Initial Margin Policy and Stochastic Volatility in the Crude Oil Futures Market

Theodore E. Day
University of Texas at Dallas

Craig M. Lewis
Vanderbilt University

This article examines the relationship between the volatility of the crude oil futures market and changes in initial margin requirements. To closely match changes in futures market volatility with the corresponding changes in margin requirements, we infer the volatility of the futures market from the prices of crude oil futures options contracts. Using a mean-reverting diffusion process for volatility, we show that changes in margin policy do not affect subsequent market volatility.

Dramatic short-run increases in the volatility of financial markets such as the extraordinary volatility following the stock market decline in October 1987 have renewed interest in the relation between initial margin requirements and market volatility. Hardouvelis (1988) argues that there is a significant negative relation between initial margin requirements on stocks and stock market volatility. This contradicts earlier

This article was presented at the Third Annual Winter Finance conference at the University of Utah, the Fifth Annual Conference of the Financial Options Research Centre at the University of Warwick, and at Vanderbilt University. We would like to thank Kaushik Amin, Cliff Ball, Don Chance, Andrew Karolyi, Andrew Lo (the editor), Raymond Melkomian, Barry Schacter, Paul Seguin, Hans Stoll, an anonymous referee, and conference participants for many helpful comments. This research was supported through the Financial Markets Research Center at Vanderbilt University by a grant from the New York Mercantile Exchange and by the Dean's Fund for Faculty Research at the Owen Graduate School of Management. Address correspondence to Theodore E. Day, School of Management, University of Texas at Dallas, Box 830688, Sta Jo 5.1, Richardson, TX 75083-0688.

The Review of Financial Studies Summer 1997 Vol. 10, No. 2, pp. 303–332

research by Moore (1966) and Officer (1973) which concludes that margin requirements have no effect on the volatility of the stock market. Hsieh and Miller (1990) present evidence suggesting that the relation between margin requirements on stock and volatility found by Hardouvelis may be spurious since his methodology is biased in favor of finding a statistical relation between margins and volatility. Instead they find that the relation between margin changes and stock market volatility is consistent with a tendency for the Federal Reserve Board to increase initial margin requirements as a response to increases in the volatility of the stock market. This interpretation of the evidence is consistent with that of Schwert (1988).

Seguin (1990) examines the relation between margin requirements and volatility by comparing the volatility of NASDAQ securities during the periods before and after each stock becomes eligible for margin trading. For each of the stocks that he examines, volatility declines in the period following the approval of margin trading. Seguin and Jarrell (1993) examine the impact of margin trading on NASDAQ securities during the stock market crash on October 19, 1987. In spite of the fact that marginable securities had greater abnormal volume during the crash, Seguin and Jarrell (1993) find that the impact of the crash on the prices of marginable securities was no greater than the corresponding impact on the prices of securities that were not eligible for margin trading.

The debate over the relation between margin policy and the volatility of financial markets extends to the initial margin requirements in futures markets. The margin requirement for stocks is set by the Federal Reserve Board. In contrast, changes in the required margins for most futures contracts are initiated by the futures clearinghouse without prior approval from the Commodity Futures Trading Commission. This permits futures margins to be used as an endogenous risk management tool by the clearinghouse [see Fenn and Kupiec (1993) and Day and Lewis (1996)]. The primary difference between initial margin requirements in the stock and futures markets is that the initial margin requirements in the stock market are set with respect to the "loan value" of the position, whereas initial margins on futures contracts are designed to guarantee that the buyer and seller will be able to fulfill their contractual obligations. For this reason, the initial margin requirement for stocks is expressed in terms of the percentage of the purchase price that must be paid in cash, while margin requirements in futures markets are set by the respective futures exchanges in terms of a fixed dollar amount per contract. These differences have led to a controversy among regulators because futures margins tend to be a much smaller percentage of the contract value than the current 50% margin required for stocks. In part, the greater leverage available from trading in stock index futures has been blamed by some for the excessive stock market volatility during October 1987. This has led to a number of proposals to consolidate the authority to regulate margins in both the stock and futures markets with the Federal Reserve Board. Such a change in the regulation of margin requirements would supposedly allow regulators to control the volatility of the markets through both the equalization of the leverage available in the stock and futures markets and through increases of initial margins to more prudent levels at the appropriate times.

The impact of margin requirements on the volatility of futures markets has been examined by Kupiec (1989) and Kupiec and Sharpe (1990). Kupiec and Sharpe model the effect of margin requirements on portfolio choice in an economy with heterogeneous investors. Their analysis shows that margin requirements may either increase or decrease the volatility of prices according to the nature of the heterogeneity across investors. Kupiec examines the relation between stock market volatility and the margin requirements for S&P 500 futures contracts and finds that high margin rates in the futures market tend to be associated with above average volatility in the (cash) market for the S&P 500 stocks. The relation between margin requirements and futures market volatility has been examined by Fishe et al. (1990) for a sample of 10 futures contracts on agricultural commodities and precious metals. As was the case in Kupiec, they find no evidence of a statistical relation between margin requirements and the subsequent volatility of the futures market.

