CEEPR Center for Energy and Environmental Policy Research

Volatility in Natural Gas and Oil Markets

by
Robert S. Pindyck

03-012 WP June 2003

A Joint Center of the Department of Economics, Laboratory for Energy
and the Environment, and Sloan School of Management

VOLATILITY IN NATURAL GAS AND OIL MARKETS*

by

Robert S. Pindyck

Massachusetts Institute of Technology
Cambridge, MA 02142

This draft: June 16, 2003

Abstract: Using daily futures price data, I examine the behavior of natural gas and crude oil price volatility since 1990. I test whether there has been a significant trend in volatility, whether there was a short-term increase in volatility during the time of the Enron collapse, and whether natural gas and crude oil price volatilities are interrelated. I also measure the persistence of shocks to volatility and discuss its implications for gas- and oil-related contingent claims.

JEL Classification Numbers: G13; L71, Q40

Keywords: Natural gas, oil markets, price volatility, Enron, commodity markets.

---

*I am grateful to M.I.T.'s Center for Energy and Environmental Policy Research for its financial support of the research leading to this paper, to Mr. Scott Byrne and the New York Mercantile Exchange for providing futures market data, and to Martin Minnoni for his outstanding research assistance.

1. Introduction

This paper examines the behavior of natural gas and crude oil price volatility since 1990. Prices of crude oil and especially natural gas rose sharply (but temporarily) during late 2000, and natural gas trading was buffeted by the collapse of Enron in late 2001 suggesting to some that volatility in these markets has increased. Whether or not this is true, volatility has been high, and (like prices themselves) fluctuates dramatically.

Understanding the behavior of price volatility in natural gas and crude oil markets is important for several reasons. Persistent changes in volatility can affect the risk exposure of producers and industrial consumers of natural gas and oil, and alter the incentives to invest in natural gas and oil inventories and facilities for production and transportation. Likewise, volatility is a key determinant of the value of commodity-based contingent claims, whether financial or "real." Thus understanding the behavior of volatility is important for derivative valuation, hedging decisions, and decisions to invest in physical capital tied to the production or consumption of natural gas or oil.

In addition, including volatility as a "market" variable can help us better understand the short-run market dynamics for natural gas, oil, and commodities in general. As I have discussed in an earlier paper (2002), volatility should in principle affect the demand for storage, and should also affect the total marginal cost of production by affecting the value of firms' operating options and thus the opportunity cost of current production. In particular, greater volatility should lead to an increased demand for storage, and an increase in both spot prices and marginal convenience yield.[1] Thus, changes in volatility may help explain changes in these other variables.

With this in mind, I address the following questions: First, has natural gas and/or crude oil price volatility increased or decreased in a significant way since 1990, and in particular, are there measurable trends in volatility? Related to this, have the events surrounding the collapse of Enron affected volatility, i.e., was there a significant short-term increase in volatility around the time of the collapse? Second, is there evidence that natural gas and crude oil volatilities are

[1] Using weekly data for the petroleum complex over the period 1984 to 2001, I show in Pindyck (2002) that the theoretical relationships between volatility and other market variables are well supported by the data for heating oil, but less so for crude oil and gasoline. The role of volatility as a determinant of the opportunity cost of current production has also been spelled out and tested by Litzenberger and Rabinowitz (1995). Finally, for a general introduction to the interrelationships between price, inventories, and convenience yields, see Pindyck (2001).

interrelated, i.e., can changes in one help predict changes in the other? Third, although volatility clearly fluctuates over time, how persistent are the changes? If changes are very persistent, then they will lead to changes in the prices of options and other derivatives (real or financial) that are tied to the prices of these commodities. If, on the other hand, changes in volatility are highly transitory, they should have little or no impact on market variables or on real and financial option values. Finally, extending the work in Pindyck (2002), I revisit the question off whether changes in volatility are predictable.

To address these questions, I use daily futures price data for natural gas and crude oil, along with data on interest rates, to infer daily spot prices and daily values of the net marginal convenience yield. From the log price changes (adjusted for non-trading days) and marginal convenience yield, I calculate daily and weekly returns from holding each commodity. I then estimate volatility in three different ways.

First, using a five-week overlapping window, I estimate weekly series for price volatility by calculating sample standard deviations of (adjusted) log price changes. As Campbell et. al. (2001) point out in their study of stock price volatility, in addition to its simplicity, this approach has the advantage that it does not require a parametric model describing the evolution of volatility over time.[2] Second, I estimate series for conditional volatility by estimating GARCH models of the weekly returns on the commodities, and I compare the volatility estimates from these models to the sample standard deviations. Third, I estimate a *daily* series for conditional volatility by estimating GARCH models of the daily returns on the commodities.

I study the behavior of volatility in two different ways. First, using the estimated weekly sample standard deviations, I test for the presence of time trends, I test whether volatility was significantly greater during the period of the Enron collapse, and I examine whether gas (oil) volatility is a significant predictor of oil (gas) volatility. I also use these series to estimate the persistence of changes in volatility. Second, I address these same questions using weekly and daily GARCH models of commodity returns. For example, I can test whether a time trend or a dummy variable for the Enron period is a significant explanator of volatility (and/or an explanator of returns) in the context of the GARCH framework. Likewise, the estimated

[2] Schwartz (1997) and Schwartz and Smith (2000) have shown how futures and spot prices can be used can be used to estimate the parameters of a mean-reverting price process and derive values of commodity-based options. That approach also yields implicit time-varying estimates of volatility.

coefficients from the variance equation of each GARCH model provide a direct estimate of the persistence of volatility shocks.

