To investigate these issues, we apply our methodology to each of the subsequent introduction dates reported in Table 1. The only difference is that each of these introductions occurs after the start of our daily crude oil price series, so we use the daily prices (rather than weekly) in this analysis. For each introduction date, we begin by fitting our stochastic volatility model to the post-introduction sample (i.e., for unleaded gas futures, the sample is December 3, 1984 to December 31, 1997). We then use the resulting parameter estimates to calibrate the Kalman filter and estimate the daily h_t series for the entire sample. Finally, we evaluate the significance of the h_t realizations during the period following the introduction date.

Table 6 reports the results.[15] The "Model Parameters" columns in Panel A contain the GMM parameter estimates of our model for each of the introduction dates. In general, these estimates are similar to those reported in Table 3 for the overall sample, and there is not much variability across introduction dates. The only differences, perhaps, are the tendency toward a lower volatility of h, s_h, over time, and the dip in the AR(1) coefficient, f_h, that occurs near the middle dates. As noted earlier, however, these differences are not statistically significant.

The remaining columns of Panel A show the p-values for the average volatility realized $k = 1, 20, 40, \dots , 100$ days after the introduction date. For the first introduction, unleaded gas futures, the average volatility is less than expected for the entire 100-day period, in contrast to our findings for the introduction of crude oil futures. The source of this pattern is apparent from the p-values reported in Panel B for the realized volatilities and the volatility innovations. After 20 days, the volatility level is abnormally low, but subsequent volatilities conform more closely with expectations. The only other marginally abnormal shock (p-value = 0.063) occurs 120 days after the introduction. This shock, and the general trend of lower than expected volatilities for several months, is consistent with the long-term increase in volatility that we hypothesized earlier. Unlike our earlier results, however, volatility decreases initially after the introduction. This is inconsistent with the directional effect for crude oil futures, although the evidence here is less conclusive.

---

[15] Table 6 does not include results for the introduction of natural gas options because our GMM approach does not converge for lag lengths of $l = 10, 20, 30,$ or 40. A possible explanation for this is the small sample size. This introduction date (October 2, 1992) is the closest one to the end of our sample, and the post-introduction period includes only 1345 observations (versus 3641 observations in the full sample).

The introduction effects are even less apparent for the other introduction dates examined in Table 6. Few of the average or realized volatilities for these introductions are significantly different from what we expect. The primary exception is for natural gas futures, but the run-up in crude oil volatility following this date (April 3, 1990) can be attributed to Iraq's invasion of Kuwait. Comparing the results across all introductions reveals no systematic patterns within post-introduction periods and no trends in the effects across introductions. This evidence provides little support for the hypothesis that the volatility effects should gradually disappear with subsequent introductions. Instead, the effects are ambiguous for the first introduction after crude oil futures, and they are not at all detectable for any others.

V. Analysis of Futures Trading Depth and Liquidity Effects

A. Methodology

In this section, we provide further evidence on the impact of derivatives on the crude oil market by examining the effect of futures trading on the market depth and liquidity. Specifically, we assess the relation between spot market volatility and changes in the size of the futures market (as represented by open interest) and trading volume. As Figure 3 illustrates, both volume and open interest in NYMEX crude oil futures have increased dramatically since the inception of the contract. By 1990, the barrels of oil represented by NYMEX futures trades in one year actually exceeded the annual world production of oil. Figure 4 shows that this increasing trend has been accompanied by substantial variability in daily trading activity. We focus on the effect of this variability.

Table 7 provides summary statistics for daily futures trading activity and spot volatility. The volume and open interest data represent aggregate amounts across all open NYMEX crude oil contracts and the spot prices are for WTI sweet Cushing crude oil. We estimate the spot volatility by first fitting our stochastic volatility model to the daily Cushing returns, and then we use the parameter estimates in the Kalman filter to estimate the stochastic volatility time-series. The parameter estimates using these data ($m_h = -8.6140$, $s_h = 1.4801$, $f_h = 0.9853$) are similar to those reported in Table 3 for the WTI near price series with 40 lags.

The returns and volatilities reported in Table 7 exhibit the same general patterns as those for the WTI near series reported in Table 2. The volume and open interest statistics show the rapid growth in oil futures trading through the 1980s. In the first year of trading, average

daily volume represented 1.7 million barrels of oil and average daily open interest represented 8.8 million barrels. Both series peaked in 1994 with volume of 106.8 million barrels and open interest of 411.6 million barrels. The standard deviations for both series substantially increased over this period as well. Finally, the autocorrelation statistics reveal strong persistence in the trading activity and volatility data.

To analyze the relation between futures trading activity and spot market volatility, we regress unexpected spot volatility ($UVOL_t$) on the expected and unexpected components of futures volume and open interest (A_{ti}),

$$UVOL_t = a + \sum_{i=1}^{6} b_i A_{ti} + \sum_{j=1}^{4} h_j d_j + \sum_{k=1}^{10} g_k UVOL_{t-k} + e_t. \tag{13}$$

We include daily dummy variables (d_j) and lagged volatility shocks ($UVOL_{t-k}$) to control for day-of-the-week effects and volatility persistence. We proxy for $UVOL_t$ by subtracting the one-sided, contemporaneous Kalman filter estimate (realized volatility on day t) from the one-step ahead Kalman filter estimate (expected volatility on day $t-1$). We distinguish between the expected and unexpected components of volume and open interest due to the high persistence in these variables. Following Bessembinder and Seguin (1992), we first detrend each series by subtracting its 100-day moving average, and then we fit an ARIMA model to estimate its expected and unexpected components. For both variables, the optimal fit is an ARIMA(0,1,21) which incorporates about one month of data. We use the expected component from this model as a proxy for the predictable level of trading activity, and we use the unexpected component to proxy for the daily shock. We also include the 100-day moving average in the regression to represent longer-term shifts in trading activity. Note that summing these three components yields the original trading activity series.

B. Volume-Volatility and Open Interest-Volatility Relations

The first set of columns in Table 8 reports the regression results for the raw trading activity series over the full sample. The raw series are scaled so the underlying unit is one million futures contracts. The results indicate that the lagged unexpected volatilities and daily dummy variables in the regression are not significant. This is expected because we accounted for seasonalities and volatility persistence in constructing our unexpected volatility estimate. All of the trading activity variables, however, are highly significant. The moving average, expected, and unexpected components of volume are each significant at the 5% level. The coefficient

estimates indicate that the effect of unexpected volume on volatility is by far the strongest, nearly three times greater than the effect of the moving average component and nearly two times that of the expected component. This strong volume-volatility relation, in part, is influenced by the effect of spot market volume on volatility. We would expect a strong link between spot and futures market volumes, and we cannot isolate the marginal impact of futures volume without controlling for spot volume.[16]

In contrast to the volume coefficients, the coefficients on open interest are all significantly negative. Again, the magnitude of the coefficient on the unexpected component is much larger than the coefficients on the two predictable components (nearly five times the moving average component and two times the expected component). These estimates indicate that, conditional on futures volume, the long-term increase in open interest is related to lower spot market volatility, and that unexpected increases in open interest correspond to negative volatility shocks. Therefore, the volatility shock associated with a given volume is less when market depth increases. This finding is consistent with the results obtained by Bessembinder and Seguin (1992, 1993) for other markets, and supports the idea that futures trading improves depth and liquidity in the underlying market rather than destabilizing the market.

