
VOLATILITY AND COMMODITY PRICE DYNAMICS

ROBERT S. PINDYCK

Commodity prices are volatile, and volatility itself varies over time. Changes in volatility can affect market variables by directly affecting the marginal value of storage, and by affecting a component of the total marginal cost of production, the opportunity cost of producing the commodity now rather than waiting for more price information. I examine the role of volatility in short-run commodity market dynamics and the determinants of volatility itself. I develop a structural model of inventories, spot, and futures prices that explicitly accounts for volatility, and estimate it using daily and weekly data for the petroleum complex: crude oil, heating oil, and gasoline.
© 2004 Wiley Periodicals, Inc. Jrl Fut Mark 24:1029–1047, 2004

INTRODUCTION

Most commodity markets are volatile, and volatility itself fluctuates over time. This paper examines the short-run dynamics of commodity prices and inventories, focusing on the behavior and role of volatility. I show

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 Denny Ellerman, Richard Farmer, Kenneth French, Christopher Gilbert, Philip Verleger, Jr., and an anonymous referee for helpful comments and suggestions, and to Nicola Stafford for her outstanding research assistance.

For correspondence, Robert S. Pindyck, Massachusetts Institute of Technology, Sloan School of Management, 50 Memorial Drive, Cambridge, MA 02142; e-mail: rpindyck@mit.edu

Received October 2003; Accepted January 2004

-
- *Robert S. Pindyck is a Bank of Tokyo–Mitsubishi Professor of Economics and Finance at the Sloan School of Management at Massachusetts Institute of Technology in Cambridge, Massachusetts.*

how changes in volatility affect spot prices, futures prices, and inventories, and measure the magnitudes of these effects.

Volatility affects prices, production, and inventories in two principal ways. First, it directly affects the marginal value of storage (the *marginal convenience yield*), i.e., the flow of benefits from an extra unit of inventory. When prices—and hence production and demand—are more volatile, there is a greater demand for inventories, which are needed to smooth production and deliveries and reduce marketing costs. Thus an increase in volatility can lead to inventory build-ups and raise prices in the short run.

Second, for a depletable resource like oil, volatility affects the total marginal cost of production via the “option premium.” Producers hold operating options, with an exercise price equal to direct marginal production cost and a payoff equal to the spot price. *Total* marginal cost equals the direct marginal cost plus the opportunity cost of exercising the incremental operating option. An increase in price volatility raises the value of this option and the associated opportunity cost, and can thus result in a decrease in production. Litzenger and Rabinowitz (1995) used a two-period model to show that this option premium can cause backwardation in futures markets. Using data for crude oil, they demonstrated that consistent with the theory, production is negatively correlated, and the extent of backwardation is positively correlated, with price volatility.¹ I show how volatility and option value can be incorporated in a model of a commodity market dynamics.

I develop a weekly model that relates the dynamics of inventories, spot and futures prices, and the level of volatility, and I estimate the model using data for the three commodities that make up the petroleum complex: crude oil, heating oil, and gasoline. To estimate volatility, I use sample standard deviations of adjusted daily log changes in spot and futures prices. As Campbell et al. (2001) point out, in addition to its simplicity, this approach has the advantage that it does not require a parametric model of the evolution of volatility.

As shown below, at least for the petroleum complex, changes in price volatility are not predicted by market variables such as inventories or convenience yields, and can be viewed as exogenous. However, changes in volatility affect market variables through the marginal value of storage and by affecting price and production through the option premium,

¹Schwartz (1997) and Schwartz and Smith (2000) show how futures prices can be used to estimate a mean-reverting price process and derive values of commodity-based options. Casassus and Collin-Dufresne (2001) estimate a three-factor model of commodity prices that extends the models of Schwartz (1997) and others by allowing for time-varying risk premia. Milonas and Henker (2001) also study the behavior of crude oil convenience yield and show that it is like a call option.

although the impact is smaller than that suggested by Litzenberger and Rabinowitz (1995), Milonas and Henker (2001), and others. Also, changes in the value of storage affect production, inventories, and spot prices, so these variables are indirectly affected by volatility.²

This paper also shows how inventories adjust and affect prices in the short run. Inventories can be used to reduce costs of varying production (when marginal cost is increasing), and to reduce marketing costs by facilitating production and delivery scheduling and avoiding stockouts. Equilibrium inventory behavior is the solution to a stochastic dynamic optimization problem. Early studies of manufacturing inventories, as well as Eckstein and Eichenbaum's (1985) study of crude oil inventories, rely on linear-quadratic specifications to obtain analytical solutions. This is unrealistic for commodity markets because the cost of drawing down inventory is highly convex in the stock of inventory, rising rapidly as the stock falls toward zero, but remaining very small for moderate to high stock levels. I therefore adopt a more general specification and estimate the Euler equations that follow from intertemporal optimization, using futures market data to directly measure the marginal value of storage.³ Unlike my earlier study of commodity markets (Pindyck, 1994), here I explicitly account for price volatility as a determinant of both the marginal value of storage and the full marginal cost of production. I thereby estimate the extent to which changes in volatility affect prices and inventories, and obtain evidence on the channels through which these effects occur.