I focus on the volatility of prices, but one might ask whether that is the most relevant measure of volatility. Putting aside issues of data availability, one could instead examine the volatility of consumption, production, or inventories. That would, in fact, be appropriate if the objective was to explain the motivations for holding inventories, and in particular the role of production and/or consumption smoothing, and production-cost smoothing, as determinants of inventory demand.[3] My concern, however, is the overall market, and the spot price is the best single statistic for overall market conditions. Spot price volatility reflects the volatility of current production, consumption, and inventory demand, as well as volatility in the expected future values of these variables.[4]

The results of this study can be summarized as follows: (1) I find a statistically significant positive trend in volatility for natural gas (but not for crude oil). However, this trend is not significant in economic terms; over a ten year period, it amounts to about a 3-percent increase in volatility. (2) There is no statistically significant increase in volatility during the period of the Enron collapse. (3) The evidence is mixed as to the interrelationship between crude oil and natural gas returns and volatilities. Using daily data, crude oil returns are a significant predictor of natural gas returns (but not the other way around), and crude oil volatility is a significant predictor of natural gas volatility. Using weekly data, however, these results are less clear-cut. (4) Shocks to volatility are generally short-lived for both natural gas and crude oil. Volatility shocks decay (i.e., there is reversion to the mean) with a half-life of about 5 to 10 weeks.

In the next section, I discuss the data and the calculation of returns and weekly sample standard deviations. All of the empirical work is presented in Section 3. Section 4 concludes.

2. The Data

I begin with natural gas and crude oil futures price data covering the period May 2, 1990 through February 26, 2003. (The start date was constrained by the beginning of active trading in natural gas futures.) To obtain a weekly series for volatility, I use the sample standard deviations

---

[3] Pindyck (1994) addresses these issues; also see Eckstein and Eichenbaum's (1985) study of crude oil inventories.

[4] Furthermore, one cannot actually put aside issues of data availability. Although weekly data are available for U.S. production, consumption, and inventories of natural gas and crude oil, daily data are not.

of adjusted daily log price changes in spot and futures prices. However, I also obtain estimates of conditional volatility from GARCH models of weekly and daily returns. These calculations are discussed in detail below.

2.1. Spot Prices and Weekly Volatility

For each commodity, daily futures settlement price data were compiled for the nearest contract (often the spot contract), the second-nearest contract, and the third-nearest. These prices are denoted by $F1$, $F2$, and $F3$. The spot price can be measured in three alternative ways. First, one can use data on *cash prices*, purportedly reflecting actual transactions. One problem with this approach is that daily cash price data are usually not available. A second and more serious problem is that a cash price can include discounts and premiums that result from relationships between buyers and sellers, and need not even reflect precisely the same product (including delivery location) that is specified in the futures contract. A second approach, which avoids these problems, is to use the price on the spot futures contract, i.e., the contract that expires in month t. But this approach also has problems, because the spot contract often expires before the end of the month. In addition, active spot contracts do not always exist for each month.

The third approach, which I use here, is to infer a spot price from the nearest and the next-to-nearest active futures contracts. This is done for each day by extrapolating the spread between these contracts backwards to the spot month as follows:

$$P_t = F1_t (F1_t / F2_t)^{n_{0t}/n_1} \qquad (1)$$

where P_t is the spot price on day t, $F1_t$, and $F2_t$ are the prices on the nearest and next-to-nearest futures contracts, and n_{0t} and n_1 are the number of days from t to the expiration of the first contract, and the number of days between the expiration dates for the first and second contracts.

Given these daily estimates of spot prices, I compute weekly estimates of volatility. To do this, one must take into account weekends and other non-trading days. If the spot price of the commodity followed a geometric Brownian motion, then this could be done simply by dividing the log price changes by the square root of the number of intervening days (e.g., three days in the case of a week-end), and then calculating the sample variance. However, as is well known, on average the standard deviation of n-day log price changes is significantly less than $\sqrt{n}$ times the standard deviation of one-day log price changes, when n includes non-trading days.[5] To deal with this, I sort the daily price data by intervals, according to the number of days since the last trading day. For example, if there were no holidays in a particular period, prices for Tuesday, Wednesday, Thursday, and Friday would all be classified as having an interval of one day, since there was always trading the day before. Monday, on the other hand, would be classified as an interval of three days, because of the two-day weekend. Because of holidays, some prices could also be assigned to intervals of two, four, or even five days (the latter occurring when a weekend was followed by a two-day holiday).

For each interval set, I calculate the sample standard deviation of log price changes for the entire sample for each commodity. Let $\hat{s}_n$ denote this sample standard deviation for log price changes over an interval of n days. I then compute the "effective" daily log price change for each trading day as follows:

$$\delta_\tau = \frac{(\log P_\tau - \log P_{\tau-n})}{\hat{s}_n / \hat{s}_1}. \tag{2}$$

For each week, I then compute a sample variance and corresponding sample standard deviation using these daily log price changes for that week and the preceding four weeks:

$$\hat{\sigma}_t = \sqrt{\frac{1}{N-1}\sum_{\tau=1}^{N}(\delta_{t\tau} - \bar{\delta}_t)^2}\ , \tag{3}$$

where N is the number of "effective" days in the five-week interval. Eqn. (3) gives the sample standard deviation of *daily* percentage price changes; to put it in weekly terms, I multiply by $\sqrt{30/4} = \sqrt{7.5}$. The resulting weekly series is my measure of volatility, σ_t.

2.2. Daily and Weekly Returns

As discussed above, I obtained weekly estimates of volatility from the sample standard deviations of (adjusted) log daily price changes over five-week intervals. An important advantage of this approach (besides its simplicity) is that it does not require a parametric model describing the evolution of volatility over time. However, there are also disadvantages. The first is that the use of overlapping intervals introduces serial correlation as an artifact, which makes it

---

[5] If P_t follows a geometric Brownian notion, $p_t = \log P_t$ follows an arithmetic Brownian notion, so that $\text{var}\,(p_{t+n} - p_t) = n\text{var}(p_{t+1} - p_t)$.

more difficult to discern the time-series properties of volatility. A second disadvantage is that even the use of a five-week interval yields imprecise estimates of the sample standard deviation. Hence, I also estimate volatility from GARCH models of commodity returns. These models can include parameters that test for time variation (such as trends or an “Enron effect”), and have the additional advantage that the time-series properties of volatility (the ARCH and GARCH components, which determine the persistence of shocks to volatility) are estimated along with the volatility itself.