A change in margin requirements following a perceived change in the volatility of a given market is usually implemented quickly, often within the same trading session in futures markets. However, the need to use a time series of daily returns to estimate the volatility of the market has forced most previous research to examine the relation between volatility and margin requirements over relatively long time intervals [e.g., Hsieh and Miller (1990) use a monthly time interval]. In highly volatile markets, this can obscure the true relation between margin requirements and volatility. Although this problem can be circumvented by using transactions data to focus on the intraday volatility of the market in question, the approach taken here is to infer both the spot volatility and the parameters of a mean-reverting diffusion process for volatility from the prices of call options on futures contracts. The joint estimation of the spot volatility and the stochastic process for volatility has several advantages. First, since option prices should respond instantaneously to changes in volatility, inferring volatility from option prices permits us to closely match changes in the volatility of the futures market with both the corresponding changes in margin requirements and with any feedback effects on volatility from the change in margin requirements. Further, to the extent that increases in the volatility of a particular market are transitory or mean reverting, the parameters of the stochastic process can provide an estimate of the time required for the volatility of the market to return to more normal levels.

This article extends previous research by examining the relation between changes in margin requirements and changes in consensus forecasts of the volatility of the crude oil futures market. This market is unique with respect to the frequency with which the New York Mercantile Exchange adjusts the initial margin requirements for the most actively traded futures contracts (the nearest to delivery or spot month contract). Since the beginning of trading in options on crude oil futures in late 1986, there have been 30 changes in the initial margin requirements for the spot month futures contract. In contrast, the margin requirement for stocks has been changed only 19 times since 1945 and has been constant at the current level of 50% since 1974. Consequently the market for crude oil futures provides a unique opportunity to study the behavior of ex ante forecasts of futures market volatility during the time interval surrounding changes in initial margin requirements.

The focus of our analysis on the interaction between changes in initial margin requirements and the ex ante volatility of the futures markets represents a departure from previous work, which focuses on the relation between margin requirements and realized volatility. Building on the work of Ball and Roma (1994), Heston (1993), Hull and White (1987), Johnson and Shanno (1987), Stein and Stein (1992), and Wiggins (1987), we examine the impact of changes in initial margin policy on ex ante volatility within an asset pricing framework that permits volatility to be stochastic. Given the assumption that the prices of futures options accurately reflect a consensus forecast of future volatility, generalized method of moments estimation can be used to obtain estimates of the short-run or *spot* volatility of the futures market. This procedure also provides estimates of the parameters of the diffusion process for the volatility of the crude oil futures market. The impact of changes in initial margin requirements on ex ante volatility can be examined by using these parameter estimates and the implied spot volatility of the crude oil futures market to create a time series of consensus forecasts of future volatility. Since implied stochastic volatilities from the options market reflect the expectations of market participants concerning the potential range of future price movements, the stochastic volatilities implicit in the prices of crude oil futures options provide new evidence concerning the effect of changes in margin requirements on the volatility of futures markets.

This article is organized as follows. Section 1 describes the estimation of risk in markets with stochastic volatility and characterizes the asset pricing framework. In Section 2 we discuss the estimation of the stochastic volatility model from the prices of call options on crude oil futures contracts using generalized method of moments. Section 3 describes the data. The estimated parameters of the mean-reverting process for stochastic volatility are discussed in Section 4. The behavior of forecasts of forward volatility from the options market during the period surrounding changes in initial margin requirements is examined in Section 5. Section 6 examines the causal relation between margin requirements and futures market volatility using Granger causality tests [see Granger (1969)]. Section 7 concludes the article.

1. The Estimation of Risk in Markets with Stochastic Volatility

Many stochastic volatility models have been proposed in the literature. In this article we assume that the variance of crude oil futures follows a square-root diffusion process. The mean-reverting nature of this process is attractive for several reasons. First, Day and Lewis (1993) show that volatility shocks in the crude oil futures market are persistent and mean reverting. Second, the relation between the spot volatility and the long-run volatility can be examined directly. Finally, this process is analytically tractable. Cox et al. (1985) implicitly solve for the moment-generating function of the average of this process in the derivation of their formula for the price of a discount bond. Ball and Roma (1994) use this result to derive a simple closed-form expression for the expected value of average future volatility. Given this result, estimates of the parameters of the diffusion process for futures market volatility can be used to determine the market forecasts of volatility implicit in the prices of call options on crude oil futures contracts. These parameter estimates also can be used to generate the term structure of expected average volatility and forward expected average volatilities.

Note that pricing options using the standard techniques for risk-neutral pricing is not possible when volatility is a state variable, since a perfect hedge for the risk associated with changes in the volatility of the underlying asset does not exist. Consequently the valuation of derivative securities is no longer preference free and the market price of volatility risk, $\phi(\nu)$, is required to determine the price of an option on a futures contract.