A MODEL OF PRICES, INVENTORIES, AND VOLATILITY

Costs

The total economic cost of commodity production, marketing, and storage is:

$$TC = C(x) + \Omega(x; \sigma, r) + \Phi(N, P, \sigma) + kN \quad (1)$$

and has four components: (1) $C(x)$ is the direct cost of producing output x . (2) $\Omega(x; \sigma, r)$ is the opportunity cost of producing x now, rather

²The exogeneity of volatility is consistent with informational efficiency in the spot and futures markets.

³See Pindyck (1993, 1994). This approach has also been used in studies of manufacturing inventories, e.g., Miron and Zeldes (1988) and Ramey (1991). Considine (1997) and Considine and Heo (2000) estimated Euler equation models of inventory behavior for petroleum products, focusing on the joint production characteristics of refining. McDonald and Shimko (1998) estimate the marginal value of storage for gold.

than waiting to see how prices evolve; it is the cost of exercising firms' "operating options," and it depends on the level of price volatility, σ and the risk-free rate r . (3) $\Phi(N, P, \sigma)$ is total marketing cost, i.e., the cost of production and delivery scheduling and stockout avoidance, and is decreasing in the level of inventories N . (4) k is the per-unit storage cost, which I assume is constant.

The first and last components of cost (the direct production cost and the cost of storage) are standard. The second component, $\Omega(x; \sigma, r)$, is the total opportunity cost of producing output x . The opportunity cost of producing the marginal unit is $\omega = \partial \Omega / \partial x$; price will exceed direct marginal cost by this premium. The third component, $\Phi(N, P, \sigma)$, is the total marketing cost, which includes actual or opportunity costs of activities facilitated by inventories, e.g., costs of adjusting production over time, delivery scheduling, and stockout avoidance.⁴ The value of the marginal unit of inventory (the marginal convenience yield) is $\psi = -\partial \Phi / \partial N$. Note that the net (of storage costs) marginal convenience yield can be measured from spot and futures prices:

$$\psi_t - k = (1 + r)P_t - F_{1t} \quad (2)$$

where F_{1t} is the futures price at time t for a contract maturing at time $t + 1$, r is the one-period interest rate, and k is the one-period cost of storage.⁵ Thus ψ_t is the value of the flow of production- and delivery-facilitating services from the marginal unit of inventory, a value that should be greater the greater is the volatility of price.

Euler Equations

Taking prices as given, firms choose production and inventory levels to maximize the present value of the expected flow of profits:

$$\max \mathcal{E}_t \sum_{\tau=0}^{\infty} R_{\tau,t} (P_{t+\tau} Q_{t+\tau} - TC_{t+\tau}) \quad (3)$$

⁴If marginal production costs are increasing with the rate of output and if demand is fluctuating, producers can reduce costs by selling from inventory during high-demand periods, and replenishing inventories during low-demand periods. Industrial consumers also hold inventories to facilitate their production processes.

⁵To see why Equation (2) must hold, note that the (stochastic) return from holding a unit of the commodity for one period is $\psi_t + (P_{t+1} - P_t) - k$. Suppose that one also shorts a futures contract. The return on this futures contract is $F_{1t} - F_{1,t+1} = F_{1t} - P_{t+1}$, so one would receive a total return equal to $\psi_t + (F_{1t} - P_t) - k$. No outlay is required for the futures contract, and this total return is non-stochastic, so it must equal the risk-free rate times the cash outlay for the commodity, i.e., rP_t , from which Equation (2) follows. Because futures contracts are marked to market, strictly speaking, F_{1t} should be a forward price. For most commodities, however, the difference between the futures and forward prices is negligible.

where $R_{\tau,t}$ is the τ -period discount factor, Q is sales, TC is given by Equation (1), and the maximization is subject to the accounting identity

$$\Delta N_t = x_t - Q_t \quad (4)$$

(The maximization is subject to the additional constraint that $N_{t+\tau} \geq 0$ for all τ , but because $\Phi \rightarrow \infty$ as $N \rightarrow 0$, this constraint will never be binding.) To obtain first-order conditions, first maximize with respect to x_t , holding N_t fixed so that $\Delta x_t = \Delta Q_t$:

$$P_t = \frac{\partial \text{TC}_t}{\partial x_t} \quad (5)$$

Equation (5) simply equates price with full marginal cost, where the latter includes the opportunity cost of exercising the marginal operating option.