The negative coefficient on unexpected open interest indicates that an increase in open interest mitigates the impact of a volume shock on volatility. We can estimate the magnitude of this effect by comparing the coefficients on unexpected open interest and unexpected volume. The marginal impact of an unexpected volume of one million crude oil contracts on volatility is 1.8391 ± 0.7539 (or $\pm 41.0\%$) depending on whether open interest unexpectedly increases or decreases. This effect of open interest on the volume-volatility relation may reflect the nature of trades that increase end-of-the-day open interest. As Bessembinder and Seguin (1993) argue, open interest may not only proxy for market depth but also for uninformed trading. Many speculators are "day-traders" who exit their positions overnight, so open interest tends to reflect uninformed trading initiated by hedgers. To the extent this argument holds, we can distinguish between the price effects generated by informed versus uninformed trading in the crude oil market. Specifically, if an unexpected increase in volume is accompanied by an unexpected increase in open interest, more of the unexpected volume is attributable to hedgers and, therefore, the price revisions are smaller.

---

[16] We cannot include spot volume in the regression because reliable spot volume data is not available.

C. Robustness Checks

The summary statistics reported in Table 7 suggest some evidence of nonstationarity in the volume and open interest series across our sample. Figure 4 shows the daily volume (Panel A) and open interest (Panel B) over this period, revealing a pattern of increasing variance in both series. Detrending the series by the 100-day moving average removes nonstationarity in the mean, but not in the variance. Therefore, as a sensitivity check, we repeat the analysis using the natural logarithms of volume and open interest. Again, after taking logs, we decompose each series into its expected and unexpected components. The regression results are reported in the second set of columns of Table 8. For the most part, these results are quite similar to those for the raw series. The coefficients for the futures volume components are all positive and significant, and the coefficients for the open interest components are all negative although the coefficient on unexpected open interest is now insignificant.

Given this conflicting evidence on the relation between unexpected open interest and volatility, we repeat the analysis using a reduced sample beginning on April 4, 1988, five years after the contract was introduced. Figure 4 and Table 7 suggest that this subsample may avoid the nonstationarity evident in the entire sample. The final two sets of columns in Table 8 report the regression results for the reduced sample using both the raw series and log transformations. In both cases, the original results are confirmed. The positive volume-volatility relation is apparent in the reduced sample, as is the negative open interest-volatility relation. For both the raw and log series, the magnitude of the coefficient on unexpected open interest is even larger than in the full sample.

D. Asymmetries in the Volume-Volatility and Open Interest-Volatility Relations

Many empirical studies have documented volatility asymmetries. Schwert (1989, 1990) for example, finds that expected volatility increases more with negative stock market returns than it decreases with equal-sized positive returns. Bessembinder and Seguin (1993) find asymmetries in the relations between spot volatility and unexpected futures volume and open interest. To assess whether these asymmetries are apparent in the crude oil futures market, we include interactive dummy variables in our regression to allow the effects of unexpected volume and open interest on volatility to vary with the sign of the volume or open interest shock. These dummy variables equal zero for negative shocks or one for positive shocks. Table 9 reports the

results. The coefficient for the unexpected series represents the marginal impact of a negative trading activity shock. To estimate the marginal impact of a positive shock, the coefficient on the interactive term is added to the coefficient on the corresponding unexpected activity series.

The results indicate no significant asymmetry for unexpected open interest, but we do find asymmetry in the relation between volatility and unexpected volume. Specifically, the coefficient estimates indicate that the volatility increase associated with an unexpected increase in volume is 80% larger than the decrease in volatility associated with an equivalent unexpected decrease in volume. These findings are generally unchanged if we instead use either the log series of the trading activity variables or the reduced sample period.

VI. Conclusions

Our empirical results address three aspects of the impact of energy derivatives trading on the crude oil market. First, we examine the effect of introducing crude oil futures on the structure of oil market volatility. Second, we assess whether this effect differs with subsequent derivative introductions, including crude oil options and derivatives on related energy commodities. Finally, we evaluate the on-going relation between oil futures trading activity and the depth and liquidity of the crude oil market.

Our results indicate large unexpected increases in volatility for three consecutive weeks after the introduction of crude oil futures. Under our stochastic volatility model, we expect volatility to increase over this period from 6.87% to 8.14%, but realized volatility increases to 13.16%. The probability of such a large increase is just 0.2%. We also find evidence of a longer-term (more than a year) volatility increase which coincides with the growth of the energy derivative markets. It is inappropriate, however, to attribute this effect to derivatives. Derivatives activity grew over this period as a means of managing increased volatility induced by deregulation of the U.S. energy markets. Given this linkage, we cannot conclude that derivatives caused this volatility.

Following the introduction of crude oil futures, there is little evidence that subsequent derivative introductions had any effect on crude oil volatility. In particular, we find no effects following the introduction of crude oil options and no pattern in the effects across the time-series of introductions of other energy derivatives. These results are counter to the idea that subsequent derivative introductions gradually complete the market. Instead, the effects are apparent following the first introduction but disappear for subsequent introductions.

Our findings regarding the relation between futures trading activity and spot market volatility indicate that deep and liquid futures markets have a mitigating effect on volatility in the underlying market. We find a positive relation between futures volume and volatility, but we cannot determine the marginal impact of futures versus spot market volume because reliable spot volume data are unavailable. The relation between open interest and volatility, on the other hand, is large and negative. We find that the impact of volume on volatility is inversely related to both the unexpected change and long-term predictable component of open interest. Our estimates indicate that the volatility increase associated with an unexpected increase in volume is approximately 40% lower when accompanied by an unexpected increase in open interest than when the unexpected change in open interest is zero. These findings suggest that futures trading improves depth and liquidity in the underlying market, and they contradict the idea that derivatives destabilize the market.

References

Abramowitz, Milton and Irene A. Stegun, 1972, *Handbook of Mathematical Functions*, Washington, D.C.: National Bureau of Standards.

The American Petroleum Institute, 1997, *Basic Petroleum Data Book* XVIII.

Andrews, Donald W.K., 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.

Bessembinder, Hendrik, and Paul J. Seguin, 1992, Futures-trading activity and stock price volatility, *Journal of Finance* 47, 2015–2034.

Bessembinder, Hendrik, and Paul J. Seguin, 1993, Price volatility, trading volume, and market depth: Evidence from futures markets, *Journal of Financial and Quantitative Analysis* 28, 21–39.

Clark, Peter K., 1973, A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41, 135–156.

Conrad, Jennifer, 1989, The price effect of option introduction, *Journal of Finance* 44, 487–498.

Cox, Charles C., 1976, Futures trading and market information, *Journal of Political Economy* 84, 1215–1238.

Damodaran, Aswath, and J. Lim, 1991, Put listing, short sales and return generating processes, Working paper, Stern School of Business, New York University.

Damodaran, Aswath, and Marti G. Subrahmanyam, 1992, The effects of derivative securities on the markets for the underlying assets in the United States: A survey, *Financial Markets, Institutions, and Instruments* 1(5), 1–22.

Danthine, Jean-Pierre, 1978, Information, futures prices, and stabilizing speculation, *Journal of Economic Theory* 17, 79–98.

Detemple, Jerome, and Philippe Jorion, 1990, Option listing and stock returns: An empirical analysis, *Journal of Banking and Finance* 14, 781–802.