Assume that the dynamics of the futures price are given by

$$dF = \mu F\,dt + \sqrt{\nu} F\,dz_1(t), \tag{1}$$

where μ and ν are respectively the instantaneous mean and variance of the futures price and $z_1(t)$ is a standard Wiener process. The instantaneous (spot) variance of the futures price is assumed to evolve stochastically according to the diffusion process,

$$d\nu = \alpha(\beta - \nu)\,dt + \zeta\sqrt{\nu}\,dz_2(t), \tag{2}$$

where α determines the speed at which the instantaneous variance reverts to its long-run mean β and $z_2(t)$ is a standard Wiener process such that the instantaneous correlation between increments in $z_1(t)$ and $z_2(t)$ is ρ.

Let the price of a call option on a futures contract be denoted by C. Given that the cost of carrying a futures contract is zero, the partial differential equation that must be satisfied by the price of a futures option is

$$\begin{aligned}\frac{\partial C}{\partial t} &+ \frac{1}{2}\nu F^2\frac{\partial^2 C}{\partial F^2} + \rho\nu\zeta F\frac{\partial^2 C}{\partial F\partial\nu} + \frac{1}{2}\zeta^2\nu\frac{\partial^2 C}{\partial\nu^2}\\ &+ (\alpha(\beta-\nu) - \phi\nu)\frac{\partial C}{\partial\nu} - rC = 0.\end{aligned} \tag{3}$$

This differential equation can be solved to obtain the price of an option on a futures contract by applying the appropriate boundary conditions and specifying ϕ. Note that since crude oil futures contracts are American options, the boundary conditions must incorporate the possibility of early exercise at each moment prior to expiration. Since there is no analytic solution for this problem, Equation (3) is solved numerically.[1]

The solution to the partial differential equation is a function of the spot volatility (a state variable) and five parameters that define the bivariate system of stochastic differential equations specified by Equations (1) and (2), that is, $C = C(\nu, \alpha, \beta, \zeta, \rho, \phi)$. Equation (3) shows that there is a linear relation among α, β, and ϕ that does not permit their unique identification. Because the market price of

[1] We use finite difference methods to estimate the value of call options on crude oil futures contracts. The computation of a numerical solution to a partial differential equation of the type represented by Equation (4) is complicated by the fact that the partial differential equation includes a mixed partial derivative. For partial differential equations having mixed derivatives, Gourlay and McKee (1977) show that the "line hopscotch" approach developed in Gourlay (1970) provides more accurate solutions than do techniques such as "ordered odd-even" hopscotch, alternating direction implicit, and local one dimensional methods. This technique, which we use to compute call option prices, has previously been used by Kuwahara and Marsh (1992) and Wiggins (1987) to price options on financial instruments having a stochastic volatility. The partial differential equation is solved after transforming the variables to make the price of an option a function of the natural logarithm of the futures price ($\ln F$) and the natural logarithm of the spot volatility ($\ln \nu$).

risk, ϕ, cannot be separated from the terms that govern the drift of the volatility diffusion, we estimate a "risk-neutralized" version of the partial differential equation given by Equation (3). That is,

$$\frac{\partial C}{\partial t} + \frac{1}{2}\nu F^2 \frac{\partial^2 C}{\partial F^2} + \rho\nu\zeta F \frac{\partial^2 C}{\partial F \partial \nu} + \frac{1}{2}\zeta^2 \nu \frac{\partial^2 C}{\partial \nu^2} + \alpha^*(\beta^* - \nu)\frac{\partial C}{\partial \nu} - rC = 0. \tag{4}$$

Under this specification [see Heston (1993)], the instantaneous variance of the futures price follows the "risk-neutralized" square root diffusion

$$d\nu = \alpha^*(\beta^* - \nu)\, dt + \zeta\sqrt{\nu}\, dz_2(t), \tag{5}$$

where $\alpha^* = \alpha + \phi$ and $\beta^* = \alpha\beta/(\alpha + \phi)$. Although the long-run mean volatility and the speed of reversion to the mean cannot be separately identified using the parameters of the risk-neutral diffusion implicit in the prices of futures options, the risk-neutral partial differential equation [Equation (4)] can be used to exactly determine the implied value for the spot volatility of the futures price. In the next section we discuss the statistical procedures that are used to jointly estimate the spot volatility and the parameters of the risk-neutral diffusion process.

2. Parameter Estimation of the Stochastic Volatility Model

The parameters of the stochastic volatility model and the spot volatility of the crude oil futures are estimated using Hansen's (1982) generalized methods of moments (GMM). Conceptually our implementation of this procedure is similar to the estimation of an implied volatility using the Black–Scholes model. However, under the assumption that volatility is stochastic, the price of the call option is a nonlinear function of the spot volatility, ν, and four parameters—α^*, β^*, ζ, and ρ—that are assumed to be constant over the sample period. Since the price of each call option must be calculated numerically, estimation of the underlying parameters is computationally intensive. To ease the burden of computation, we implement a two-stage estimation approach. In the first stage, a subset of the data is used to estimate $\Theta_p = (\alpha^*, \beta^*, \zeta, \rho)$ along with the spot volatility for each of the days included in the first-stage estimation sample. This produces a parameter vector $\mathbf{\Theta} = (\Theta_p, \Lambda)$, where $\mathbf{\Lambda}$ is a vector of spot volatilities for each day in the estimation sample, $\mathbf{\Lambda} = (\nu_1, \ldots, \nu_{49})$. In the second stage, we estimate the spot volatilities for the remaining days in the sample under the constraint that the parameter vector $\mathbf{\Theta}_p$ equals the estimate obtained in the first-stage of the procedure.[2]