Next, maximize Equation (3) with respect to N_t , holding Q_t and N_{t+1} fixed:

$$\mathcal{E}_t \left(\frac{R_{t+1}}{R_t} P_{t+1} \right) = P_t + \frac{\partial \text{TC}_t}{\partial N_t} \quad (6)$$

Over a one-week time period, $R_{t+1}/R_t \approx 1$, so that Equation (6) can be rewritten as:

$$\mathcal{E}_t P_{t+1} = P_t + \frac{\partial \text{TC}_t}{\partial N_t} \quad (7)$$

Finally, to make this easier to interpret, substitute Equation (5) for P_t and rearrange:

$$\frac{\partial \text{TC}_t}{\partial N_t} = \mathcal{E}_t \left(\frac{\partial \text{TC}_{t+1}}{\partial x_{t+1}} \right) - \frac{\partial \text{TC}_t}{\partial x_t} \quad (8)$$

Equation (8) describes the trade-off between selling out of inventory versus producing. Consider producing an extra unit now, holding it in inventory for one period, and then next period selling it and producing one unit less. The left-hand side of the equation is the savings in marketing costs from the extra unit of inventory, net of storage costs (i.e., net marginal convenience yield). This should equal the expected change in production cost (the increase from producing an extra unit now, minus the expected decrease next period from producing one unit less). Expected changes in production cost can come from expected changes in input prices and opportunity costs, and changes due to convexity of the cost function.

Empirical Specification

The components of cost are modelled as follows. I assume that the direct cost of production is quadratic. For crude oil, direct cost is:

$$C(x) = (c'_0 + \eta_t)x_t + \frac{1}{2}c_1x_t^2 \quad (9)$$

where η_t is a random shock. Note that there are no input cost variables (such as wage rates) in Equation (9); such variables cannot be measured—and are unlikely to vary much—on a weekly basis. For heating oil and gasoline, however, the price of crude oil ($P_{C,t}$) is a large component of direct production cost, and must be accounted for:

$$C(x) = (c'_0 + \eta_t)x_t + \frac{1}{2}c_1x_t^2 + c_2P_{C,t}x_t \quad (10)$$

Marketing cost should be roughly proportional to price, and should be increasing in the level of price volatility, which is a proxy for market volatility in general. Ideally, the marketing cost function Φ should be derived from a dynamic optimizing model that accounts for stockout costs and costs of scheduling and managing production and shipments, but that is beyond the scope of this work. Instead, I assume that this function is isoelastic in price, the variance of log price changes, and the total inventory level:

$$\Phi(N, P, \sigma) = \frac{1}{\alpha_3 - 1} \exp\left(b_0 + \sum_{j=1}^{11} b_j \text{DUM}_{jt}\right) P_t^{\alpha_1} (\sigma_t^2)^{\alpha_2} N_t^{1-\alpha_3} \quad (11)$$

where DUM_{jt} are monthly time dummies and $\alpha_3 > 1$. This implies that the marginal value of storage (marginal convenience yield), $\psi_t = -\partial\Phi/\partial N$, can be written as:

$$\log \psi_t = b_0 + \sum_{j=1}^{11} b_j \text{DUM}_{jt} + \alpha_1 \log P_t + \alpha_2 \log \sigma_t^2 - \alpha_3 \log N_t \quad (12)$$

To model the marginal opportunity cost $\omega_t = \partial\Omega_t/\partial x_t$, we need the value of the option to produce a marginal unit of the commodity, and the optimal price P^* at which that option should be exercised. The difference between P^* and the direct marginal cost $C'(x)$ is the opportunity cost of exercising the option to produce the marginal unit. Valuing this option requires assumptions about the stochastic dynamics of price. Because commodity prices tend to be mean-reverting, I assume the following continuous-time reduced form price process:

$$dP/P = \lambda(\mu - P) dt + \sigma dz \quad (13)$$

Here, μ is the “normal” price to which P_t tends to revert and λ is the speed of reversion. I treat σ as a constant because allowing for stochastic volatility precludes a closed-form solution for the option value. Furthermore, it should not affect the way in which the option value depends on volatility, although it will affect its magnitude (overstating it). To account for this, I include a scaling coefficient that is estimated as part of the model.⁶

If the price process follows Equation (13) and direct marginal cost is non-stochastic, a series solution can be found for the value of the option to produce. I use a quadratic approximation to this solution. As shown in the Appendix, letting r denote the risk-free rate and ρ the commodity’s risk-adjusted expected return, the opportunity cost ω_t can be written as:

$$\omega_t = \frac{1}{\sqrt{\gamma_1 \gamma_2}} - \mu \quad (14)$$

where

$$\gamma_1 = \frac{\lambda \theta}{\lambda \mu + \theta \sigma^2}, \quad \gamma_2 = \frac{\lambda(\theta + 1)}{2\lambda \mu + (2\theta + 1)\sigma^2} \quad (15)$$

and

$$\theta = \frac{1}{2} + \frac{(\rho - r - \lambda \mu)}{\sigma^2} + \sqrt{\left[\frac{(r - \rho + \lambda \mu)}{\sigma^2} - \frac{1}{2} \right]^2 + \frac{2r}{\sigma^2}} \quad (16)$$

I include a scaling coefficient, so that $c_3 \omega_t$ is the marginal opportunity cost. Note that the estimated value of c_3 should be close to 1.