Marginal Convenience Yield. To calculate the total return on the physical commodity, I need to know the net marginal convenience yield at each point in time, i.e., the value of the flow of production- and delivery-facilitating services from the marginal unit of inventory, net of storage costs. Denoting net marginal convenience yield by ψ_t, it can be measured from spot and futures prices as follows:

$$\psi_t = (1 + r_t)P_t - F_{1t}, \tag{4}$$

where F_{1t} is the futures price at time t for a contract maturing at time t + 1, and r_t is the one-period riskless interest rate. I calculate values of ψ_t for every trading day using the futures price corresponding as closely as possible to a 1-month interval from the spot price. (When there are few or no trades of the nearest futures contract, as sometimes occurs with natural gas, the next-to-nearest contract is used instead.) Also, I use the yield on 3-month Treasury bills, adjusted for the number of days between P_t and F_{1t}, for the interest rate r_t.

In what follows, I will use both daily and weekly series for the marginal convenience yield. I therefore convert the net marginal convenience yields calculated above into daily terms, i.e., dollars per unit of commodity per day. For days followed by another trading day (e.g., a Monday), I do this simply by dividing the values of ψ_t calculated above by the number of days between P_t and F_{1t}. For days followed by n non-trading days, I multiply these values by n+1. (Thus for a Friday, which is typically followed by n = 2 non-trading days, the convenience yield is the flow of value from holding a marginal unit of inventory over the next 3 days.) As explained below, this daily series is used to compute daily returns from holding the commodity.

To obtain a weekly series, I use the calculated values of ψ_t for the Wednesday of each week, and multiply those values by 7 so that the convenience yield is measured in dollars per unit of commodity per week.[6] I then use this weekly series to calculate weekly returns from holding the commodity.

Calculating Returns. The total return from holding a unit of a commodity over one period is the capital gain or loss over that period, plus the "dividend," which is the net marginal convenience yield, i.e., the flow of benefits to producers or consumers from holding the marginal unit of inventory, net of storage costs.[7] I calculate a series of daily (weekly) returns by summing the "effective" daily log price changes over each day (week) and adding to this the estimate of daily (weekly) convenience yield. The weekly return, for example, is calculated as:

$$R_t = \sum_{\tau=1}^{T} \delta_\tau + \psi_t \tag{5}$$

where δ_τ is given by eqn. (2), and T is the number of days in the week. A series for the daily return is calculated by using the effective daily log price change for each effective trading day and adding the *daily* flow of marginal convenience yield.[8]

3. The Behavior of Volatility and Prices

In this section, I examine the behavior of natural gas and crude oil price volatility using the weekly time series of (overlapping) sample standard deviations of adjusted log price changes. I also discuss the behavior of spot prices themselves, and the relationship of price levels to volatility. Based on these time series alone, there is little evidence of a trend in volatility, nor is there evidence of a significant increase in volatility for natural gas or crude oil during the period of the Enron collapse. In addition, changes in volatility appear to be highly transitory, with a half-life of several weeks.

As an alternative way of measuring volatility, I estimated GARCH models of the weekly returns to holding the commodity, and from these models estimated conditional standard deviations on a weekly basis. I test for changes in volatility over time by introducing a time trend and an Enron dummy variable in the variance equations of the GARCH models. I show that the results are similar to those obtained from the weekly sample standard deviations.

---

[6] If Wednesday is a holiday, I use Thursday's price.

[7] Thus, the price of a storable (and stored) commodity should equal the present value of the expected flow of marginal convenience yield. This model of price has been tested in Pindyck (1993).

[8] Note that because I use effective trading days, my daily series will have about 20 data points per month.

Finally, I used the original daily adjusted return series to estimate daily GARCH models. These provide estimates of conditional standard deviations on a daily basis, and are also used to test for time trends and an Enron effect, and to estimate the persistence of changes in volatility.

3.1. Weekly Sample Standard Deviations

Figures 1 and 2 show the weekly series for the spot price and volatility of natural gas and crude oil, where volatility is measured as the sample standard deviations of adjusted log price changes. Note that for both commodities, volatility is high, and is itself volatile. The mean values of volatility are 12.8 percent per week for natural gas and 5.9 percent per week for crude oil; the corresponding standard deviations are 7.0 percent for natural gas and 3.2 percent for crude oil. Natural gas and crude oil volatilities are correlated, but only weakly so; the coefficient of correlation for the two series is 0.169. As expected, both volatility series have high degrees of skewness and kurtosis; the skewness coefficient and degree of kurtosis are 1.60 and 6.99 respectively for natural gas, and 1.76 and 7.77 for crude oil. For the log of volatility, these coefficients are –0.46 and 3.99 respectively for natural gas, and 0.23 and 2.84 for crude oil. For both commodities, these coefficients are roughly consistent with a normal distribution for the log of volatility. However, a Jarque-Bera test rejects normality in both cases at the 1-percent level).

As Figures 1 and 2 illustrate, periods of unusually high volatility tend to accompany sharp increases in the spot price. In the case of crude oil, for example, volatility was high in late 1990 and early 1991 following the Iraqi invasion of Kuwait, as spot prices reached $40 per barrel. However, there were also periods of high volatility that accompanied unusually low spot prices, e.g., during 1998 for both commodities. Overall, volatility and price are moderately correlated; the correlation (in levels) is .27 for natural gas and .37 for crude oil.