2.1 First-stage estimation

To ensure that the dynamics of the volatility during the sample period are accurately captured by the diffusion parameters that we estimate, the subset of data that is used in the first stage of the estimation process spans the entire sample period. This subset is created by selecting the at-the-money option for each expiration series on every 20th day of the sample, resulting in a sample of 121 observations from 49 different days.[3]

The first-stage estimate of the parameter vector, $\mathbf{\Theta}$, is determined using the vector of sample moments,

$$g_T(\Theta) = \frac{1}{\hat{T}} \sum_{t=1}^{T} \sum_{n=1}^{N_t} (C_t(X, \tau_n) - \hat{C}(X, \tau_n, \Theta)) Z_{tn},$$

where $C_t(X, \tau_n)$ is the date t market price of the at-the-money call option with strike price X and time to expiration τ_n, $\hat{C}(X, \tau_n, \Theta)$ is the model price at time t, Z_{tn} is the vector of partial derivatives of $\hat{C}(X, \tau_n, \Theta)$ with respect to the parameter vector ($\mathbf{\Theta}$) weighted by the proportion of the day t trading volume in option n, T is the total number of days in the first-stage estimation sample ($T = 49$), N_t is the number of expiration series trading on day t, and $\hat{T}$ is the total number of observations ($\hat{T} = 121$).[4] The estimation of Θ using the GMM procedure requires the minimization of the quadratic form,

$$J_T(\Theta) = g_T(\Theta)' \Omega_T^{-1} g_T(\Theta), \tag{6}$$

where

$$\Omega_T = E[((C_t(X, \tau_n) - \hat{C}(X, \tau_n, \Theta)) Z_{tn})((C_t(X, \tau_n) - \hat{C}(X, \tau_n, \Theta)) Z_{tn})'].$$

---

[2] Our sample includes every option that has at least 7 days to expiration and trading volume of at least 100 contracts during the day. The requirement that trading volume exceed 100 contracts eliminates thinly traded contracts, for which the closing trade in the option is very likely to occur prior to the closing trade in the underlying futures contract. By using only options having at least 7 days to expiration, we eliminate options for which the option premium is too small to provide a precise estimate of the implied volatility of the futures market.

[3] The first-stage sample is restricted to the at-the-money option for each expiration series in order to reduce the computational burden required to estimate the parameters of the stochastic volatility model. Since the at-the-money option is usually the most actively traded contract within a given expiration series, using at-the-money options helps to reduce the estimation error arising from a lack of simultaneity between the closing futures price and the final price quotation for the call option.

[4] At-the-money options are used because trading volume is concentrated in the at-the-money options. This further mitigates problems associated with nonsynchronous price data. At-the-money options also are the most sensitive to changes in volatility.

Hansen (1982) demonstrates that the asymptotic distribution of the GMM estimator is

$$\sqrt{\hat{T}}(\Theta_T - \Theta_0) \overset{a}{\sim} N(0, (D_t \Omega_T^{-1} D_T)^{-1}),$$

where D_T is the Jacobian matrix of $g_T(\Theta)$ with respect to Θ evaluated at Θ_T.

The specification of the model can be tested using the minimized value of the quadratic form in Equation (6). Under the null hypothesis that the model is true, this value is asymptotically distributed as χ^2 with degrees of freedom equal to $\hat{T}$ less the number of parameters, that is,[5]

$$\hat{T} g_T(\Theta)' \Omega_T^{-1} g_T(\Theta) \overset{a}{\sim} \chi^2_{68}.$$

The estimation is implemented using the standard two-step procedure in Hansen and Singleton (1982). The first step estimates the parameter vector by minimizing Equation (6), with Ω_T^{-1} equal to the identity matrix. This provides a consistent estimator of the parameter vector, which then is used to compute the optimal weighting matrix. The second step reestimates the parameter vector using the optimal weighting matrix.

2.2 Second-stage estimation

In the second stage of the estimation procedure, the daily spot volatilities are estimated using the estimate of the parameter vector $\mathbf{\Theta}_p$ from the first step. Given an initial estimate of the spot volatility, ν_{t0}^n, an updated estimate, ν_t^n, is obtained using a Newton–Raphson procedure. At each iteration of the process, the new estimate of ν_t^n is given by

$$\nu_t^n = \nu_{t0}^n + \omega(\partial C_t^n / \partial \nu)^{-1} \varepsilon_t^n,$$

where ω is the step size for updating the spot volatility, $\partial C_t^n / \partial \nu$ is the numerical partial derivative of $C(X, \tau_n, \Theta_p)$ with respect to the spot volatility evaluated at ν_{t0}^n, and ε_t^n is equal to $C_t(X, \tau_n)$ minus $\hat{C}(X, \tau_n, \Theta_p)$. The estimate of ν_t^n is taken as acceptable when ε_t^n converges to within the desired tolerance level. If the estimate is not within the desired tolerance level, the procedure is repeated using ν_t^n in place of ν_{t0}^n.