Estimating Equations

Substitute these functions for the components of cost into Equations (5) and (8). Including the crude oil input price, $c_2 P_{C,t}$, Equation (5) becomes:

$$P_t = c'_0 + c_1 x_t + c_2 P_{C,t} + c_3 \omega_t + \eta_t \quad (17)$$

(This equation applies to heating oil and gasoline; for crude oil the $P_{C,t}$ term is dropped.)

It is convenient to eliminate production and write the model in terms of prices and inventories. In the short run consumption should be

⁶The numerical analyses of Hull and White (1987) suggest that treating σ as non-stochastic makes little quantitative difference. See, also, Franks and Schwartz (1991) and Ball and Torous (1991).

very price inelastic, so I model it as:

$$Q_t = \bar{Q} + \sum_{j=1}^{11} d_j \text{DUM}_{jt} + c_4 \text{HDD}_t + c_5 \text{CDD}_t + c_6 T_t + \epsilon_t \quad (18)$$

where the DUM_{jt} are monthly dummies, HDD and CDD are, respectively, heating and cooling degree days, and T is a time trend. Thus I assume that consumption fluctuates seasonally and in response to changes in temperature, is subject to (possibly serially correlated) random shocks (ϵ_t), but is insensitive to price. Substituting for Q_t in Equation (4) and rearranging:

$$x_t = \Delta N_t + \bar{Q} + \sum_{j=1}^{11} d_j \text{DUM}_{jt} + c_4 \text{HDD}_t + c_5 \text{CDD}_t + c_6 T_t + \epsilon_t \quad (19)$$

Thus Equation (17) can be rewritten as:

$$P_t = c_0 + c_1 \Delta N_t + c_2 P_{C,t} + c_3 \omega_t + c'_4 \text{HDD}_t + c'_5 \text{CDD}_t + c'_6 T_t + \sum_{j=1}^{11} d'_j \text{DUM}_{jt} + c_1 \epsilon_t + \eta_t \quad (20)$$

where $c_0 = c'_0 + c_1 \bar{Q}$, $c'_4 = c_1 c_4$, $c'_5 = c_1 c_5$, $c'_6 = c_1 c_6$, and $d'_j = c_1 d_j$.⁷

Making the appropriate substitutions into Equation (8) and denoting $\Delta^2 N_{t+1} \equiv \Delta N_{t+1} - \Delta N_t$ the second first-order condition becomes:

$$0 = \mathcal{E}_t \left[c_1 \Delta^2 N_{t+1} + \psi_t - k + c_2 \Delta P_{C,t+1} + c_3 \Delta \omega_{t+1} + \sum_{j=1}^{11} d'_j \Delta \text{DUM}_{j,t+1} + c'_4 \Delta \text{HDD}_{t+1} + c'_5 \Delta \text{CDD}_{t+1} + c'_6 + \Delta \eta_{t+1} + c_1 \Delta \epsilon_{t+1} \right] \quad (21)$$

To estimate the model, I substitute Equation (12) for ψ_t in Equation (21). Also, because estimation is by GMM, I drop the expectation operator and use actual values of variables dated at $t + 1$:

$$0 = c_1 \Delta^2 N_{t+1} + \exp \left(b_0 + \sum_{j=1}^{11} b_j \text{DUM}_{jt} \right) P_t^{\alpha_1} (\sigma_t^2)^{\alpha_2} N_t^{-\alpha_3} - k + c_2 \Delta P_{C,t+1} + c_3 \Delta \omega_{t+1} + \sum_{j=1}^{11} d'_j \Delta \text{DUM}_{j,t+1} + c'_4 \Delta \text{HDD}_{t+1} + c'_5 \Delta \text{CDD}_{t+1} + c'_6 + \Delta \eta_{t+1} + c_1 \Delta \epsilon_{t+1} \quad (22)$$

⁷Equation (20) is an expanded version of a model that has been used by a number of other authors. See, for example, Williams and Wright (1991), Routledge, Seppi, and Spatt (2000), and Schwartz and Smith (2000).

The model is closed by including Equation (12) for the marginal convenience yield. Together, Equations (20), (22), and (12) describe the evolution of the state variables P_t , N_t , and ψ_t .

DATA AND ESTIMATION

The Data

Estimation uses weekly data for January 1, 1984, to January 31, 2001, for crude and heating oil, and January 1, 1985, to January 31, 2001, for gasoline (reflecting the later start of gasoline futures trading). Daily futures settlement prices for the nearest contract (often the spot contract), the second-nearest, and the third-nearest were used to estimate spot prices and volatility.

The spot price is estimated from the nearest and next-to-nearest active futures contracts by extrapolating the spread between these contracts backwards to the spot month:

$$P_t = F1_t(F1_t/F2_t)^{n_{0t}/n_1} \quad (23)$$

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 first and second contracts.