Detemple, Jerome, and Larry Selden, 1991, A general equilibrium analysis of option and stock market interactions, *International Economic Review* 32, 279–304.

Edwards, Franklin R., 1988, Futures trading and cash market volatility: Stock index and interest rate futures, *Journal of Futures Markets* 8, 421–440.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 1998a, Information and volatility linkages in the stock, bond, and money markets, *Journal of Financial Economics*, Forthcoming.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 1998b, Measuring the impact of stochastic volatility on short-horizon investment and risk management decisions, Working paper, Jones Graduate School, Rice University.

Foster, Dean P., and Dan B. Nelson, 1996, Continuous record asymptotics for rolling sample variance estimators, *Econometrica* 64, 139–174.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Grossman, Sanford J., 1988, An analysis of the implications for stock and futures price volatility of program trading and dynamic hedging strategies, *Journal of Business* 61, 275–298.

Hakansson, Nils H., 1982, Changes in the financial market: Welfare and price effects and the basic theorems of value conservation, *Journal of Finance* 37, 977–1004.

Hamilton, James D., 1994, *Time Series Analysis*, Princeton: Princeton University Press.

Hansen, Lars P., 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.

Harris, Lawrence, 1989, S&P 500 cash stock price volatilities, *Journal of Finance* 44, 1155–1176.

Harvey, Andrew C., Esther Ruiz, and Neil Shephard, 1994, Multivariate stochastic variance models, *Review of Economic Studies* 61, 247–264.

Kumar, Raman, Atulya Sarin, and Kuldeep Shastri, 1998, The impact of options trading on the market quality of the underlying security: An empirical analysis, *Journal of Finance* 53, Forthcoming.

Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.

Peck, Anne E., 1985, The economic role of traditional commodity futures markets, in Anne E. Peck, ed., *Futures Markets: Their Economic Role*, Washington, DC: American Enterprise Institute for Public Policy Research, 1–81.

Poterba, James M., and Lawrence H. Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review* 76, 1142–1151.

Powers, Mark J., 1970, Does futures trading reduce price fluctuations in the cash markets?, *American Economic Review* 60, 460–464.

Ross, Stephen A., 1976, Options and efficiency, *Quarterly Journal of Economics* 90, 75–90.

Ross, Stephen A., 1989, Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy, *Journal of Finance* 44, 1–17.

Schwert, G. William, 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1154.

Schwert, G. William, 1990, Stock volatility and the crash of '87, *Review of Financial Studies* 3, 77–102.

Skinner, Douglas J., 1989, Options markets and stock return volatility, *Journal of Financial Economics* 23, 61–78.

Skinner, Douglas J., 1990, Options markets and the information content of accounting earnings releases, *Journal of Accounting and Economics* 13, 191–211.

Stein, Jeremy C., 1987, Informational externalities and welfare-reducing speculation, *Journal of Political Economy* 95, 1123–1145.

Stoll, Hans R., and Robert E. Whaley, 1985, The new option markets, in Anne E. Peck, ed., *Futures Markets: Their Economic Role*, Washington, DC: American Enterprise Institute for Public Policy Research, 205–282.

Tauchen, George E., and Mark Pitts, 1983, The price variability-volume relationship on speculative markets, *Econometrica* 51, 485–505.

Taylor, Stephen J., 1994, Modelling stochastic volatility: A review and comparative study, *Mathematical Finance* 4, 183–204.

Taylor, Gregory S., and Raymond M. Leuthold, 1974, The influence of futures trading on cash cattle price variations, *Food Research Institute Studies* 13, 29–35.

Working, Holbrook, 1960, Price effects of futures trading, *Food Research Institute Studies* 1, 3–31.

Table 1
Introduction Dates for Energy Derivative Contracts

The table reports the introduction dates for each of the primary exchange-traded energy derivative contracts. Each of these contracts is traded at either the New York Mercantile Exchange (NYMEX) or the International Petroleum Exchange (IPE). The dates reported in the table (MM/DD/YY) indicate the first day of trading for each futures or futures option contract. Note that the heating oil futures and gasoil futures contracts (marked with an *) were introduced prior to the start of our crude oil price series.

| | | Introduction Dates | |
Commodity	Exchange	Futures	Futures Options
Crude Oil	NYMEX	3/30/83	11/14/86
Heating Oil	NYMEX	11/14/78*	6/26/87
Unleaded Regular Gas	NYMEX	12/03/84	3/13/89
Brent Crude Oil	IPE	6/23/88	5/11/89
Gasoil	IPE	4/06/81*	7/20/89
Natural Gas	NYMEX	4/03/90	10/02/92

Table 2
Summary Statistics for Crude Oil Prices

The table summarizes the price levels, returns, and volatility of returns for crude oil. We report the results based on both daily and weekly observations, and by year of the sample. The "initial price" is the closing price level on the last day of the prior year, the "return" is the total return during the year computed as the log of the closing and initial price levels, and the "standard dev." is the sample standard deviation of the daily or weekly returns during the year. We also report the average two-sided Foster/Nelson volatility estimate ("average vol.") computed using the procedure described in the text. All of the returns and volatilities are reported as annualized amounts on the basis of 252 trading days (or 52 weeks) during a year. The daily sample begins September 1, 1983 and ends December 31, 1997. The weekly sample begins February 5, 1982 and ends December 26, 1997.

Sample	Daily Observations					Weekly Observations				
	Obs.	Initial Price	Return (%)	Standard Dev. (%)	Average Vol. (%)	Obs.	Initial Price	Return (%)	Standard Dev. (%)	Average Vol. (%)
Overall	3641		−4.13	36.95	31.00	829		−3.81	33.36	29.00
1982						47	33.70	−8.17	22.88	19.55
1983	82	32.00	−30.78	9.09	9.93	52	31.30	−7.81	12.39	12.71
1984	251	28.95	−11.18	12.03	11.77	52	28.95	−11.13	11.33	11.55
1985	250	25.90	2.69	22.57	19.74	52	25.90	3.79	18.06	23.02
1986	250	26.60	−41.06	68.57	63.14	52	26.90	−45.31	68.45	62.87
1987	253	17.70	−5.79	25.89	25.11	53	17.10	−2.32	29.17	27.18
1988	255	16.70	0.59	39.02	37.07	52	16.70	0.60	34.31	34.50
1989	252	16.80	26.28	28.47	28.54	52	16.80	26.28	23.71	25.00
1990	254	21.85	26.01	59.71	57.67	52	21.85	23.54	48.47	50.72
1991	256	28.40	−41.66	55.43	38.72	52	27.65	−38.58	55.47	41.83
1992	256	18.60	5.26	19.99	18.83	53	18.80	4.19	17.96	19.30
1993	256	19.62	−30.24	23.95	23.81	52	19.62	−30.72	22.30	22.47
1994	256	14.43	20.05	29.62	28.96	52	14.43	20.37	30.31	29.41
1995	256	17.69	9.34	20.37	21.00	52	17.69	9.49	22.37	24.08
1996	257	19.45	25.71	39.09	38.21	52	19.45	25.62	32.75	31.27
1997	257	25.28	−35.45	28.51	28.43	52	25.13	−31.44	28.23	27.93

Table 3
GMM Estimation of the Stochastic Volatility Model

The table reports the GMM parameter estimates and overidentifying test statistics (*J*-statistics) for our model of the log volatility (h_t) in the crude oil market. The estimation procedure uses the moment conditions implied by the model for seasonally-adjusted, log squared returns (y_t) to estimate the mean (m_h), variance (s_h), and AR(1) parameter (f_h) of the log volatility process. The table reports the coefficient estimates (Coef.) and standard errors (S.E.) for lag lengths of $l = 10, 20, 30, 40$ for daily returns (Panel A) and $l = 12, 16, 20, 24$ for weekly returns, as well as the *J*-statistic which is distributed c_{l-1}^2 under the model. The sample period for the daily estimation is September 1, 1983 through December 31, 1997 (3641 observations) and the sample for the weekly estimation (Friday observations) is February 5, 1982 through December 26, 1997 (829 observations).

