VOLATILITY IN NATURAL GAS AND OIL MARKETS

*Robert S. Pindyck**

Introduction

This paper examines the behavior of natural-gas and crude-oil price volatility in the United States 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 volatility in natural gas and crude oil markets is important for several reasons. Persistent changes in volatility can affect the risk exposure of

*Robert S. Pindyck, the Bank of Tokyo-Mitsubishi Professor of Economics and Finance in the Sloan School of Management at the Massachusetts Institute of Technology (M.I.T.), is also a Research Associate of the National Bureau of Economic Research and a Fellow of the Econometric Society. The author, who received a Ph.D. in economics from M.I.T., has been a visiting professor at both Tel Aviv and Harvard Universities, a coeditor of *The Review of Economics and Statistics*, and a consultant to a large number of public and private organizations. Dr. Pindyck's research and writing have covered topics in microeconomics and industrial organization, the behavior of resource and commodity markets, financial markets, and econometric modeling and forecasting. Along with numerous academic journal articles, he has authored or coauthored seven books, including the widely used texts *Econometric Models and Economic Forecasts* (McGraw-Hill), *Microeconomics* (Prentice-Hall), and *Investment under Uncertainty* (Princeton University Press). His recent work in economics and finance has examined the role of research and development and the value of patents, determinants of market structure and market power, the dynamics of commodity spot and futures markets, and criteria for investing in risky projects. The author is 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.

The Journal of Energy and Development, Vol. 30, No. 1

producers and industrial consumers of natural gas and oil; they can 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, the behavior of volatility is significant 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, volatility plays a role in the short-run market dynamics for natural gas and oil. As discussed in my recent paper, volatility can affect the demand for storage and also can impact the total marginal cost of production by affecting the value of firms' operating options and thereby the opportunity cost of current production.[1] In particular, greater volatility should lead to an increased demand for storage and an increase in both spot prices and marginal convenience yield.[2] Thus, changes in volatility can help explain changes in these other variables.

With this in mind, the following questions are addressed. First, has natural-gas and/or crude-oil price volatility changed significantly 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, are natural gas and crude oil volatilities 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. However, if 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 an earlier analysis,[3] the question is revisited of whether changes in volatility are predictable.

To address these questions, daily futures price data for natural gas and crude oil are used 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, daily and weekly returns from holding each commodity are calculated. Volatility then is estimated in three different ways.

First, using a five-week overlapping window, weekly series for price volatility are estimated by calculating sample standard deviations of (adjusted) daily log price changes. As J. Campbell et al. 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.[4] Second, series for conditional volatility are estimated by estimating generalized autoregressive conditional heteroscedasticity (GARCH) models of the weekly returns on the commodities; the volatility estimates from these models are compared to the sample standard deviations.[5] Third, a *daily* series for conditional volatility is estimated by using GARCH models of the daily returns on the commodities.[6]

The behavior of volatility is studied 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. These series also are used to estimate the persistence of changes in volatility. Second, these same questions are addressed using weekly and daily GARCH models of commodity returns. For example, I test whether a time trend or a dummy variable for the Enron period is a significant explanator of volatility (and/or returns) in the GARCH framework. Likewise, the estimated coefficients from the variance equation of each GARCH model provide a direct estimate of the persistence of volatility shocks.

While the focus here is on prices, there are other measures of volatility. Putting aside issues of data availability, one could examine instead the volatility of consumption, production, or inventories. That would indeed be useful if the objective was to explain the determinants of inventory demand, e.g., the role of production and/or consumption smoothing and production-cost smoothing.[7] My concern, however, is with the overall market, and the spot price is the best single statistic for market conditions. Spot price volatility reflects the volatility of current as well as expected future values of production, consumption, and inventory demand.[8]

My results can be summarized as follow. (1) There is a statistically significant positive trend in volatility for natural gas (but not for crude oil). However, this trend is of little economic importance; over a 10-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, the data and the calculation of returns and weekly sample standard deviations are discussed. All of the empirical work is presented in the subsequent section on the behavior of volatility and prices, which is followed by the conclusions.