From these daily spot prices, I compute weekly estimates of volatility. To do this, one must account for non-trading days. If the spot price followed a geometric Brownian motion, one could simply divide log price changes by the square root of the number of intervening days and then calculate the sample variance. However, as is well known, on average the variance of n -day log price changes is less than n times the variance of one-day changes, when n includes non-trading days. 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, with no holidays, prices for Tuesday through Friday would be assigned an interval of one day, but Monday would be assigned an interval of three days, because of the two-day weekend. With holidays, some prices could be assigned to intervals of two, four, or even five days. For each interval set, I calculate the sample standard deviation of log price changes for the entire 16- or 17-year sample for each commodity. Let \hat{s}_n denote this sample standard deviation for an interval of n days. The “effective” daily log price change for each trading

day is then computed as:

$$r_\tau = \frac{(\log P_\tau - \log P_{\tau-n})}{\hat{s}_n/\hat{s}_1} \quad (24)$$

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

$$\hat{\sigma}_t = \sqrt{\frac{1}{N-1} \sum_{\tau=1}^N (r_{t\tau} - \bar{r}_t)^2} \quad (25)$$

where N is the number of “effective” days in the five-week interval. Equation (25) 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 the measure of volatility, σ_t .

The T -period net marginal convenience yield, $\psi'_{T,t} = \psi_{T,t} - k_T$, is computed weekly from Equation (2) using the futures price and estimated spot price for the Wednesday of each week:

$$\psi'_{T,t} = (1 + R_{T,t})P_t - F_{T,t} \quad (26)$$

I use the futures price corresponding as closely as possible to a three-month interval from the spot price, and the three-month Treasury bill rate for the interest rate $R_{T,t}$. These net marginal convenience yields are converted to weekly terms.

For each commodity, there are periods when $\psi'_{T,t}$ is negative. By definition, gross marginal convenience yield must always be positive, so I estimate k for each commodity as $\hat{k} = |\min \psi'_t|$, and then compute gross marginal convenience yield as $\psi_t = \psi'_t + \hat{k}$.

To obtain a weekly series for the opportunity cost ω_t from Equation (14), I need estimates of μ and λ , and the average value of σ , for each commodity. I estimate these parameters from an OLS regression of the discrete-time version of Equation (13):

$$\Delta \log P_t = \alpha - \lambda P_{t-1} + \sigma \epsilon_t \quad (27)$$

so that $\hat{\mu} = \hat{\alpha}/\hat{\lambda} + \hat{\sigma}^2/2\hat{\lambda}$. The resulting estimates of μ , λ , and σ for crude oil, heating oil, and gasoline respectively are: $\hat{\mu} = \$20.44, 57.2¢$, and $58.6¢$, $\hat{\lambda} = 0.00114, 0.00050$, and 0.00071 ; and $\hat{\sigma} = 0.050, 0.052$, and 0.059 .

Finally, the U.S. Department of Energy, *Monthly Energy Review*, provides weekly production and inventory data. The numbers are

announced on Tuesday evenings, so that the information is incorporated in the prices and convenience yields each Wednesday.

Augmented Dickey-Fuller unit root tests on P_t , N_t , ψ_t , and σ_t (run with a constant, and then a constant and trend, and with six lags in the equation) reject a unit root in all cases. Thus in much of the empirical analysis that follows, I work with variables in levels.

Estimation Method

I estimate Equations (20), (22), and (12) using Generalized Method of Moments (GMM). GMM is an instrumental variables procedure that minimizes the correlation between variables known at time t and the equation residuals, and is thus a natural estimator for an Euler equation model such as this one.

Equations (20) and (22) include the “structural” error terms η_t and ϵ_t , which represent unobserved shocks to cost and demand. These errors may be serially correlated, and appear in differenced form in Equation (22). Also, actual values for variables at time $t + 1$ are used in place of expectations, which introduces expectational errors. Thus the equations will have composite error terms with a possibly complex autocorrelation structure. The GMM procedure uses an autocorrelation-robust weighting matrix and yields autocorrelation-robust standard errors. However, the error structure has implications for the choice of instruments.

By definition, the expectational errors are uncorrelated with any variable known at time t . The structural errors, however, may be correlated with endogenous variables. Hence I use instruments that can reasonably be viewed as exogenous: the seasonal dummy variables, the time trend, heating and cooling degree days, and the following unlagged, lagged once, and lagged twice: the exchange-weighted value of the U.S. dollar (EXVUS), the New York Stock Exchange Index (NYSE), the three-month Treasury bill rate (TBILL), the rate on Baa corporate bonds (BAA), and the Commodity Research Bureau’s commodity price index (CRB). I also include four endogenous variables lagged two periods: the spot price, production, inventory, and convenience yield. With the constant term, this gives a total of 34 instruments.

The minimized value of the objective function from the GMM procedure times the number of observations provides a statistic, J , which is distributed as χ^2 with degrees of freedom equal to the number of instruments times the number of equations minus the number of parameters. This statistic is used to test the model’s overidentifying

restrictions, and hence the hypothesis that agents are optimizing with rational expectations.