Was volatility unusually high during the period of Enron collapse? The Enron bankruptcy sharply reduced spot and forward trading in natural gas and electricity, and also led to speculation over net long and short positions in natural gas. This was likely to have caused increased uncertainty over natural gas prices, which could have spilled over into crude oil. Pinpointing the beginning of the Enron collapse is difficult, but clearly by September 2001 analysts began questioning Enron's valuation. (On September 26, 2001, Kenneth Lay made his famous announcement to employees that the stock is "an incredible bargain.") On October 16, 2001, Enron reported a $638 million third-quarter loss and disclosed a $1.2 billion reduction in shareholder equity. Further financial statement revisions were announced during October and November, and Enron filed for Chapter 11 bankruptcy protection on December 2.

I defined the period of the Enron collapse as August 29 to December 5, 2001, and created a dummy variable that takes on the value 1 during this period and 0 otherwise. Figure 3 shows natural gas and crude oil price volatility from the middle of 2000 through the middle of 2002, with the Enron period shaded. Natural gas volatility reached a peak during this period of 38 percent per week, and crude oil volatility was also unusually high. I examine the significance of these increases in volatility in the context of forecasting regressions.

I have shown elsewhere (Pindyck (2002)), using data for crude oil, heating oil and gasoline, that price volatility cannot be forecasted using market variables for that commodity (such as production, inventories, or convenience yields), or using macroeconomic variables (such as interest rates). As mentioned above, there is a *contemporaneous* positive correlation between volatility and the price level itself (and thus between volatility and the contemporaneous convenience yield), but little or no correlation with *lagged* prices or other market variables. As discussed below, the only variables that do have forecasting power for volatility are its own lagged values (i.e., volatility can be modeled as an ARMA process), and possibly lagged values of volatility for another commodity (e.g., crude oil in the case of natural gas).

Table 1 shows simple forecasting regressions for volatility. In columns (1) and (4), the explanatory variables are 6 lags of volatility and the Enron dummy variable. For natural gas, the Enron dummy is marginally significant, and for crude oil it is insignificant. Even for natural gas, however, it has little economic significance, temporarily adding about 1.5 percent to an average volatility of about 20 percent. In columns (2) and (5), a time trend is added; in both cases it is insignificant, and has almost no effect on the other estimated coefficients. Finally, columns (3) and (6) test whether lagged values of crude oil volatility help explain natural gas volatility, and vice versa. For natural gas, the answer is ambiguous: an F-test on the joint significance of the lagged crude oil volatility terms in column (3) has a value of 1.84, which is significant at the 10-percent level. Lagged values of natural gas volatility, however, are not significant explanators of crude oil volatility: the corresponding F-test for column (6) yields a value of 1.39.[9]

[9] Note that when lagged values of volatility for the second commodity are added to the regression, the Enron dummy becomes insignificant. However, this can simply reflect the fact that volatility for both commodities was unusually high during the Enron period.

The bottom of Table 1 shows the sum of the autoregressive coefficients for each equation, along with the implied half-life for shocks to volatility. The half-life is about five to six weeks for natural gas, and eleven to twelve weeks for crude oil. Thus, although volatility itself is quite volatile, these simple autoregressive models show shocks to volatility to be quite transitory, particularly for natural gas.

The volatility series shown in Figures 1 to 3 and used in the regressions in Table 1 suffer from two main problems. First, the sample standard deviations are estimated from daily log price changes for *overlapping* five-week intervals. Thus, the series are serially correlated by construction. Second, even with five-week intervals, each sample standard deviation is based on at most twenty-five observations. One way to get around these problems is to estimate GARCH models of the commodity returns themselves. I turn to that next.

3.2. GARCH Models of Weekly Returns

The models I estimate have the following form. The equation for the weekly return to holding the commodity is given by:

$$\mathrm{RET}_t = a_0 + a_1\mathrm{TBILL}_t + a_2\sigma_t + a_3\mathrm{ENRON}_t + a_4\mathrm{TIME}_t + \sum_{j=1}^{11} b_j\mathrm{DUM}_{jt} + \varepsilon_t, \qquad (6)$$

where DUM_{jt} are monthly dummy variables. In this equation, the treasury bill rate should affect the return because it is a large component of the carrying cost of holding the commodity. Likewise, we would expect the return to increase with its own riskiness, so σ_t, the standard deviation of the error term ε_t, is included in the equation. Finally, I also include the Enron dummy variable and a time trend to test for any systematic time variation in returns.

The second equation explains the variance of ε_t as a GARCH (p,q) process:

$$\sigma_t^2 = \alpha + \sum_{j=1}^{p} \alpha_j \varepsilon_{t-j}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 + \gamma_1\mathrm{ENRON}_t + \gamma_2\mathrm{TIME}_t \qquad (7)$$

The Enron dummy and a time trend are included to test for time variation in volatility.

Table 2 shows maximum likelihood estimates of this model. Because the return includes the current and previous week's price, the model is estimated with and without a first-order moving average error term in Equation (6). In all cases the number of lags in eqn. (7) is chosen to minimize the Akaike information criteria.

The results for crude oil (in columns 3 and 4 of Table 2) are consistent with the basic theory of commodity returns and storage. Returns have a strong positive dependence on the interest rate and on volatility (as measured by the standard deviation of ε_t). In the case of natural gas, however, both the interest rate and volatility are statistically insignificant in the returns equation. The coefficient on the time trend in the returns equation is insignificant, but the corresponding coefficient in the variance equation is positive and significant in all cases. For both commodities, the variance of returns is positively related to the Enron dummy, but in all cases the coefficient is statistically insignificant. Thus, I find a clear statistically significant positive trend in volatility for both gas and oil, but no separate impact of the Enron events. However, this trend is not economically significant. For natural gas, for example, the time trend coefficient is around 7×10^{-7}, which implies a 10-year increase in the average variance of .00035. The mean value of volatility (standard deviation of returns) is about .13 for natural gas, so the mean variance is about .017. Thus the trend represents a roughly 2-percent increase in the variance over a decade.