The Data

To begin, natural-gas and crude-oil futures price data are utilized 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, the sample standard deviations of adjusted daily log price changes in spot and futures prices are used. As discussed below, I also obtain estimates of conditional volatility from GARCH models of weekly and daily returns.

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, 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 *cash prices*, purportedly reflecting actual transactions. But daily cash price data usually are not available, and a cash price can include discounts and premiums that result from relationships between buyers and sellers; it need not reflect precisely the same product (including delivery location) specified in the futures contract. A second approach is to use the price of the spot futures contract, i.e., the contract that expires in month t. But the spot contract often expires before the end of the month, and active spot contracts do not always exist for each month.

The third approach, which is used 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 backward to the spot month as follows:

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

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, respectively; 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, respectively.

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 followed a geometric Brownian motion, 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 weekend), 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.[9] To deal with this, the daily price data are sorted 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, and Monday would be assigned an interval of three days because of the weekend. Because of holidays, some prices are assigned to intervals of two, four, or even five days (if a weekend was followed by a two-day holiday).

For each interval set, the sample standard deviation of log price changes is calculated for the entire sample for each commodity. Letting $\hat{S}_n$ denote this sample standard deviation for log price changes over an interval of n days, the "effective" daily log price change is computed as:

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

For each week, I then compute a sample variance and corresponding sample standard deviation using these "effective" 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. Equation (3) gives the sample standard deviation of daily percentage price changes; to put it in weekly terms, one multiplies by $\sqrt{30/4} = \sqrt{7.5}$. The resulting weekly series is a measure of volatility, σ_t.

Daily and Weekly Returns: An important advantage of using the weekly estimates of volatility discussed above (besides its simplicity) is that it does not require a parametric model of the evolution of volatility over time. However, there are disadvantages as well. The first is that the use of overlapping intervals introduces serial correlation as an artifact, which makes it 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 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 volatility shocks) are estimated along with the volatility itself.

Marginal Convenience Yield: One part of the total return on the commodity is the net marginal convenience yield, ψ_t, i.e., the value of the flow of production- and delivery-facilitating services from the marginal unit of inventory, net of storage costs. Net marginal convenience yield can be measured from spot and futures prices as

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

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. The values of ψ_t are calculated for every trading day using the futures price corresponding as closely as possible to a one-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.) Further, I employed the yield on three-month Treasury bills, adjusted for the number of days between P_t and F_{1t}, for the interest rate r_t.

In what follows, both daily and weekly series are used for the marginal convenience yield, so ψ_t is converted into daily terms, i.e., dollars per unit of commodity per day. For days followed by another trading day (e.g., a Monday), one simply divides 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, these values are multiplied 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 three days.) 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 seven so that the convenience yield is measured in dollars per unit of commodity per week. (If Wednesday is a holiday, Thursday's price is used.) This weekly series then is utilized 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. 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:

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

where δ_τ is given by equation (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. (Because effective trading days are utilized, the daily series will have about 20 data points per month.)

The Behavior of Volatility and Prices

To examine the behavior of natural-gas and crude-oil price volatility, I first use the weekly time series of sample standard deviations of adjusted log price changes.

These time series show little evidence of either a trend in volatility or a significant increase in volatility 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, GARCH models of the weekly returns to holding the commodity are estimated, which yields estimates of the weekly conditional standard deviations. 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, and obtain results that are similar to those obtained from the weekly sample standard deviations. The daily adjusted return series is used to estimate daily GARCH models. These provide estimates of conditional standard deviations on a daily basis that also are used to test for time trends and an Enron effect and to estimate the persistence of changes in volatility.

Weekly Sample Standard Deviations: Figures 1 and 2 show the weekly series for the spot price and volatility, 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; 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 for natural gas and 0.23 and 2.84 for crude oil, which are roughly consistent with a normal distribution. However, a Jarque-Bera test rejects normality at the 1-percent level in both cases.

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 also were periods of high volatility that accompanied unusually low spot prices for both commodities, e.g., during 1998. Overall, volatility and price are moderately correlated; the correlation (in levels) is .27 for natural gas and .37 for crude oil.