Volatility

Price, inventories, and convenience yield are endogenous, and can depend directly or indirectly on volatility, as well as on the exogenous variables. A natural question is whether these market variables, or other exogenous variables, predict volatility.

To investigate this, I estimate vector autoregressions (VARs) relating the three market variables and volatility to each other and to a set of exogenous variables, and then test the predictive power of each variable. The VAR includes six lags of each of the four variables and six lags of the following exogenous variables: the Three-month Treasury bill rate, the Baa bond rate, the exchange-weighted value of the dollar, and the monthly dummies. For each equation, I test whether all lags of a particular variable can be excluded as explanators of the dependent variable. I find that the spot price, inventories, and convenience yield all have no predictive power with respect to volatility for crude and heating oil, consistent with the view that volatility is exogenous. However, the spot price and convenience yield are significant predictors of volatility for gasoline. This could simply reflect the fact that past values of the spot price affect past values of volatility, which in turn affect its current value.

Given these results, in what follows I treat volatility as exogenous. For simulation purposes, I use a simple sixth-order autoregression (sufficient to capture the AR structure of volatility), along with six lags of the three-month Treasury bill rate, the Baa corporate bond rate, the exchange-weighted value of the dollar, the CRB commodity price index, and monthly dummy variables to generate forecasts of volatility.

GMM Estimation

Table 1 shows the results of estimating Equations (20), (22), and (12) as a system by GMM. In addition to the 10 coefficients in the table, there are another 22 coefficients (not shown) for the 11 monthly time dummies: the b_j s in the marginal convenience yield Equation (12), and the d_j s in the demand Equation (18) and thus in Equations (20) and (22). The t statistics are based on autocorrelation-consistent standard errors. The J statistics, distributed as χ^2 (71) for crude oil and χ^2 (70) for heating oil and gasoline, are all insignificant at the 5% level, so we cannot reject the overidentifying restrictions.

TABLE I
Estimates of Three-Equation System

Parameter	Crude oil	Heating oil	Gasoline	
	NOB = 881	NOB = 881	NOB = 829	
			(1)	(2)
c_0	53.63 (20.69)	-30.48 (-5.33)	5.409 (0.99)	10.98 (2.03)
c_1	0.426 (4.76)	1.478 (10.16)	1.568 (11.00)	1.541 (11.79)
c_2	—	2.927 (43.35)	2.780 (46.67)	2.728 (45.34)
c_3	-0.594 (-12.53)	0.946 (4.31)	-0.018 (-1.06)	-0.035 (-2.15)
c'_4	0.001 (0.97)	0.016 (8.74)	-0.007 (-4.75)	-0.007 (-5.13)
c'_5	0.001 (0.20)	0.001 (0.32)	0.006 (1.67)	0.006 (1.69)
c'_6	0.015 (7.97)	-0.012 (-4.38)	0.007 (2.72)	0.009 (4.10)
α_1	0.932 (22.63)	0.951 (15.4)	0.841 (21.61)	0.633 (12.7)
α_2	0.097 (8.56)	0.092 (7.10)	0.036 (3.85)	0.003 (0.25)
α_3	1.946 (9.11)	2.196 (12.03)	-0.341 (-2.73)	1.1
J	79.73	77.70	82.24	88.62

Note. Table shows GMM estimates of Equations (20), (22), and (12), with t statistics in parentheses. Estimates of 22 parameters for monthly dummy variables are not shown. For gasoline, the estimate of α_3 is negative, so the model is re-estimated with α_3 constrained to equal 1.1. The J statistics are distributed as $\chi^2(71)$ for crude oil and $\chi^2(70)$ for heating oil and gasoline; the critical 5% values are 91.40 and 90.32, respectively.

In each case, the marginal cost of production is increasing ($\hat{c}_1 > 0$ and significant). The price of crude oil is the most important determinant of marginal cost for heating oil and gasoline (\hat{c}_2 is between 2.7 and 2.9, so a \$1 per-barrel increase in the crude price, i.e., $\$1/42 = 2.4$ cents per gallon, leads to a commensurate increase in the per-gallon price of heating oil or gasoline). The opportunity cost ω_t affects total marginal cost as predicted only for heating oil: c_3 is close to 1 and significant for heating oil, but negative for crude and gasoline.

Apart from the constant and time dummies, the marginal value of storage, ψ , is characterized by the coefficients α_1 , α_2 , and α_3 in Equations (12) and (22). For crude and heating oil, the estimates of these coefficients are all positive and significant, and consistent with a well-behaved marginal value of storage function: $\hat{\alpha}_3 > 1$, $\hat{\alpha}_1$, the price

elasticity of ψ is close to 1, and $\hat{\alpha}_2 > 0$, i.e., ψ is increasing with volatility. For gasoline, $\hat{\alpha}_3$ was negative, so the model was re-estimated with α_3 constrained to equal 1.1. The estimates of c_1, \dots, c_6 are largely unchanged, but $\hat{\alpha}_1$ drops from 0.84 to 0.63, and $\hat{\alpha}_2$ becomes insignificant.