Table 2 also shows estimates of the half-life of volatility shocks. This is determined by the sum of the ARCH and GARCH coefficients in the variance equation, i.e.,

$$\text{Half-life} = \log(.5) / \log(\sum \alpha_j + \sum \beta_j). \tag{8}$$

The half-life of volatility shocks for natural gas is about seven to ten weeks, and for crude oil is about seven to eight weeks. These numbers differ slightly from the estimates in Table 1. For crude oil, the estimated half-life is about seven to eight weeks, which is shorter than the estimates in Table 1. Overall, however, shocks to volatility again appear transitory for both natural gas and crude oil.

We can compare the volatility estimates from these GARCH models (i.e., the conditional standard deviation of ε_t) with the sample standard deviations. Using the GARCH models that include the moving average term, i.e., columns 2 and 4 of Table 2, the simple correlation of the two volatility series is .593 for natural gas and .665 for crude oil. Figure 4 shows the two volatility series for natural gas. The two series generally track each other, but the GARCH volatility is lower on average (a mean of 8.7 percent vs. 12.8 percent for the sample standard deviation) and has a higher degree of kurtosis.

3.3. GARCH Models of Daily Returns

An advantage of estimating GARCH models of weekly returns is that the resulting estimates of the conditional standard deviations can be compared to the weekly estimates of the sample standard deviations. However, these weekly models do not make use of all of the available daily data. I turn now to the use of that data for the estimation of GARCH models of daily returns. These models also take the form of eqns. (6) and (7), except that I do not include monthly dummy variables in the returns equation. Once again, in each regression the number of lags is chosen to minimize the Akaike information criterion.

The results are shown in Table 3. As with the weekly GARCH models, the results for crude oil are consistent with the theory of commodity returns and storage, but the results for natural gas are not. Crude oil returns have a strong positive dependence on the interest rate and on volatility, but both the interest rate and volatility are statistically insignificant in the equation for natural gas returns. And as with the weekly models, there is no statistically significant impact of the Enron events on volatility for either commodity. The time trend for volatility is now only marginally significant for natural gas, and insignificant for crude oil, but even for natural gas, it is only of marginal economic significance. (Using an average estimate of 5.35×10^{-8} for the trend coefficient, the 10-year trend increase in the variance of daily returns would be .00020, which is about 9 percent of the mean daily variance of .00228.)

The estimates of the half-life of volatility shocks vary across the different specifications, but overall are not very different from the results in Tables 1 and 2. The half-life is about 6 to 9 weeks for natural gas, and 3 to 11 weeks for crude oil. Once again, shocks to volatility appear to be fairly transitory for both commodities.

3.4. Returns and Volatilities Across Markets

I turn next to the interrelationship between crude oil and natural gas returns and volatilities. The results in Table 1, based on the 5-week sample standard deviations, provided some evidence that crude oil volatility has some predictive power with respect to natural gas volatility (but not the other way around). Here I explore this further by running Granger causality tests between gas and oil using the sample standard deviations, and the weekly and daily volatilities from the GARCH models. I also run these test on weekly and daily gas and oil returns. These tests are simply F-tests of the exclusion restrictions $b_1 = b_2 = ... = b_L = 0$ in the regression equation

$y_t = a_0 + \sum_{i=1}^{L} a_i y_{t-i} + \sum_{i=1}^{L} b_i x_{t-i}$. A failure to reject these exclusion restrictions is a failure to reject the hypothesis that x_t Granger-causes y_t. When running these tests, I use 2, 4, and 6 lags for the weekly regressions, and 4, 6, 10, 14, 18, and 22 lags for the daily regressions.

The results are shown in Table 4. The first two panels show tests for the weekly and daily returns. The weekly returns show no evidence of causation in either direction, but for the daily returns, I can reject the hypothesis that there is no causality from oil to gas. Given that oil prices are determined on a world market, if there is causality in either direction we would expect it to run from oil to gas, and not the other way around.

The next three panels show test results for volatility. The tests based on the weekly sample standard deviations and the daily GARCH models show evidence of causality from oil to gas, and not from gas to oil, as expected. However, the results using the volatility estimates from the weekly GARCH models show just the opposite. But note that the simple correlations of the oil and gas volatilities are much higher for the weekly sample standard deviations and the daily GARCH estimates (.170 and .146, respectively) than for the weekly GARCH estimates (.092), so I discount these latter results. Overall, these tests (along with the regressions in Table 1) provide some evidence that crude oil volatility is a predictor of natural gas volatility.

4. Summary and Conclusions

Using daily futures price data, I examined the behavior of price volatility for natural gas and oil over the period May 1990 to February 2003. The results can be summarized as follows.

First, there is evidence of a statistically significant positive time trend in volatility for natural gas, and to a lesser extent for oil. This trend, however, is small, and not of great economic significance. Given the fairly limited length of my sample, there are certainly no conclusions that can be drawn about long-term trends. As for the events surrounding the demise of Enron, they do not appear to have contributed to any significant increase in volatility.

Second, there is some evidence that crude oil volatility and returns has predictive power for natural gas volatility and returns, but not the other way around. But this predictive power is quite limited; for practical purposes, volatility can be modeled as a pure ARMA process.

Third, although volatility fluctuates considerably, shocks to volatility are short-lived, with a half-life on the order of 5 to 10 weeks. This means that fluctuations in volatility could certainly affect the values of financial gas- or oil-based derivatives (such as options on futures contracts), because such derivatives typically have a duration of only several months. But fluctuations in volatility should not have any significant impact on the values of most real options (e.g., options to invest in gas- or oil-related capital), or the related investment decisions. Of course these fluctuations might lead one to think that financial or real options should be valued using a model that accounts for stochastic volatility. However, the numerical analyses of Hull and White (1987), among others, suggests that treating volatility as non-stochastic will make little quantitative difference for such valuations.