Panel A: Daily Observations

Parameter	10 lags Coef.	S.E.	20 lags Coef.	S.E.	30 lags Coef.	S.E.	40 lags Coef.	S.E.
m_h	−8.5431	0.046	−8.5368	0.045	−8.5407	0.045	−8.5411	0.045
s_h	1.3628	0.109	1.3171	0.094	1.3119	0.088	1.3050	0.083
f_h	0.9821	0.007	0.9895	0.003	0.9913	0.002	0.9925	0.001
J-statistic	8.98		13.90		17.24		19.94	
p-value	0.4390		0.7897		0.9583		0.9951	

Panel B: Weekly Observations

Parameter	12 lags Coef.	S.E.	16 lags Coef.	S.E.	20 lags Coef.	S.E.	24 lags Coef.	S.E.
m_h	−6.9281	0.099	−6.9267	0.098	−6.9198	0.099	−6.9154	0.098
s_h	1.4884	0.280	1.4661	0.275	1.5656	0.266	1.5290	0.258
f_h	0.9667	0.014	0.9702	0.010	0.9638	0.010	0.9673	0.008
J-statistic	10.66		11.21		12.96		13.51	
p-value	0.4723		0.7375		0.8406		0.9400	

Table 4
Kalman Filter Model Specification Diagnostics

The table reports diagnostic statistics for our stochastic volatility model based on the Kalman filter estimates of conditional volatility. We use these estimates to construct the seasonally-adjusted returns (r) and two sets of residuals under the model: those based on the one-sided volatility estimates (z) and those based on the smoothed estimates (z^*). Under the model, the residuals should be normally distributed with mean zero and variance one. We report the mean, variance, skewness, and excess kurtosis for each series, as well as the first-order autocorrelations of the series (x), its absolute values ($|x|$), and its squared values (x^2). Below each of the residual statistics, we report the corresponding probability value. These values are based on 5,000 simulations of the data under our model calibrated with our GMM parameter estimates. We report the probability of realizing a value lower than that observed given that the model is valid. The daily sample period (Panel A) is September 1, 1983 through December 31, 1997 (3641 observations) and the weekly sample period (Panel B) is February 5, 1982 through December 26, 1997 (829 observations).

Panel A: Daily Observations

Series	Mean	Variance	Skewness	Excess Kurtosis	Autocorrelations x	$|x|$	x^2
r	—	—	3.7072	41.3037	–0.0093	0.0703	0.2706
z	–0.0022	1.1569	0.2062	6.6813	0.0061	0.1452	0.0812
	.4517	1.0000	1.0000	1.0000	0.6347	0.9944	0.9786
z^*	0.0051	1.1563	0.0704	4.8386	–0.0130	0.0832	0.0269
	.6305	1.0000	1.0000	1.0000	0.2188	0.9984	0.9850

Panel B: Weekly Observations

Series	Mean	Variance	Skewness	Excess Kurtosis	Autocorrelations x	$|x|$	x^2
r	—	—	0.8093	10.9447	–0.1138	0.2447	0.3602
z	–0.0001	1.0765	0.5870	3.5688	–0.0279	0.0281	0.0277
	.4877	0.8988	0.9970	0.9746	0.2430	0.1710	0.2154
z^*	0.0029	1.1468	0.7080	5.7331	–0.0595	0.0091	–0.0041
	.5249	0.9998	0.9996	0.9992	0.0504	0.5611	0.7373

Table 5
Crude Oil Volatility After the Introduction of NYMEX Oil Futures Contracts

The table reports the likelihood under our stochastic volatility model of the Kalman filter volatility estimates realized just after the introduction of crude oil futures. We calibrate the model using the GMM parameter estimates over the sample of weekly observations from April 1, 1983 to December 26, 1997 (770 observations). We then use the deseasonalized returns over the entire sample (beginning February 5, 1982) to generate Kalman filter estimates of conditional volatility under the model. The table reports the actual and expected realizations for three different measures of volatility: the conditional volatility given the volatility level just prior to the introduction ($s_t|s_0$), the conditional volatility given the preceding week's (or month's, if $t>4$) volatility level ($s_t|s_{t-1}$), and the average volatility from time 1 to t given the volatility level at time 0 ($\bar{s}_{1;t}|s_0$). The expected volatilities and probability values are determined by simulations. We draw 5,000 52-week sequences (with replacement) from the sample of estimated volatility innovations (u_t) beginning one year after the introduction date. We use these innovations to construct a corresponding h_t-series to compute the distributions for the conditional, step-ahead, and average volatilities.

Weeks in Advance (t)	Conditional Volatility					Average Volatility		
	s_t	$E[s_t|s_0]$	p-value	$E[s_t|s_{t-1}]$	p-value	$\bar{s}_{1;t}$	$E[\bar{s}_{1;t}|s_0]$	p-value
0	6.87							
1	8.70	7.28	0.9704					
2	10.61	7.70	0.9912	9.15	0.9128	9.66	7.60	0.9884
3	13.16	8.14	0.9980	11.04	0.9692	10.82	7.84	0.9970
4	14.52	8.57	0.9966	13.53	0.6825	11.75	8.08	0.9978
8	13.78	10.30	0.8206	15.89	0.2641	13.64	9.06	0.9882
12	6.27	12.07	0.0434	15.25	0.0012	12.23	10.04	0.8212
16	7.80	13.71	0.0772	7.98	0.4185	11.03	11.01	0.5345
20	7.87	15.08	0.0602	9.59	0.1982	10.37	11.95	0.3421
24	7.25	16.51	0.0306	9.62	0.1288	9.69	12.86	0.1966
28	7.59	17.68	0.0318	8.94	0.2364	9.29	13.72	0.1180
32	9.25	18.78	0.0526	9.35	0.4429	9.23	14.53	0.0814
36	16.39	19.78	0.3211	10.97	0.9698	9.70	15.28	0.0784
40	10.84	20.76	0.0708	17.66	0.0374	10.12	15.95	0.0762
44	11.88	21.39	0.0880	12.56	0.3813	10.27	16.64	0.0650
48	11.83	22.03	0.0780	13.48	0.2709	10.43	17.25	0.0580
52	11.56	22.45	0.0664	13.52	0.2388	10.58	17.81	0.0530

Table 6
Crude Oil Volatility After the Introduction of Other Energy Derivative Contracts

The table reports the likelihood under our stochastic volatility model of the Kalman filter volatility estimates realized just after the introduction of each energy derivative contract. We calibrate the model using the GMM parameter estimates over the sample of daily observations from the introduction date to December 31, 1997. We then use the deseasonalized returns over the entire sample (beginning September 1, 1983) to generate Kalman filter estimates of conditional volatility under the model. The table reports the GMM parameter estimates, and the actual and expected realizations for three different measures of volatility: the conditional volatility given the volatility level just prior to the introduction ($s_t|s_0$), the conditional volatility given the preceding month's volatility level ($s_t|s_{t-20}$), and the average volatility from time 1 to t given the volatility level at time 0 ($\bar{s}_{1;t}|s_0$). The expected volatilities and probability values are determined by simulations. We draw 5,000 120-day sequences (with replacement) from the sample of estimated volatility innovations (u_t) beginning one year after the introduction date. We use these innovations to construct a corresponding h_t-series to compute the distributions for each volatility measure.