[5] The number of degrees of freedom equals 68, which is determined by the total of the 121 observations less the sum of the number of spot volatility estimates (49) and the parameters of the square-root diffusion that are constant over time (4).

3. Data Description

The data consist of daily closing prices for call options on crude oil futures and the underlying futures contracts from the beginning of trading in the options on November 14, 1986, through March 18, 1991. The risk-free rate of interest for each option expiration series is computed for each day using the average of the bid and ask discounts for the U.S. Treasury bill whose maturity is closest to the expiration date. Bid and ask discounts for U.S. Treasury bills are collected daily from the *Wall Street Journal.*

The New York Mercantile Exchange supplied the dates and margin levels for all adjustments of initial margin requirements in the nearby futures contract (the spot month futures contract). While there were a small number of changes in margin requirements prior to the inception of trading in futures options, most of the changes in margin requirements have occurred since 1986. This period is of particular interest since it includes the August 1, 1990, invasion of Kuwait by Iraq. The sample includes 19 changes in margin requirements prior to the invasion of Kuwait and 11 changes in margin requirements in the subsequent period.

4. Estimation Results

This section presents estimates of the parameters that describe the evolution of the spot volatility of the crude oil futures market. In addition, we present summary statistics for the implied spot volatilities derived from these parameter estimates.

4.1 Parameter estimates of the "risk-neutralized" stochastic volatility model

The first stage of the GMM procedure outlined in Section 3 provides estimates of the parameters of the risk-neutralized diffusion. The parameter estimates of the risk-neutralized diffusion for the volatility of crude oil futures, along with their heteroskedasticity-consistent standard errors, are presented in Table 1. With the exception of the estimate for the correlation between the stochastic component of the futures price and volatility, the estimated parameters for the model are all more than three standard errors from zero. The χ^2 value of 0.00035, corresponding to a p-value less than 0.0001, indicates that the model fits the observed option prices quite well.

Table 1 shows that the risk-adjusted long-run mean for the volatility of the crude oil futures market (β^*) has an estimated value of 0.0718, which corresponds to an annualized standard deviation of approximately 26.8%. Although this represents a lower bound for the long-run

Table 1
Parameter estimates for the risk-neutralized stochastic volatility model

Parameter	Estimated value	t-statistic
α^*	2.18053	6.075
β^*	0.07183	4.820
ζ	0.42932	30.531
ρ	−0.15684	−0.362
χ^2 value	0.003529	
Degrees of freedom	68	

The parameter estimates for the risk-neutral stochastic volatility model are obtained from the first stage of the GMM procedure outlined in Section 2. The parameters were estimated using a subset of the data that included the at-the-money options of each expiration series for every 20th day during the sample period. t-statistics are computed using heteroskedasticity-consistent standard errors.

mean of the *true* stochastic process, this estimate is reasonable given the volatility that has been typical of the crude oil futures markets in recent years. The estimated value of the mean-reversion parameter (α^*) has an estimated value of 2.181, with a t-statistic of 6.08. The standard error of the estimate for α^* indicates that there is significant mean-reversion in the volatility of the crude oil futures market. The estimated value for the parameter that determines the time-varying standard deviation of the square-root diffusion (ζ) is 0.429. Given the dynamics for the spot volatility of the futures price, this implies that when the spot volatility is at the long-run mean (of the risk-neutral distribution), the standard deviation of the spot volatility is 11.5%. Although the estimated value for ρ is −0.157, the standard error of the estimate is 0.433, which indicates that the correlation between the futures price and random changes in volatility is not statistically significant.

4.2 Sample estimates of the implied spot volatilities

Table 2 presents summary statistics for the daily estimates of the spot volatility for the crude oil futures market. Each of the daily estimates within this series represents a weighted average of the implied spot volatilities for the three futures options nearest to expiration, where the weights are based on the numerical partial derivatives of the call price with respect to spot volatility (normalized to sum to one). The time series of implied spot volatilities has a mean value of 0.1704 and a median value of 0.0811. These mean and median values of the spot variance respectively correspond to annualized standard deviations of 41.28% and 28.48%. Table 2 indicates that the distribution of the estimates of spot volatility is highly skewed and has fat tails. This is attributable to several dramatic increases in the volatility of the crude

Table 2
Summary data for estimates of spot volatility from crude oil futures call options for the period November 14, 1986, to March 18, 1991

Statistic	Estimated value
Mean	0.1704
Median	0.0811
Standard error	0.2980
Minimum value	0.0117
Maximum value	3.1925
Skewness	5.01
Kurtosis	34.17

The sample period includes 1086 trading days.

oil futures market during the sample period. For example, during the period following the invasion of Kuwait by Iraq, the increase in the volatility of the crude oil futures is both more dramatic and more prolonged than the volatility observed during the period following the stock market crash of 1987.