Sharp (but temporary) increases in the prices of crude oil and natural gas, along with the collapse of Enron, have created a perception that volatility has increased significantly, increasing the risk exposure of energy producers and consumers. This does not seem to be the case. The increases in volatility that I measure are too small to have economic significance, and fluctuations in volatility are generally short-lived.

References

Campbell, John Y., Burton Malkiel, Martin Lettau, and Yexiao Xu, "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, Feb. 2001, **56**, 1-43.

Eckstein, Zvi, and Martin S. Eichenbaum, "Inventories and Quantity-Constrained Equilibria in Regulated Markets: The U.S. Petroleum Industry, 1947—1972," in T. Sargent, ed., *Energy, Foresight, and Strategy*, Washington, D.C.: Resources for the Future, 1985.

Hull, John, and Alan White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, June 1987, **42**, 281—300.

Litzenberger, Robert H., and Nir Rabinowitz, "Backwardation in Oil Futures Markets: Theory and Empirical Evidence," *Journal of Finance*, 1995, **50**, 1517—1545.

Pindyck, Robert S., "The Present Value Model of Rational Commodity Pricing," *The Economic Journal*, 1993, **103**, 511—530.

Pindyck, Robert S., "Inventories and the Short-Run Dynamics of Commodity Prices," *The RAND Journal of Economics*, Spring 1994, **25**, 141—159.

Pindyck, Robert S., "The Dynamics of Commodity Spot and Futures Markets: A Primer," *The Energy Journal*, 2001, **22**, 1—29.

Pindyck, Robert S., "Volatility and Commodity Price Dynamics," M.I.T. Center for Energy and Environmental Policy Research Working Paper, August, 2002.

Schwartz, Eduardo S., "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging," *The Journal of Finance*, July 1997, **52**, 923-973.

Schwartz, Eduardo S., and James E. Smith, "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science*, 2000.

Table 1: Forecasting Equations for Volatility

Dep. Var.	(1) **NG**	(2) **NG**	(3) **NG**	(4) **CRUDE**	(5) **CRUDE**	(6) **CRUDE**
Const.	0.0149	0.014	0.015	0.0032	0.0027	0.0018
	(5.71)	(4.62)	(4.38)	(3.40)	(2.40)	(1.40)
NGSIG (-1)	1.0550	1.0540	1.0445			-0.0245
	(28.12)	(28.06)	(27.69)			(-1.67)
NGSIG (-2)	-0.0906	-0.0904	-0.0799			0.0281
	(-1.74)	(-1.73)	(-1.53)			(1.38)
NGSIG (-3)	-0.0606	-0.0607	-0.0574			0.0141
	(-1.17)	(-1.17)	(-1.11)			(0.70)
NGSIG (-4)	0.1969	0.1968	0.1906			-0.004
	(3.81)	(3.81)	(3.68)			(-0.20)
NGSIG (-5)	-0.5118	-0.5115	-0.5119			0.0065
	(-9.82)	(-9.81)	(-9.81)			(0.32)
NGSIG (-6)	0.2930	0.2916	0.2924			-0.0118
	(7.84)	(7.79)	(7.80)			(-0.80)
ENRON	0.0149	0.0142	0.0132	0.0046	0.0042	0.0035
	(2.00)	(1.88)	(1.73)	(1.63)	(1.44)	(1.18)
TIME		3.54E-06	4.18E-06		-2.25E-06	1.83E-06
		(0.60)	(0.70)		(0.99)	(0.79)
CSIG (-1)			0.2158	1.082	1.0804	1.0794
			(2.33)	(30.10)	(30.03)	(29.86)
CSIG (-2)			-0.0915	-0.1548	-0.1546	-0.1480
			(-0.68)	(-2.96)	(-2.95)	(-2.82)
CSIG (-3)			-0.0478	-0.0130	-0.013	-0.0166
			(-0.35)	(-0.24)	(-0.24)	(-0.31)
CSIG (-4)			-0.1321	0.1122	0.1121	0.1086
			(0.98)	(2.13)	(2.13)	(2.06)
CSIG (-5)			0.0905	-0.4755	-0.4756	-0.4812
			(0.67)	(-9.09)	(-9.09)	(-9.15)
CSIG (-6)			-0.0522	0.3923	0.3908	0.3956
			(-0.56)	(10.96)	(10.91)	(10.95)
R^2	0.846	0.846	0.849	0.893	0.890	0.891
Σ AR(i)	0.882	0.880	0.878	0.943	0.940	.938
Half-Life (weeks)	5.5	5.4	5.3	11.9	11.2	10.8