Panel A: Average Volatility

		Model Parameters			$\bar{s}_{1;t}|s_0$						
Date	Contract	m_h	s_h	f_h	$t=1$	$t=20$	$t=40$	$t=60$	$t=80$	$t=100$	$t=120$
12/03/84	Unl. Gas Futures	−8.33	0.895	0.988	0.054	0.038	0.058	0.157	0.190	0.152	0.119
11/14/86	Crude Oil Options	−8.31	0.594	0.979	0.432	0.142	0.143	0.077	0.101	0.081	0.057
6/26/87	Heating Oil Options	−8.28	0.569	0.986	0.819	0.453	0.711	0.767	0.641	0.500	0.480
6/23/88	Brent Crude Futures	−8.18	0.457	0.960	0.788	0.969	0.957	0.829	0.841	0.924	0.958
3/13/89	Unl. Gas Options	−8.28	0.580	0.978	0.923	0.826	0.849	0.843	0.902	0.853	0.815
5/11/89	Brent Crude Options	−8.29	0.539	0.967	0.883	0.569	0.835	0.699	0.626	0.525	0.439
7/20/89	Gasoil Futures	−8.31	0.616	0.988	0.821	0.785	0.544	0.439	0.359	0.288	0.309
4/03/90	Natural Gas Futures	−8.31	0.626	0.992	0.105	0.995	0.985	0.984	0.983	0.988	0.994

Panel B: Spot Volatility and Volatility Innovations

		$s_t|s_0$					$s_t|s_{t-20}$				
Date	Contract	$t=20$	$t=40$	$t=60$	$t=80$	$t=120$	$t=40$	$t=60$	$t=80$	$t=100$	$t=120$
12/03/84	Unl. Gas Futures	0.011	0.262	0.197	0.313	0.061	0.895	0.272	0.560	0.269	0.063
11/14/86	Crude Oil Options	0.288	0.048	0.122	0.104	0.107	0.051	0.420	0.245	0.087	0.477
6/26/87	Heating Oil Options	0.641	0.826	0.679	0.253	0.730	0.808	0.377	0.088	0.227	0.988
6/23/88	Brent Crude Futures	0.990	0.596	0.800	0.880	0.702	0.249	0.766	0.800	0.778	0.494
3/13/89	Unl. Gas Options	0.764	0.754	0.709	0.939	0.293	0.619	0.567	0.943	0.171	0.191
5/11/89	Brent Crude Options	0.633	0.738	0.585	0.113	0.292	0.691	0.435	0.064	0.611	0.340
7/20/89	Gasoil Futures	0.735	0.279	0.320	0.196	0.711	0.115	0.464	0.214	0.481	0.965
4/03/90	Natural Gas Futures	0.979	0.949	0.980	0.945	0.995	0.700	0.872	0.495	0.915	0.939

Table 7
Summary Statistics for Daily Crude Oil Futures Trading Activity and Spot Volatility Estimates

The table provides the mean, standard deviation (S.D.), and autocorrelation (ρ) estimates for the daily futures volume and open interest series and the Kalman filter spot market volatility estimates. The mean spot return for the period is also reported. The returns and volatilities are reported as annualized figures on the basis of 252 trading days during a year. The full sample period is from April 5, 1983 through December 31, 1997 (3686 observations). The reduced sample period does not begin until April 4, 1988 (2451 observations).

Sample	Obs.	% Spot Return	Futures Volume (Million Barrels)					Open Interest (Million Barrels)					Spot Volatility Series (%)				
			Mean	S.D.	$\rho(1)$	$\rho(2)$	$\rho(3)$	Mean	S.D.	$\rho(1)$	$\rho(2)$	$\rho(3)$	Mean	S.D.	$\rho(1)$	$\rho(2)$	$\rho(3)$
Overall	3686	−3.51	69.6	42.1	0.82	0.76	0.74	246.3	137.4	1.00	0.99	0.99	23.69	13.73	0.98	0.96	0.95
Reduced	2451	0.32	91.2	30.1	0.52	0.34	0.30	329.2	76.2	0.99	0.99	0.98	25.51	11.99	0.97	0.95	0.93
1983	188	0.68	1.7	1.7	0.79	0.67	0.68	8.8	6.0	0.99	0.99	0.99	8.7	22.96	0.96	0.93	0.88
1984	250	−11.66	7.2	3.6	0.61	0.44	0.40	28.8	4.5	0.94	0.90	0.87	9.3	43.12	0.96	0.93	0.91
1985	243	-0.38	15.2	6.3	0.29	0.17	0.17	55.1	15.1	0.97	0.95	0.94	13.7	43.74	0.93	0.86	0.79
1986	239	−38.50	32.2	13.8	0.66	0.48	0.40	99.1	20.7	0.98	0.96	0.94	46.4	17.76	0.96	0.93	0.91
1987	252	−7.19	57.9	24.1	0.64	0.44	0.30	167.0	25.4	0.97	0.95	0.93	19.5	14.97	0.92	0.83	0.76
1988	252	3.23	74.8	25.2	0.59	0.46	0.44	197.4	20.0	0.94	0.89	0.83	23.1	85.45	0.90	0.83	0.75
1989	251	23.50	81.8	23.8	0.57	0.35	0.36	227.9	18.1	0.96	0.92	0.87	26.9	98.87	0.96	0.93	0.88
1990	251	26.28	94.2	26.3	0.60	0.39	0.27	271.4	23.1	0.80	0.75	0.68	37.6	12.57	0.93	0.88	0.84
1991	253	−39.19	83.0	27.2	0.54	0.31	0.23	280.0	18.8	0.96	0.91	0.86	32.9	23.56	0.99	0.98	0.96
1992	252	2.86	83.8	25.4	0.42	0.17	0.13	324.9	20.6	0.96	0.91	0.83	17.5	85.02	0.94	0.89	0.84
1993	250	−31.85	99.5	33.9	0.52	0.32	0.32	396.5	25.5	0.97	0.93	0.88	17.3	84.14	0.89	0.82	0.75
1994	251	20.28	106.8	29.2	0.42	0.27	0.27	411.6	19.9	0.94	0.87	0.78	22.4	64.96	0.88	0.79	0.70
1995	250	11.61	94.4	30.6	0.41	0.17	0.09	351.7	17.7	0.94	0.90	0.84	20.4	45.03	0.94	0.88	0.80
1996	252	25.95	93.2	31.5	0.52	0.30	0.18	397.8	31.5	0.98	0.96	0.93	31.5	87.28	0.93	0.85	0.79
1997	252	−36.29	97.2	30.8	0.28	0.13	0.13	402.9	25.7	0.63	0.58	0.55	24.4	35.41	0.91	0.85	0.79