Was there a significant increase in volatility during the period of the Enron collapse? The Enron bankruptcy sharply reduced spot and forward trading in natural gas and electricity and led as well to speculation over net long and short positions in natural gas. This probably 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

Figure 1

NATURAL GAS: WEEKLY SPOT PRICE AND VOLATILITY, 1990-2002

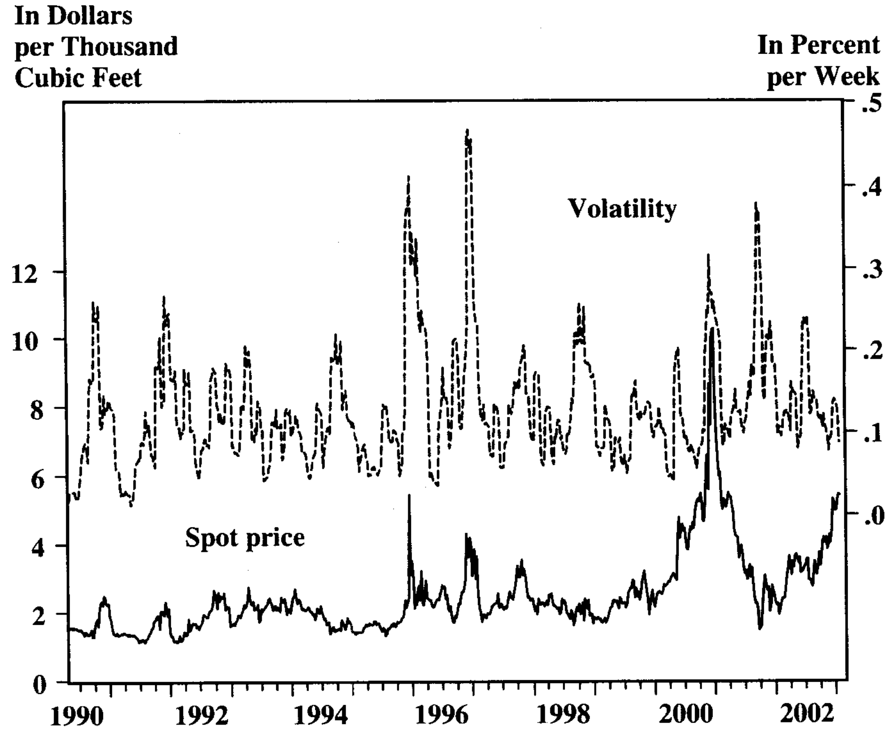

Figure 2

CRUDE OIL: WEEKLY SPOT PRICE AND VOLATILITY, 1990-2002

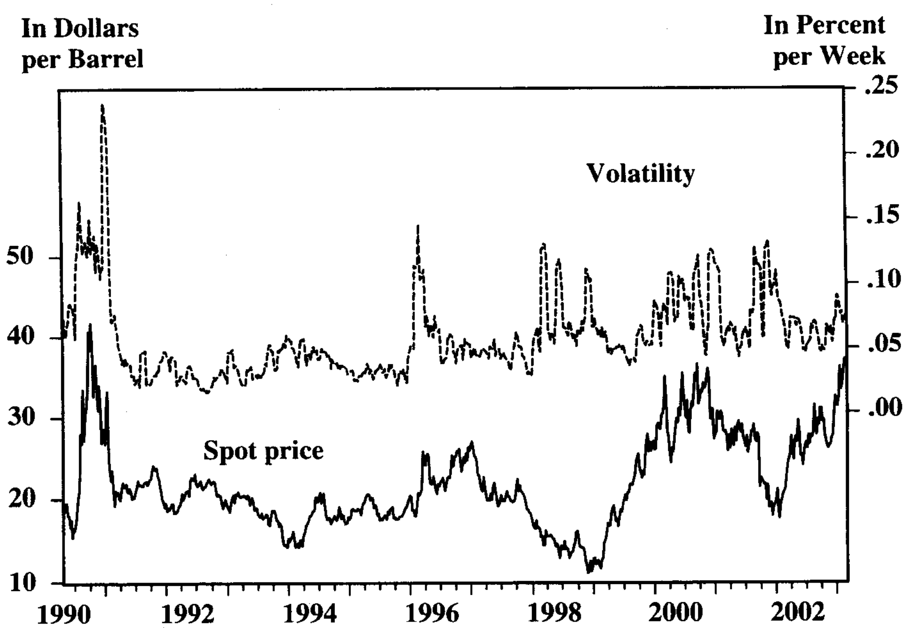

16, 2001, Enron reported a $638-million third-quarter loss and disclosed a $1.2-billion reduction in shareholder equity. Further financial statement revisions were announced during October and November; Enron filed for Chapter 11 bankruptcy protection on December 2.

I defined the period of the Enron collapse as August 29 to December 5, 2001, and created a dummy variable equal to 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; crude oil volatility also was above average. The significance of these increases in volatility is examined in the context of forecasting regressions.

Using data for crude oil, heating oil, and gasoline, I have shown elsewhere 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).[10] 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 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).[11]

Table 1 shows simple forecasting regressions for volatility. In columns (1) and (4), the explanatory variables are six lags of volatility and the Enron dummy variable. For natural gas, the Enron dummy is marginally significant; for crude oil it is insignificant. But even for natural gas 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-statistic for column (6) is 1.39.[12]

The bottom of table 1 shows the sum of the autoregressive coefficients along with the implied half-life for volatility shocks. The half-life is about five to six weeks for natural gas and 11 to 12 weeks for crude oil. Thus, although volatility itself fluctuates considerably, shocks to volatility appear 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, so that the series

Figure 3

NATURAL-GAS AND CRUDE-OIL PRICE VOLATILITY, JULY 2000-JULY 2002[a]

(in percent per week)