Thus the model fits the theory very well for heating oil, but less so for crude oil and gasoline. For both crude oil and gasoline, the net demand function is upward sloping, but does not depend on the marginal opportunity cost as predicted. Also, the unconstrained marginal value of storage function for gasoline is increasing in the level of inventories N_t , and when α_3 is set to equal 1.1, the elasticity with respect to volatility becomes zero.

SIMULATIONS

Dynamic simulations, in which Equations (20), (22), and (12) are solved as a system, can be used to evaluate the ability of the model to replicate the behavior of the endogenous variables, and to study the effects of shocks to volatility or other variables. I focus on heating oil, for which all of the coefficient estimates are consistent with the predictions of the theory. The first simulation covers the 10-week period of August 8 to October 3, 1990, when Iraq invaded Kuwait and the price of heating oil jumped from 60¢ to about \$1 per gallon. The second covers the last 10 weeks of the sample: November 29, 2000 to January 31, 2001. Both simulations are dynamic; actual values of P_t , N_t , ψ_t , and σ_t are used only prior to the starting date.

Simulated and actual values of the spot price, inventory level, and convenience yield are shown in Figures 1–3 (each figure shows both




FIGURE 1
Heating oil spot price.




FIGURE 2
Heating oil inventories.




FIGURE 3
Heating oil convenience yield.

simulations). For the 1990 period, the simulated spot price tracks the sharp increase during August 22 to 29, but not the temporary decline on September 5, and is again close to the actual over the last three weeks of the period. Actual inventories rose by 11% over this period, but the simulated series increases by much more. Finally, simulated convenience yield closely tracks the actual series throughout. In the second simulation, the actual spot price fell from about \$1.10 per gallon to about 83¢, and the simulated series closely tracks this decline. The model again over-predicts inventories. Finally, the actual convenience yield fluctuated widely, and the simulated series replicates the directional movements but not the magnitude of the fluctuations.

Overall, the model replicates the dynamics of the heating oil spot price and convenience yield well, given the sharp movements that occurred during the two simulation periods. The model does not, however, capture the dynamics of inventories very well. This is not surprising; Equations (20) and (22) explain first and second differences of inventories, so that prediction errors in the level of inventories will accumulate over time.

Finally, I repeated the second simulation, adding a shock to volatility. I increased the entire trajectory of volatility by 0.0448, which is one standard deviation of its level over the 1984–2001 sample, and then re-solved the model. As seen in Figures 1–3, this volatility shock has a substantial effect on convenience yield, but only a small effect on the spot price and inventories. Convenience yield increases because the value of storage depends directly on volatility, and the higher convenience yield raises inventories. The increase in volatility also increases the opportunity cost of production and thus the spot price, but only slightly.

CONCLUSIONS

In principle, volatility should affect market variables through the marginal value of storage and through the opportunity cost component of marginal cost. For the petroleum complex, and in particular for heating oil, changes in volatility do influence market variables, although the effects are not large. Most of the impact is on the convenience yield, and to a lesser extent, inventories. Thus, accounting for changes in volatility can help explain changes in the spot-futures spread, but not changes in the spot price itself. As for volatility, market variables do little to explain its behavior, and it can be viewed as exogenous.

These results give partial support to the theory of commodity price dynamics presented at the outset. For heating oil, the results fit the theory very well, but for crude oil and gasoline the results are less clear cut. There are several possible reasons for this. First, my use of a quadratic approximation to calculate the marginal opportunity cost may be unwarranted. Second, the intertemporal optimization that underlies the model may be inconsistent with high-frequency data. Changes in market variables may affect production decisions more slowly than can be captured by the weekly differences that appear in the estimating equations.⁸ Finally, it is unclear how much of a commodity's short-run price

⁸Also, as Borenstein and Shepard (2002) show, gasoline prices adjust slowly to crude oil price shocks, which may reflect market power in gasoline markets.

movements can be explained by rational optimizing behavior. Price variation may be partly the result of speculative noise trading or herd behavior, rather than fundamentals.⁹

APPENDIX: DERIVATION OF OPPORTUNITY COST

In this Appendix, I derive Equation (14) for the opportunity cost of production, assuming that the spot price P follows the mean-reverting process given by Equation (13).

Let $V(P)$ be the value of the option to produce a unit of the commodity. It is easily shown that $V(P)$ must satisfy the following equation (see Dixit & Pindyck, 1994):

$$\frac{1}{2}\sigma^2 P^2 V''(P) + [r - \rho + \lambda(\mu - P)] PV'(P) - rP = 0 \quad (\text{A.1})$$

where r is the risk-free rate and ρ is the risk-adjusted return on the commodity. Thus the expected return “shortfall” is $\delta = \rho - \lambda(\mu - P)$. Also, the solution must satisfy the boundary conditions

$$V(P^*) = P^* - c \quad (\text{A.2})$$

$$V'(P^*) = 1 \quad (\text{A.3})$$

where P^* is the critical price that triggers production of an incremental unit, and c is marginal cost. The solution to Equation (A.1) is:

$$V(P) = AP^\theta h(P) \quad (\text{A.4})$$

where θ is given by Equation (16), and, letting $b = 2\theta + 2(r - \rho + \lambda\mu)/\sigma^2$, $h(P) = H(\frac{2\lambda}{\sigma^2}P; \theta, b)$. Here, $H()$ is the confluent hypergeometric function:

$$H(x; \theta, b) = 1 + \frac{\theta}{b}x + \frac{\theta(\theta + 1)}{b(b + 1)2!}x^2 + \frac{\theta(\theta + 1)(\theta + 2)}{b(b + 1)(b + 2)3!}x^3 + \dots \quad (\text{A.5})$$