Table 8
Regression of Unexpected Spot Volatility on Futures Volume and Open Interest

The table reports the estimation results for a regression of unexpected spot market volatility (%) on futures trading activity variables. We proxy for unexpected spot volatility (UVOL) by differencing the step-ahead and contemporaneous volatilities obtained from the Kalman filter. For the trading activity variables, we use the 100-day moving average and the expected and unexpected components of volume (MVM, EVM, and UVM, respectively) and open interest (MOI, EOI, UOI), where the expected and unexpected components are estimated using an ARIMA(0,1,21). The regression also includes day-of-the-week dummy variables (MON, TUE, THU, and FRI) and 10 lagged unexpected volatility terms. The table reports separate results using the raw trading activity variables and the variables expressed in logs. The reported p-values are based on the hypothesis that a given coefficient equals zero, and they are computed using heteroskedasticity and autocorrelation consistent standard errors. The full sample period is from April 5, 1983 through December 31, 1997 (3686 observations) and the reduced sample excludes the period before April 4, 1988 (2451 observations).

| | Full Sample | | | | Reduced Sample | | | |
| | Raw Series | | Log Series | | Raw Series | | Log Series | |
Parameter	Coef.	p-value	Coef.	p-value	Coef.	p-value	Coef.	p-value
Intercept	−0.0009	0.9003	0.0586	0.1242	−0.0425	0.0943	−0.3834	0.1329
Futures Volumes								
MVM	0.6440	0.0295	0.0298	0.0221	1.3563	0.0076	0.0984	0.0153
EVM	0.9961	0.0001	0.0522	0.0001	1.0265	0.0001	0.0660	0.0034
UVM	1.8391	0.0001	0.1060	0.0001	1.6922	0.0001	0.1368	0.0001
Futures Open Interest								
MOI	−0.1579	0.0300	−0.0309	0.0241	−0.2222	0.0078	−0.0572	0.0138
EOI	−0.4124	0.0019	−0.0680	0.0086	−0.3535	0.0094	−0.1055	0.0264
UOI	−0.7539	0.0117	−0.0549	0.4530	−1.1233	0.0002	−0.3495	0.0002
Lagged Volatilities								
$UVOL_{t-1}$	−0.0538	0.0881	−0.0555	0.0769	−0.0523	0.1317	−0.0501	0.1500
$UVOL_{t-2}$	−0.0327	0.2395	−0.0359	0.1988	−0.0020	0.9480	−0.0018	0.9551
$UVOL_{t-3}$	0.0037	0.8995	−0.0004	0.9896	−0.0222	0.4690	−0.0223	0.4672
$UVOL_{t-4}$	−0.0050	0.8278	−0.0071	0.7605	0.0047	0.8549	0.0052	0.8376
$UVOL_{t-5}$	−0.0325	0.2454	−0.0328	0.2378	−0.0629	0.0553	−0.0605	0.0656
$UVOL_{t-6}$	−0.0023	0.9325	−0.0049	0.8532	−0.0132	0.6426	−0.0140	0.6218
$UVOL_{t-7}$	−0.0390	0.1699	−0.0404	0.1519	−0.0081	0.7874	−0.0087	0.7718
$UVOL_{t-8}$	−0.0002	0.9921	−0.0027	0.9100	−0.0109	0.7100	−0.0098	0.7407
$UVOL_{t-9}$	0.0153	0.5568	0.0132	0.6156	−0.0150	0.5400	−0.0150	0.5398
$UVOL_{t-10}$	−0.0005	0.9858	−0.0040	0.8825	0.0037	0.9028	0.0020	0.9467
Daily Dummies								
MON	0.0134	0.0974	0.0128	0.1214	0.0150	0.1293	0.0136	0.1718
TUE	0.0018	0.8225	0.0019	0.8158	−0.0065	0.5079	−0.0089	0.3672
THU	0.0015	0.8523	0.0008	0.9263	0.0014	0.8862	−0.0001	0.9893
FRI	−0.0059	0.5062	−0.0054	0.5473	−0.0112	0.3241	−0.0113	0.3180
Regression R^2	0.0600		0.0473		0.0675		0.0583	

The table reports the estimation results for a regression of unexpected spot market volatility (%) on futures trading activity variables, with the coefficients for positive and negative shocks estimated separately. We proxy for unexpected spot volatility (UVOL) by differencing the step-ahead and contemporaneous volatilities obtained from the Kalman filter. For the trading activity variables, we use the 100-day moving average and the expected and unexpected components of volume (MVM, EVM, and UVM, respectively) and open interest (MOI, EOI, UOI), where the expected and unexpected components are estimated using an ARIMA(0,1,21). The regression includes two interactive terms to capture asymmetry in the effects of trading activity shocks, UVM×DUVM and UOI×DUOI, where DUVM (DUOI) equals one if unexpected volume (open interest) is positive and zero otherwise. We also include day-of-the-week dummy variables (MON, TUE, THU, and FRI) and 10 lagged unexpected volatility terms. The table reports separate results using the raw trading activity variables and the variables expressed in logs. The reported p-values are based on the hypothesis that a given coefficient equals zero, and they are computed using heteroskedasticity and autocorrelation consistent standard errors. The full sample period is from April 5, 1983 through December 31, 1997 (3686 observations) and the reduced sample excludes the period before April 4, 1988 (2451 observations).

| | Full Sample | | | | Reduced Sample | | | |
| | Raw Series | | Log Series | | Raw Series | | Log Series | |
Parameter	Coef.	p-value	Coef.	p-value	Coef.	p-value	Coef.	p-value
Intercept	−0.0022	0.7563	0.0086	0.8199	−0.0470	0.0675	−0.5075	0.0490
Futures Volumes								
MVM	0.5711	0.0548	0.0381	0.0038	1.3558	0.0078	0.1195	0.0036
EVM	0.9234	0.0002	0.0638	0.0001	0.9770	0.0002	0.0855	0.0001
UVM	1.2383	0.0001	0.0370	0.0139	1.1293	0.0001	0.0439	0.0316
UVM×DUVM	1.0237	0.0224	0.1527	0.0001	0.9629	0.0399	0.2062	0.0001
Futures Open Interest								
MOI	−0.1719	0.0187	−0.0359	0.0095	−0.2395	0.0046	−0.0682	0.0037
EOI	−0.3823	0.0040	−0.0723	0.0048	−0.3217	0.0176	−0.1037	0.0248
UOI	−1.1604	0.0057	−0.1335	0.2160	−1.2535	0.0016	−0.3497	0.0040
UOI×DUOI	1.1177	0.1681	0.1719	0.2625	0.5082	0.5026	0.0802	0.7348
Lagged Volatilities								
$UVOL_{t-1}$	−0.0539	0.0876	−0.0565	0.0694	−0.0521	0.1338	−0.0485	0.1598
$UVOL_{t-2}$	−0.0340	0.2193	−0.0381	0.1683	−0.0029	0.9264	−0.0018	0.9527
$UVOL_{t-3}$	0.0031	0.9154	−0.0001	0.9975	−0.0228	0.4579	−0.0210	0.4948
$UVOL_{t-4}$	−0.0049	0.8311	−0.0085	0.7139	0.0045	0.8610	0.0038	0.8834
$UVOL_{t-5}$	−0.0328	0.2405	−0.0342	0.2155	−0.0634	0.0531	−0.0624	0.0567
$UVOL_{t-6}$	−0.0016	0.9514	−0.0031	0.9086	−0.0125	0.6607	−0.0110	0.6969
$UVOL_{t-7}$	−0.0393	0.1640	−0.0412	0.1423	−0.0081	0.7847	−0.0082	0.7846
$UVOL_{t-8}$	−0.0002	0.9931	−0.0025	0.9162	−0.0106	0.7177	−0.0096	0.7448
$UVOL_{t-9}$	0.0154	0.5544	0.0137	0.5980	−0.0151	0.5351	−0.0151	0.5340
$UVOL_{t-10}$	−0.0007	0.9777	−0.0034	0.9004	0.0038	0.8986	0.0045	0.8827
Daily Dummies								
MON	0.0122	0.1324	0.0117	0.1560	0.0134	0.1766	0.0120	0.2226
TUE	0.0020	0.8006	0.0035	0.6640	−0.0062	0.5301	−0.0065	0.5062
THU	0.0016	0.8469	0.0007	0.9357	0.0017	0.8689	0.0015	0.8848
FRI	−0.0072	0.4163	−0.0058	0.5138	−0.0127	0.2577	−0.0124	0.2689
Regression R^2	0.0617		0.0572		0.0687		0.0705	