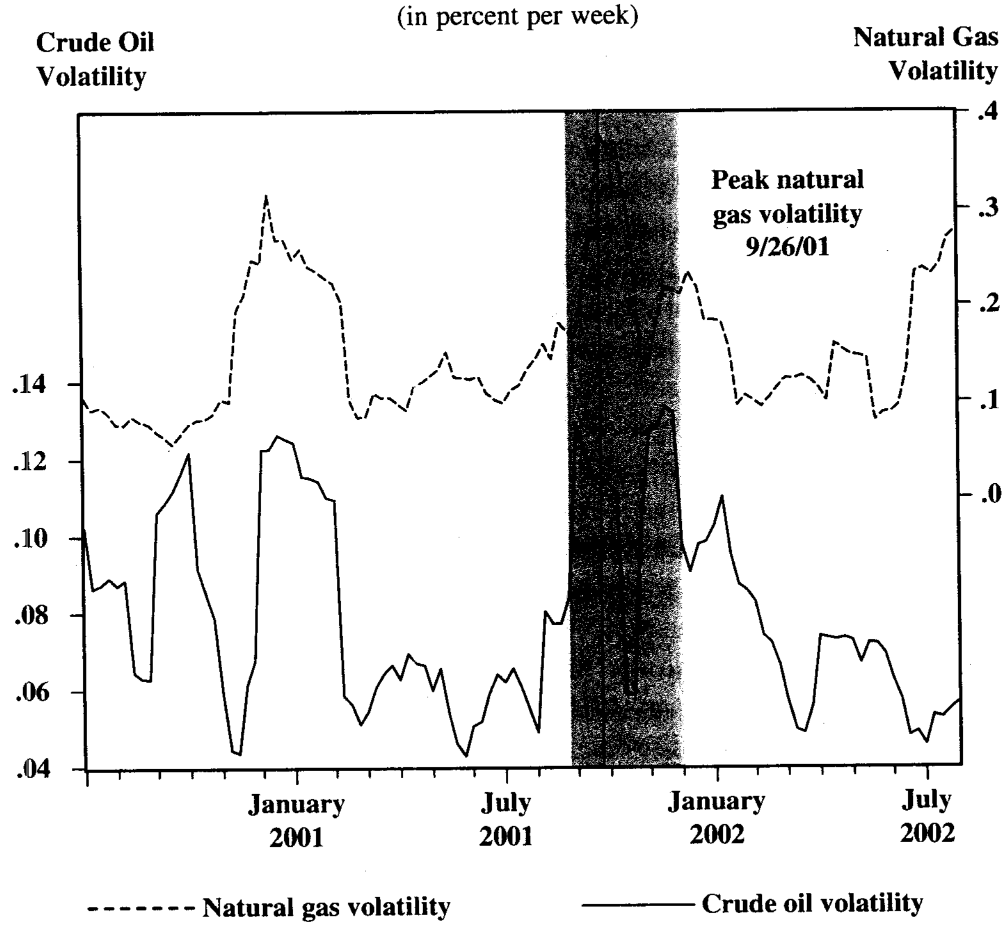

[a]Shaded area is period of Enron collapse.

are serially correlated by construction. Second, even with five-week intervals, each sample standard deviation is based on at most 25 observations. One way to get around these problems is to estimate GARCH models of the commodity returns themselves.

GARCH Models of Weekly Returns: Models are estimated of the following form. The weekly return to holding the commodity is:

$$RET_t = a_0 + a_1 TBILL_t + a_2 \sigma_t + a_3 ENRON_t + a_4 TIME_t + \sum_{j=1}^{11} b_j DUM_{jt} + \varepsilon_t \,, \qquad (6)$$

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

Table 1
FORECASTING EQUATIONS FOR NATURAL GAS (NG) AND CRUDE OIL VOLATILITY

Dependent Variable	(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

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:

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

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 equation (7) is chosen to minimize the Akaike information criteria.

The results for crude oil [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 (i.e., the standard deviation of ε_t). For natural gas, however, both the interest rate and volatility are statistically insignificant in the returns equation. For both commodities, the time trend is insignificant in the returns equation but is positive and significant in the variance equation, and the Enron dummy is positive but statistically insignificant in the variance equation. Thus, I find a 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, the time trend coefficient is about $7\mathrm{x}10^{-7}$, 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; 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.,

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

The half-life of volatility shocks is about seven to 10 weeks for natural gas, and seven to eight weeks for crude oil. These numbers differ slightly from the estimates in table 1, but overall, shocks to volatility again appear transitory for both commodities.