The first-order autocorrelation of the daily spot volatility series is 0.95. This autocorrelation reflects the persistence in volatility that has been characterized by ARCH-type models of conditional volatility. Figure 1 illustrates that the autocorrelation is quite high even at longer lags. The autocorrelation structure of the first differences is also presented in Figure 1. Although the results are not reported, Dickey–Fuller tests reject the possibility of a unit root in favor of the alternative hypothesis that the time series of estimated spot volatilities represents a stationary time series.[6]

4.3 The parameter estimates of the "true" stochastic volatility model

It is important to note that since the elastic force and long-run mean (α^* and β^*) of the risk-neutral diffusion given by Equation (6) are functions of the parameters of the *true* stochastic process (of α, β, and ϕ), the *risk-neutral* stochastic process must be consistent with the true sample path of the volatility of the crude oil futures market. Therefore, given the functional forms for α^* and β^*, the parameters of the true stochastic process for volatility can be identified from the parameters of the risk-neutral process if an independent estimate of one of the true diffusion parameters can be made. Since the mean-reverting dif-

[6] To account for the possibility that misspecification of the time-series process for volatility causes the results of our statistical tests to be spurious, the stationarity of the spot volatility process is examined for autoregressive processes having up to 40 lags. At the 1% significance level, every test rejects the null hypothesis of a unit root in favor of the alternative hypothesis of stationarity. Similar results have been noted by Diz and Finucane (1992), Flemming (1993), and Stein (1989) for implied volatilities from stock index options.

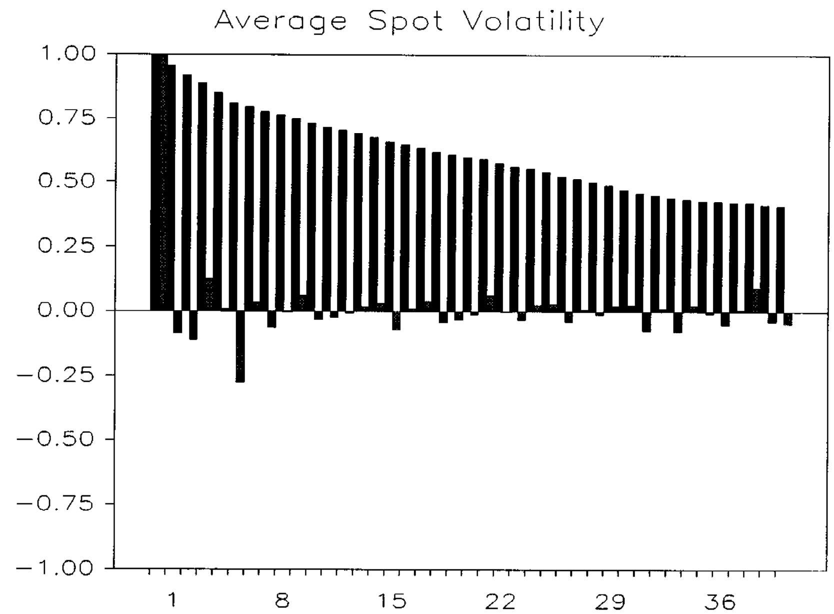

Figure 1
Autocorrelation functions
This figure presents the autocorrelation structure of the levels and the first differences of daily implied crude oil futures options spot volatilities.

fusion for the spot volatility given by Equation (2) has a steady-state distribution that asymptotically approaches a gamma distribution with a mean of β, it is tempting to estimate the long-run average volatility using the sample average for the implied spot volatilities. However, the extreme volatility of the futures market following the invasion of Kuwait suggests that the sample average of 0.1704 may overstate the true long-run volatility. To compensate for the disproportionate number of extreme observations included among our sample estimates, we use the sample median of 0.0811. This estimate is close to the sample average for the preinvasion period, which is equal to 0.0827. Note that our estimate of the long-run average volatility is consistent with the requirement that the long-run mean of the risk-neutral distribution be less than the long-run mean of the true distribution.

Given our independent estimate of the long-run mean for the true process for volatility, the risk-neutral estimates of α^* and β^* can be used to derive estimates of both the rate of mean reversion for the true volatility process and the market's required compensation per unit of volatility risk. The resulting parameter estimates for Equation (2) are

$$dv = 1.931(0.0811 - v)\,dt + 0.429\sqrt{v}\,dz_2(t),$$

where the elastic force that causes the spot volatility to revert to the long-run mean, α, has an estimated value of 1.931.[7]

To provide some indication of the speed at which the estimated mean-reversion parameter causes volatility to revert to the long-run mean, we use the estimated value for α to infer the spot volatility's "half-life." The half-life is defined as the time required for the expected future spot volatility to revert halfway to the long-run mean. The half-life is determined by finding the date, t_s, for which

$$E(v_{t_s}|v_t) = \frac{1}{2}(v_t + \beta). \tag{7}$$

Following Cox, Ingersoll, and Ross (1985), the estimate for the expected future spot volatility is given by

$$E(v_{t_s}|v_t) = v_t e^{-\alpha(t_s - t)} + \beta(1 - e^{-\alpha(t_s - t)}). \tag{8}$$

Examination of Equations (7) and (8) indicates that the half-life is determined by setting $e^{-\alpha\tau}$ equal to one-half and solving for τ. Given an estimate for α of 1.931, the expected time for an arbitrary volatility of v to revert halfway to its long-run mean is 90 trading days.