Figure 1
Rolling Volatility Estimates for the Crude Oil Market

The figure shows the time-series of conditional volatility estimates for the oil market obtained using deseasonalized returns and the Foster/Nelson (1996) procedure described in the text. Panel A shows the results using daily observations over the sample period from September 1, 1983 through December 31, 1997 (3641 observations) and Panel B is based on weekly observations over the period February 5, 1982 to December 26, 1997 (829 observations). For each series, we plot both the one-sided (forward looking) and two-sided (smoothed) estimates. The estimates are annualized on the basis of 252 trading days (for the daily estimates) or 52 weeks (for the weekly estimates) per year.

Panel A: Daily Observations
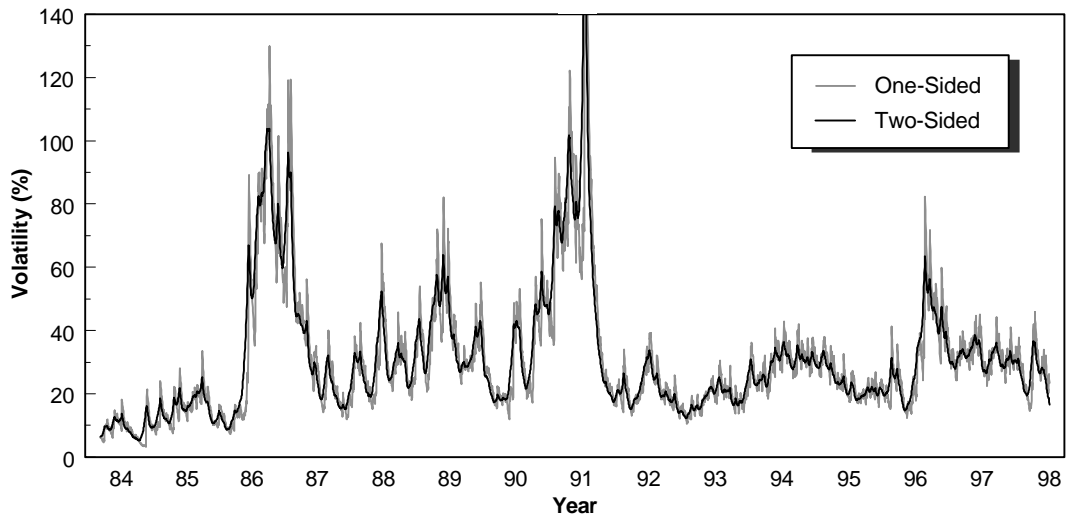

Panel B: Weekly Observations
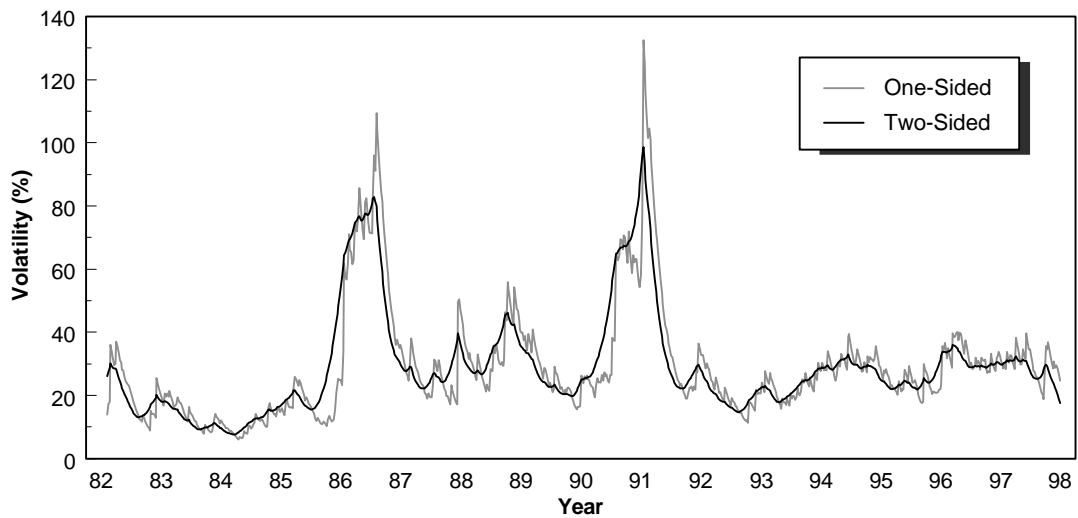
Figure 2
Stochastic Volatility Estimates for the Crude Oil Market

The figure shows the time-series of conditional volatility estimates for the oil market obtained by applying the Kalman filter to our stochastic volatility model. Panel A shows the results using daily observations over the sample period from September 1, 1983 through December 31, 1997 (3641 observations) and Panel B is based on weekly observations over the period February 5, 1982 to December 26, 1997 (829 observations). For each series, we plot both the one-sided (forward looking) and two-sided (smoothed) estimates. The estimates are annualized on the basis of 252 trading days (for the daily estimates) or 52 weeks (for the weekly estimates) per year.