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.

Table 2
GENERALIZED AUTOREGRESSIVE CONDITIONAL HETEROSCEDASTICITY (GARCH) MODELS OF NATURAL-GAS (NG) AND CRUDE-OIL WEEKLY RETURNS[a]

Dependent Variable	(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
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)

(continued)

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 sample standard deviations. However, these weekly models do not make use of all of the available daily data. Thus, GARCH models of daily returns also are estimated. These models take the form of equations (6) and (7); monthly dummy variables in the returns equation are not included. The number of lags is again chosen to minimize the Akaike information criterion.

As with the weekly GARCH models, the results for crude oil, but not natural gas, are consistent with the theory of commodity returns and storage (see table 3). Crude oil returns have a strong positive dependence on the interest rate and on volatility, but both variables are 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

Table 2 (continued)
GENERALIZED AUTOREGRESSIVE CONDITIONAL HETEROSCEDASTICITY (GARCH) MODELS OF NATURAL-GAS (NG) AND CRUDE-OIL WEEKLY RETURNS[a]

Dependent Variable	(1) NG	(2) NG	(3) CRUDE	(4) CRUDE
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

[a]Regression equations for weekly returns include monthly dummy variables, which are not reported. Numbers of ARCH and GARCH terms were chosen to minimize Akaike information criterion.

now only marginally significant for natural gas and insignificant for crude oil, but even for natural gas it is only of marginal economic importance. (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 close to those in tables 1 and 2. The half-life is about six to nine weeks for natural gas, and 3 to 11 weeks for crude oil. Once again, shocks to volatility appear to be largely transitory.