I use a quadratic approximation to $h(P)$:

$$h(P) \approx 1 + \gamma_1 P + \gamma_2 P^2 \quad (\text{A.6})$$

⁹For example, Roll (1984) found that only a small fraction of price variation for frozen orange juice can be explained by fundamentals such as the weather, and Pindyck and Rotemberg (1990) found high levels of price correlation across commodities that are inconsistent with prices driven solely by fundamentals.

where γ_1 and γ_2 are given by Equation (15). Thus $V(P) \approx AP^\theta(1 + \gamma_1 P + \gamma_1\gamma_2 P^2)$. Substituting into boundary conditions (A.2) and (A.3) gives two equations in P^* and the constant A . Divide one by the other to eliminate A and rearrange, yielding:

$$P - c = \frac{P(1 + \gamma_1 P + \gamma_1\gamma_2 P^2)}{\theta + \gamma_1(\theta + 1)P + \gamma_1(\gamma_2\theta + 2\gamma_2 + \gamma_1)P^2 + 3\gamma_1^2\gamma_2 P^3 + 2\gamma_1^2\gamma_2^2 P^4} \quad (\text{A.7})$$

Next, expand the right-hand side of this equation in a Taylor series around $P = \mu$, take a quadratic approximation, and set $c = \mu$ to obtain Equation (14).

BIBLIOGRAPHY

- Ball, C. A., & Torous, W. N. (1999). The stochastic volatility of short-term interest rates: Some international evidence. *The Journal of Finance*, 54, 2339–2359.
- Borenstein, S., & Shepard, A. (2002). Sticky prices, inventories, and market power in wholesale gasoline markets. *RAND Journal of Economics*, Spring 2002, bf 33, 116–139.
- Campbell, J. Y., Malkiel, B., Lettau, M., & Xu, Y. (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance*, 56, 1–43.
- Casassus, J., & Collin-Dufresne, P. (November 2001). ‘Maximal’ convenience yield model implied by commodity futures. Unpublished working paper.
- Considine, T. J. (1997). Inventories under joint production: An empirical analysis of petroleum refining. *Review of Economics and Statistics*, 1997, 493–502.
- Considine, T. J., & Heo, E. (2000). Price and inventory dynamics in petroleum product markets. *Energy Economics*, 22, 527–547.
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.
- Eckstein, Z., & Eichenbaum, M. S. (1985). 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.
- Franks, J., R., & Schwartz, E. (1991). The stochastic behaviour of market variance implied in the prices of index options. *The Economic Journal*, 101, 1460–1475.
- Hull, J., & White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42, 281–300.
- Litzenberger, R. H., & Rabinowitz, N. (1995). Backwardation in oil futures markets: Theory and empirical evidence. *The Journal of Finance*, 50, 1517–1545.

- McDonald, R., & Shimko, D. (June 1998). The convenience yield of gold. Unpublished working paper, Northwestern University.
- Milonas, N. T., & Henker, T. (2001). Price spread and convenience yield behavior in the international oil markets. *Applied Financial Economics*, 11, 23–26.
- Miron, J. A., & Zeldes, S. P. (1988). Seasonality, cost shocks, and the production smoothing model of inventories. *Econometrica*, 56, 877–908.
- Pindyck, R. S. (1993). The present value model of rational commodity pricing. *The Economic Journal*, 103, 511–530.
- Pindyck, R. S. (1994). Inventories and the short-run dynamics of commodity prices. *RAND Journal of Economics*, 25, 141–159.
- Pindyck, R. S., & Rotemberg, J. J. (1983). Dynamic factor demands and the effects of energy price shocks. *American Economic Review*, 73, 1066–1079.
- Ramey, V. A. (1991). Nonconvex costs and the behavior of inventories. *Journal of Political Economy*, 99, 306–334.
- Roll, R. (1984). Orange juice and the weather. *American Economic Review*, 74, 861–880.
- Routledge, B. R., Seppi, D. J., & Spatt, C. S. (2000). Equilibrium forward curves for commodities. *Journal of Finance*, 55, 1297–1338.
- Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52, 923–973.
- Schwartz, E. S., & Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46, 893–911.
- Williams, J. C., & Wright, B. D. (1991). *Storage and commodity markets*. Cambridge, UK: Cambridge University Press.