Panel A: Daily Observations
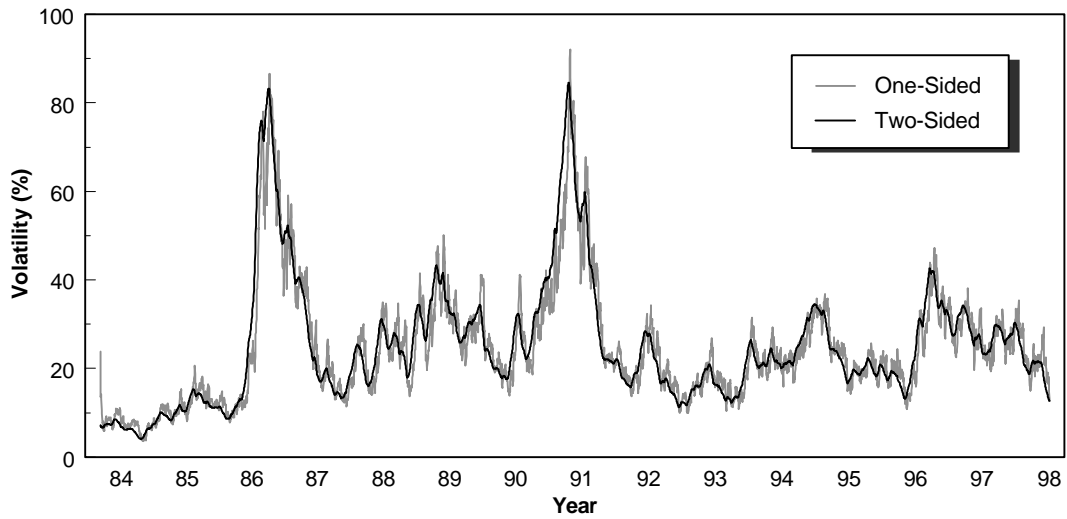
Panel B: Weekly Observations
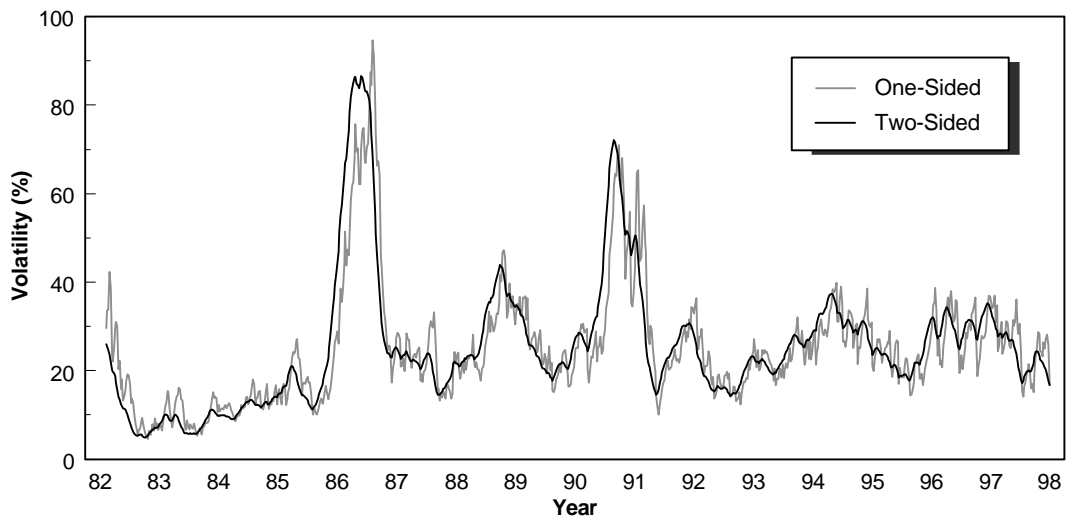
Figure 3
Annual Volumes of Crude Oil Production and NYMEX Crude Oil Futures Trading Activity
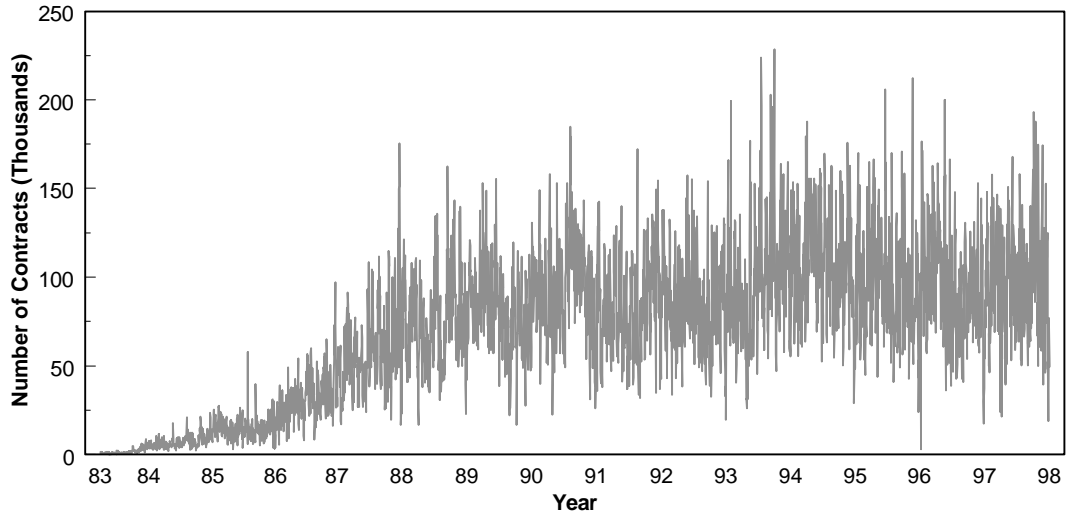
The figure shows the annual world production of crude oil and the annual trading volume of NYMEX crude oil futures contracts, reported by year since 1980. Both the production and trading volumes are stated in billions of barrels. For the futures contract, the volume amounts are aggregated across all of the available delivery months and the data begin on the contract's introduction date, March 31, 1983. The production volume for 1997 is an estimated amount.
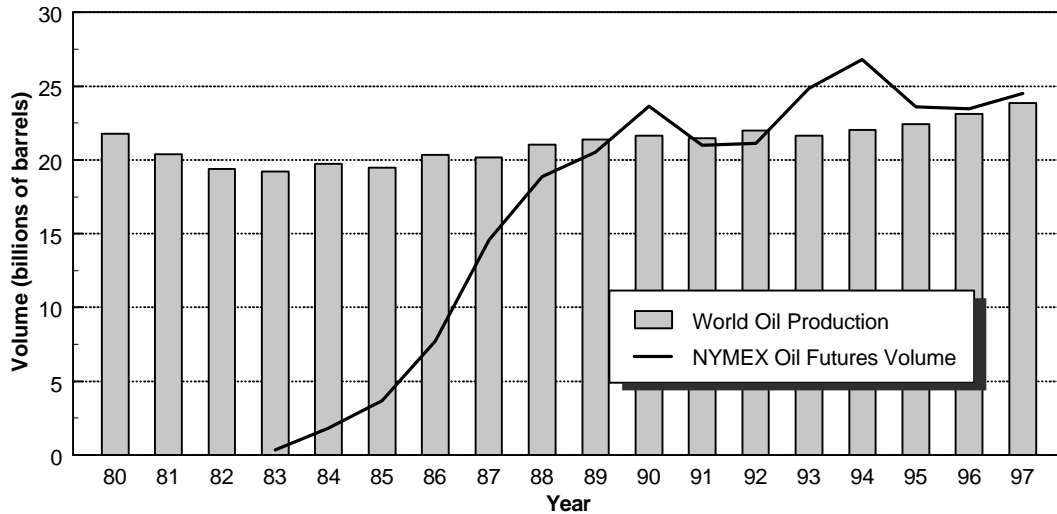
Figure 4
Daily NYMEX Futures Contract Trading Volume and Open Interest

The figure shows the daily trading volume (Panel A) and open interest (Panel B) for NYMEX crude oil futures contracts. The data are from the inception of trading on March 31, 1983 through December 31, 1997. Both trading volume and open interest amounts are aggregated across all of the available delivery months. The futures contract size is 1000 barrels.

Panel A: Trading Volume

Panel B: Open Interest