Returns and Volatilities across Markets: Turning to the interrelationship between crude oil and natural gas returns and volatilities, the results in table 1, based on the five-week sample standard deviations, provide some evidence that crude oil volatility has predictive power with respect to natural gas volatility (but not the other way around). To explore this further, I run Granger causality tests between

Figure 4
NATURAL-GAS PRICE VOLATILITY, 1990-2002
(in percent per week)

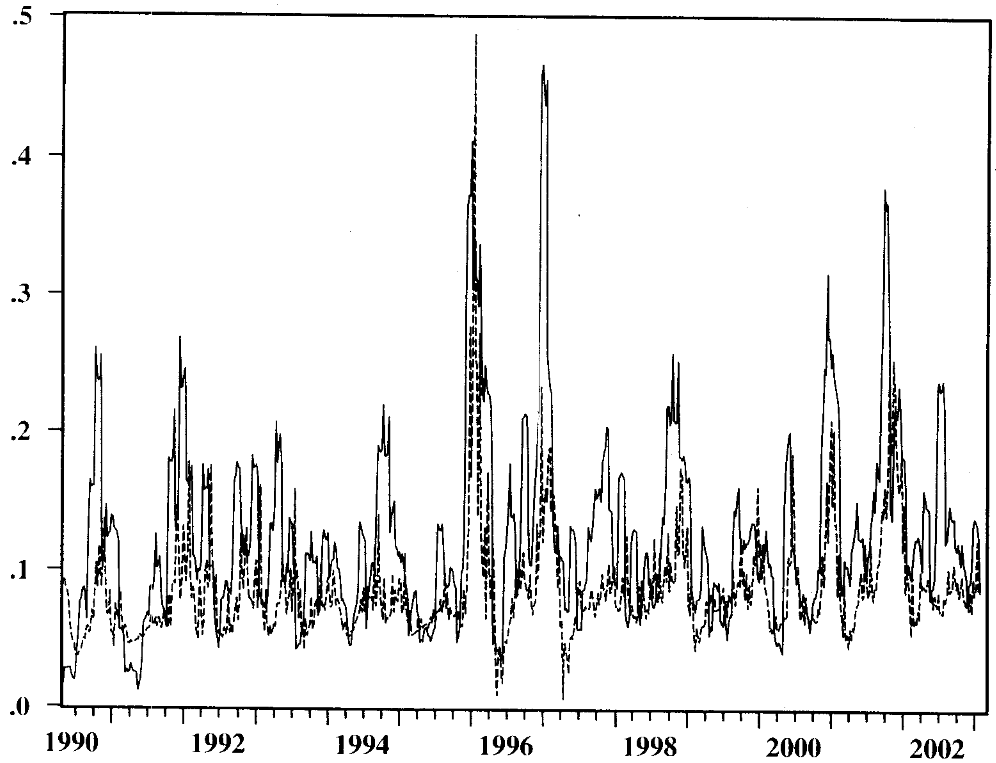

—— **Sample standard deviation** ------ **Generalized autoregressive conditional heteroscedasticity (GARCH) estimates**

gas and oil using the sample standard deviations and the weekly and daily volatilities from the GARCH models. I also apply these tests on weekly and daily gas and oil returns. These tests are simply F-tests of the exclusion restrictions $b_1 = b_2 = \ldots = b_L = 0$ 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 two, four, and six 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

Table 3

GENERALIZED AUTOREGRESSIVE CONDITIONAL HETEROSCEDASTICITY (GARCH) MODELS OF NATURAL-GAS (NG) AND CRUDE -OIL DAILY RETURNS[a]

Dependent Variable	(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

[a]Numbers of ARCH and GARCH terms were chosen to minimize Akaike information criterion. ARCH and GARCH coefficients are not shown.

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—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),

Table 4
GRANGER CAUSALITY TESTS: NATURAL GAS (NG) AND CRUDE OIL[a]

Variable	Lags	NG → Crude	Crude → NG
Weekly returns	2	No	No
(Simple correlation = .095)	4	No	No
	6	No	No
Daily returns	4	Yes*	Yes*
(Simple correlation = .028)	6	No	Yes**
	10	No	Yes**
	14	No	Yes**
	18	No	Yes*
	22	No	No
Weekly volatility,	2	No	Yes*
Sample standard deviation	4	No	Yes*
(Simple correlation = .170)	6	No	No
Weekly volatility,	2	Yes**	No
GARCH[b]	4	Yes**	No
(Simple correlation = .092)	6	Yes*	No
Daily volatility,	4	No	No
GARCH[b]	6	No	Yes*
(Simple correlation = .146)	10	No	No
	14	No	Yes**
	18	No	Yes**
	22	No	Yes**