Table 2: GARCH Models of Weekly Returns

Dep. Var.	(1) NG	(2) NG	(3) CRUDE	(4) CRUDE
Const.	0.0160	0.0150	-0.0577	-0.0498
	(1.12)	(1.11)	(-9.55)	(-5.75)
σ	-0.1005	-0.1085	0.3673	0.2978
	(-0.71)	(-0.90)	(5.21)	(3.28)
TBILL	-0.1303	-0.1255	0.7694	0.7204
	(-1.43)	(-1.52)	(12.24)	(8.01)
ENRON	-0.0622	-0.0577	-0.0270	-0.0211
	(-1.56)	(-1.63)	(-0.87)	(-0.45)
TIME	1.42E-05	1.67E-05	-1.53E-05	-6.66E-07
	(1.05)	(1.28)	(-1.39)	(0.04)
MA (1)		-0.1196		0.3432
		(-2.96)		(10.87)
VARIANCE EQUATION				
CONST.	0.0005	0.0004	9.31E-05	7.03E-05
	(4.78)	(3.93)	(0.57)	(0.68)
ARCH (1)	0.1237	0.0999	0.2434	0.0400
	(2.34)	(2.08)	(4.79)	(1.32)
ARCH (2)	-0.0644	-0.0596	0.1488	0.2072
	(-1.14)	(-1.13)	(3.58)	(3.15)
ARCH (3)	0.0292	0.0500	0.1446	-0.0429
	(0.62)	(0.99)	(3.96)	(-1.02)
ARCH (4)	0.2458	0.2638	0.2504	0.2838
	(1.72)	(1.48)	(5.60)	(5.06)
ARCH (5)	-0.2217	-0.2463		
	(-1.87)	(-1.63)		
GARCH (1)	0.8427	0.9359	0.1174	0.5585
	(9.08)	(8.79)	(1.40)	(3.13)
GARCH (2)	0.0510	-0.0127	-0.1699	-0.5046
	(0.32)	(0.07)	(-2.00)	(-2.52)
GARCH (3)	-0.1301	-0.1470	-0.3591	0.1020
	(-1.01)	(-0.98)	(-4.71)	(0.48)
GARCH (4)	0.2778	0.2629	0.5343	0.2690
	(3.65)	(2.16)	(7.44)	(2.09)
GARCH (5)	-0.2422	-0.2134		
	(-3.40)	(-2.39)		
ENRON	0.0028	0.0022	0.0027	0.0035
	(1.06)	(0.98)	(0.77)	(1.09)
TIME	7.56E-07	6.40E-07	2.15E-06	1.46E-06
	(4.98)	(4.20)	(2.58)	(4.37)
Half-Life (weeks)	7.5	10.1	7.3	7.6

Note: Regression equations for weekly returns include monthly dummy variables, which are not reported. Number of ARCH and GARCH terms chosen to minimize Akaike information criterion.

Table 3: GARCH Models of Daily Returns

Dep. Var.	(1) NG	(2) NG	(3) NG	(4) CRUDE	(5) CRUDE	(6) CRUDE
CONST	-0.0004	0.0031	-0.0004	-0.0012	-0.0013	-0.0014
	(-0.15)	(1.35)	(-0.13)	(-1.46)	(-1.52)	(-1.77)
σ	-0.0960	-0.0838	-0.0867	0.2196	0.2660	0.2266
	(-1.79)	(-1.70)	(-1.60)	(3.98)	(6.51)	(4.15)
TBILL	0.0254	0.0254	0.0237	0.0642	0.0570	0.0610
	(0.47)	(0.51)	(0.51)	(2.52)	(2.24)	(2.39)
ENRON		-0.0071	-0.0106		-0.0167	-0.166
		(-0.80)	(-1.19)		(-4.95)	(-5.05)
TIME	2.54E-06		2.26E-06	4.86E-07		7.04E-07
	(1.98)		(1.81)	(0.96)		(1.44)
VARIANCE EQUATION: GARCH (p,q)						
(p,q)	(8,7)	(4,8)	(5,8)	(5,9)	(4,9)	(4,9)
CONST	2.08E-05	4.91E-05	2.06E-05	1.78E-06	9.57E-06	2.57E-06
	(111.93)	(54.78)	(1136.09)	(0.53)	(3.63)	(0.80)
ENRON		0.0005	0.0007		0.0002	0.0002
		(1.78)	(1.57)		(0.97)	(0.91)
TIME	7.32E-08		3.37E-08	1.16E-08		1.08E-08
	(2.43)		(1.55)	(1.19)		(1.16)
Half-Life (weeks)	8.5	7.8	5.8	3.2	10.7	2.9

Note: Number of ARCH and GARCH terms chosen to minimize Akaike information criterion. ARCH and GARCH coefficients are not shown.

Table 4: Granger Causality Tests

Variable	Lags	NG → Crude	Crude → NG
Weekly Returns	2	No	No
(Simple Corr. = .095)	4	No	No
	6	No	No
Daily Returns	4	Yes*	Yes*
(Simple Corr. = .028)	6	No	Yes**
	10	No	Yes**
	14	No	Yes**
	18	No	Yes*
	22	No	No
Weekly Volatility,	2	No	Yes*
Sample Stand. Dev.	4	No	Yes*
(Simple Corr. = .170)	6	No	No
Weekly Volatility,	2	Yes**	No
GARCH	4	Yes**	No
(Simple Corr. = .092)	6	Yes*	No
Daily Volatility,	4	No	No
GARCH	6	No	Yes*
(Simple Corr. = .146)	10	No	No
	14	No	Yes**
	18	No	Yes**
	22	No	Yes**

Note: Test of x → y is an F-test of the exclusion restrictions $b_1 = b_2 = \ldots = b_L = 0$ in the regression $y_t = a_0 + \sum_{i=1}^{L} a_i y_{t-i} + \sum_{i=1}^{L} b_i x_{t-i}$. A "no" implies a failure to reject the hypothesis that the b_i's equal 0, and a "yes" implies rejection at the 5% (*) or 1% (**) level.