The estimate of α for the true stochastic process can be used to check the adequacy of our specification for the dynamics of the volatility of the crude oil futures market. Since the theoretical first-order autocorrelation for the mean-reverting diffusion given by Equation (2) is $e^{-\alpha dt}$, the first-order autocorrelation for the time series of implied spot volatilities can be checked against the theoretical autocorrelation based on the estimated value of α. Given that our estimate for α is 1.931, the theoretical first-order autocorrelation in the daily series of implied spot volatilities (based on 250 trading days per year) is 0.9923. As noted previously, the estimate for the autocorrelation of the daily series of implied volatilities is 0.95, with an approximate standard error of 0.05. Since the implied spot volatilities measure the true spot volatility of the futures market with error, this estimate of the first-order autocorrelation is a downward-biased estimate of the autocorrelation in the true process. Therefore it is not surprising to find that the estimated first-order autocorrelation is somewhat less than the theoretical autocorrelation based on the parameter estimate of α. In spite of this bias, the first-order autocorrelation observed in the im-

[7] Cox et al. (1985) demonstrate that if $2\alpha\beta/\zeta^2 > 1$, the drift rate is high enough to prevent nonpositive spot volatilities. This condition is satisfied for our parameter estimates ($2\alpha\beta/\zeta^2 = 1.702$).

plied volatility series is remarkably consistent with the autocorrelation implied by the estimate for the elastic force of the mean-reverting diffusion process, suggesting that the mean-reverting diffusion process fits the data reasonably well.

Since the risk associated with the changing volatility of the futures market cannot be diversified away, the market price of risk, ϕ, is required to value options on crude oil futures contracts. The estimate for the market price of risk is 0.25, which enters Equation (3) through the term $\phi\nu$, a measure of the total compensation for the risk attributable to stochastic volatility. To illustrate the importance of volatility risk to the required returns on a futures contract, suppose that the spot volatility of the futures market is equal to the long-run mean of 0.0811. For this case the risk premium associated with volatility risk is 0.25×0.0811, which represents a risk premium of approximately 2.0% per year. Since the average annual returns in the futures and stock markets are similar [e.g., Bodie and Rosansky (1980)], a reasonable guess for the size of the average risk premium would be about 8%. Therefore the compensation for volatility risk appears to be a significant component of the risk premia in the futures market.

4.4 Expected average volatilities and forward volatilities

Following Hull and White (1987) and Merton (1973) it has been common to interpret the implied volatility from the Black–Scholes price of a call option as the expected average of the spot volatility of the underlying security over the time remaining until the expiration of the option.[8] We depart from this tradition in that we use the estimated parameters for the mean-reverting diffusion of the volatility for the crude oil futures market to directly compute estimates of the expected average volatility. Note that while implied volatilities from static volatility models are limited to providing forecasts of the volatility over the remaining life of the option, the estimated parameters from the dynamic model permit us to derive forecasts of volatility over arbitrary future time periods. This feature of the stochastic volatility framework permits greater flexibility in designing empirical tests than is possible using traditional methods for inferring volatility from static volatility option pricing models.

[8] Hull and White (1987) and Merton (1973) discuss the conditions under which the implied volatility from the price of a call option can be interpreted as the expected average of the spot volatility of the underlying security over the remaining life of the option. While Merton examines the case where the spot volatility may be at most a known function of time, Hull and White permit volatility to follow a mean-reverting diffusion process similar to Equation (2). When volatility follows a mean-reverting diffusion, Hull and White show that this interpretation of the implied volatility is correct only when (1) volatility and price shocks are uncorrelated and (2) the market requires no compensation for volatility risk.

Ball and Roma (1994) show that the current spot volatility (ν_0) can be used to estimate the expected average volatility over the next τ periods using the formula

$$E\left(\frac{1}{\tau}\int_0^\tau \nu_t\, dt\right) = \omega(\tau)\nu_0 + (1-\omega(\tau))\beta, \tag{9}$$

where

$$\omega(\tau) = \frac{e^{\alpha\tau}-1}{\alpha\tau e^{\alpha\tau}}.$$

Equation (9) shows that the expected average volatility is a weighted average of the current spot volatility and the long-run mean volatility of the futures market. It is easy to show that as the length of the forecast horizon (τ) increases, the weight on the current spot volatility decreases, reflecting the tendency for the spot volatility to revert to its long-run mean of β.

Given our estimates of α and β, Equation (9) is used to transform each of our daily estimates of the spot volatility into a series of forecasts of the expected average volatility over forecast horizons ranging up to 2 years. The time series of term structures for the expected average volatility of the crude oil futures market is presented in Figure 2, which presents plots of these term structures for both the period prior to the invasion of Kuwait (panel A) and the period immediately following the invasion (panel B). The term structure plots in Figure 2 illustrate the mean-reverting nature of the volatility of the crude oil futures market. For example, note that during the preinvasion period shown in panel A both upward and downward sloping term structures are evident, reflecting the tendency for the spot volatility to fluctuate around its long-run mean during this subperiod. In contrast, panel B shows that the period following the invasion of Kuwait is characterized by downward sloping term structures, consistent with the fact that spot volatility remained well-above the long-run mean during the entire postinvasion period.