[a]Test of $x \longrightarrow y$ is an F-test of the exclusion restrictions $b_1 = b_2 = \ldots = b_L = 0$ in the regression

$$y_t = a_0 + \sum_{i=1}^{L} a_i y_{t-1} + \sum_{i=1}^{L} b_i x_{t-i} .$$

A "no" implies a failure to reject the hypothesis that the b_i's equal 0, and a "yes" implies rejection at the 5-percent (*) or 1-percent (**) level.

[b]Generalized autoregressive conditional heteroscedasticity.

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.

Summary and Conclusions

My results can be summarized as follow. First, there is a statistically significant positive time trend in volatility for natural gas and, to a lesser extent, for oil. The trends, however, are small, and not of great economic importance. Given the limited length of my sample, there are certainly no conclusions that can be drawn about long-term trends. As for the demise of Enron, it does not appear to have contributed to any significant increase in volatility.

Second, there is some evidence that crude oil volatility and returns have predictive power for natural gas volatility and returns, but not the other way around. Nonetheless, 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 certainly could 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 on 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 J. Hull and A. White, among others, suggests that treating volatility as non-stochastic will make little quantitative difference for such valuations.[13]

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.

NOTES

[1]Robert S. Pindyck, "Volatility and Commodity Price Dynamics," *The Journal of Futures Markets*, forthcoming 2004.

[2]Using weekly data for the petroleum complex, Robert S. Pindyck, "Volatility and Commodity Price Dynamics," shows that the theoretical relationships between volatility and other variables are well supported for heating oil, but less so for crude and gasoline. The role of volatility in the opportunity cost of production also is spelled out and tested by Robert H. Litzenberger and Nir Rabinowitz, "Backwardation in Oil Futures Markets: Theory and Empirical Evidence," *Journal of Finance*, December 1995, pp. 1517-545. For an introduction to the interrelationships among price, inventories, and convenience yields, see Robert S. Pindyck, "The Dynamics of Commodity Spot and Futures Markets: A Primer," *The Energy Journal*, vol. 22, no. 3 (2001), pp. 1-29.

[3]Robert S. Pindyck, "Volatility and Commodity Price Dynamics."

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

[5]For an introduction to generalized autoregressive conditional heteroscedasticity (GARCH) models and their use, see R. Pindyck and D. Rubinfeld, *Econometric Models and Economic Forecasts*, 4th ed. (Columbus, Ohio: McGraw-Hill, 1998), chapter 10.

[6]See Eduardo S. Schwartz, "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging," *The Journal of Finance*, July 1997, pp. 923-73 and Eduardo S. Schwartz and James E. Smith, "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science*, July 2000, where futures and spot prices are used to estimate a mean-reverting price process and value commodity-based options, an approach that also yields implicit time-varying estimates of volatility. M. Haigh and M. Holt estimate GARCH models to study volatility spillovers across the components of the petroleum complex (crude oil, heating oil, and gasoline) in Michael S. Haigh and Matthew T. Holt, "Crack Spread Hedging: Accounting for Time-Varying Volatility Spillovers in the Energy Futures Markets," *Journal of Applied Econometrics*, May-June 2002, pp. 269-89.

[7]These issues also are addressed in Robert S. Pindyck, "Inventories and the Short-Run Dynamics of Commodity Prices," *The RAND Journal of Economics*, spring 1994, pp. 141-59 and Zvi Eckstein 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).

[8]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.

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

[10]Robert S. Pindyck, "Volatility and Commodity Price Dynamics."

[11]See John H. Herbert, "Trading Volume, Maturity and Natural Gas Futures Price Volatility," *Energy Economics*, October 1995, pp. 293-99, which shows that natural gas futures price volatility can be explained partly by the volume of trading in the futures contract.

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

[13]John Hull and Alan White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, June 1987, pp. 281-300.