Figure 1
NATURAL GAS: WEEKLY SPOT PRICE AND VOLATILITY

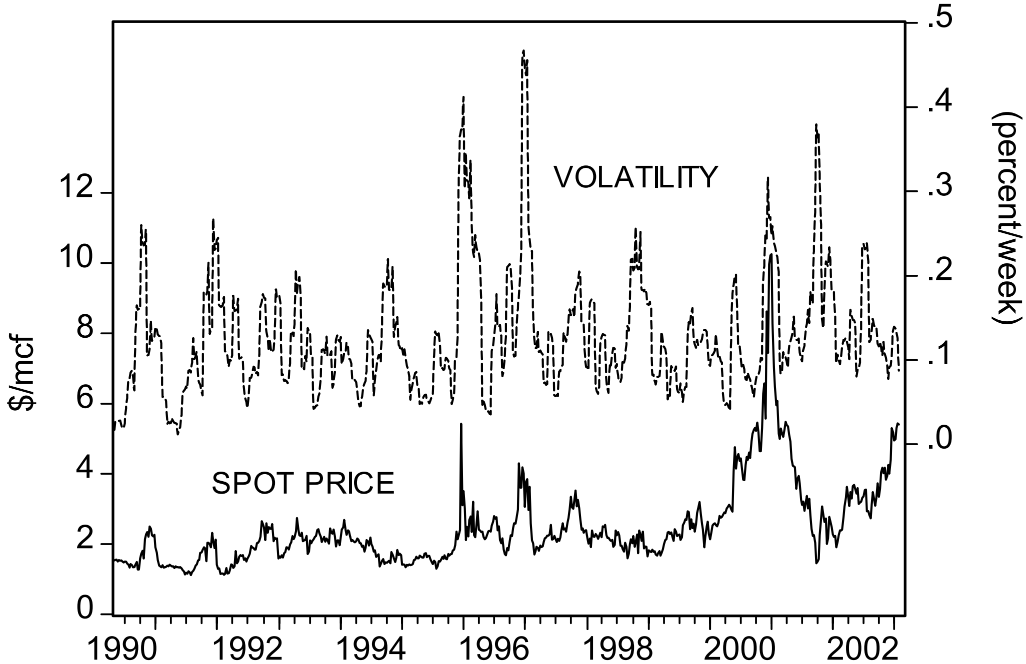

Figure 2
CRUDE OIL: WEEKLY SPOT PRICE AND VOLATILITY

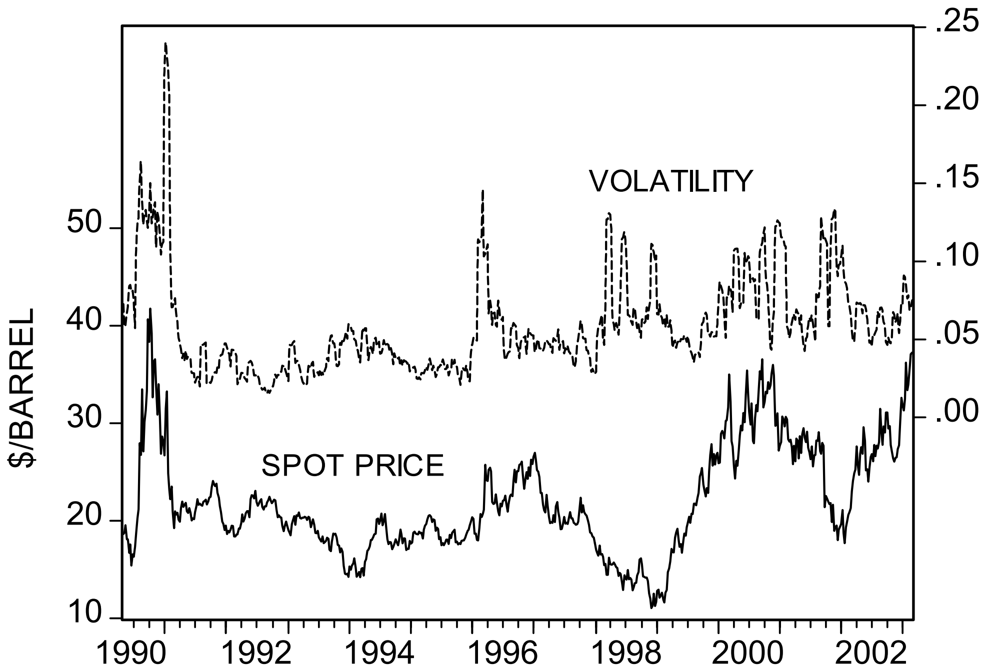

Figure 3
NATURAL GAS AND CRUDE OIL PRICE VOLATILITY

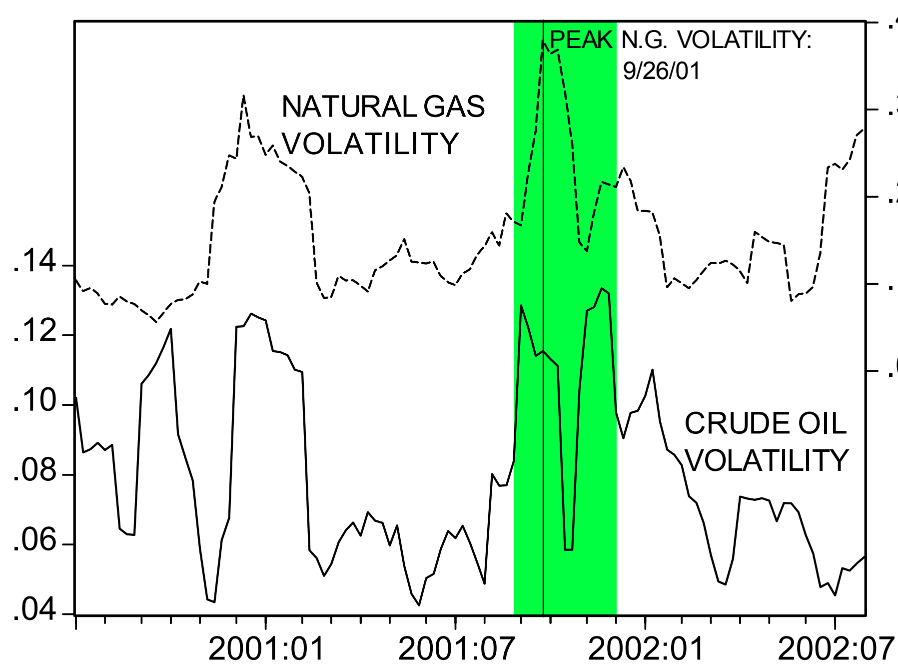

Figure 4
NATURAL GAS PRICE VOLATILITY