The term structure of expected average volatilities obtained from Equation (7) can be used to construct forecasts of future volatility similar to the forward rates implicit in the term structure of interest rates. Given the estimates for the long-run mean volatility and the current spot volatility, the expected average volatility over a distant forecast horizon (having a length of τ_2) can be expressed as a weighted average of the expected average volatility over the near-term forecast horizon (having a length of τ_1) plus a forecast of the expected average volatility for a time interval beginning τ_1 periods in the future and having a length of $\tau_2 - \tau_1$. By comparing the expected average volatilities for two forecast periods ending respectively in τ_1 and τ_2

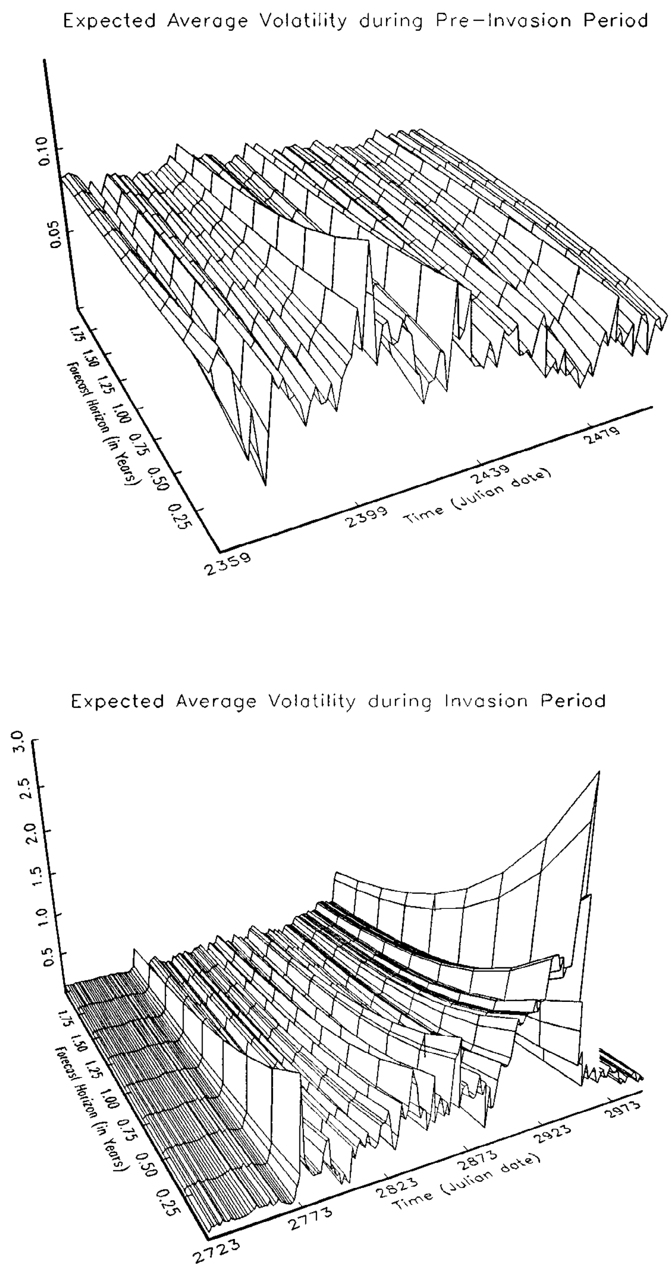

Figure 2
Expected average colatilities from crude oil futures options
This figure plots the term structure of expected average volatilities from crude oil futures options. The top and bottom figures illustrate a time period immediately preceding the invasion period and the invasion period, respectively.

periods, the forward volatility for the period beginning at τ_1 can be expressed as

$$E\left(\left(\frac{1}{\tau_2-\tau_1}\right)\int_{\tau_1}^{\tau_2}\nu_t\,dt\right)=\beta+\frac{1}{\tau_2-\tau_1}(\nu_0-\beta)\frac{e^{\alpha\tau_2}-e^{\alpha\tau_1}}{\alpha e^{\alpha(\tau_2+\tau_1)}}. \quad (10)$$

The forward volatilities defined by Equation (10) are particularly useful in examining the relation between changes in margin requirements and volatility because they provide an ex ante forecast of future volatility that reflects the anticipated decay of shocks to short-run volatility that are coincident to changes in margin requirements.

5. Empirical Tests of the Relation Between Margin Requirements and Volatility

The empirical tests that we use are designed to distinguish between two alternative hypotheses concerning the impact of margin requirements on the volatility of futures markets. According to the first hypothesis, the required margin on a futures contract is a control variable that can be used to dampen excessive volatility. Alternatively, as argued by Day and Lewis (1996) and Fenn and Kupiec (1993), changes in margin requirements are used as a risk management tool by the futures clearinghouse to control for the impact of volatility on the counterparty risk implicit in futures trading. Each of these hypotheses suggests that increases (decreases) in margin requirements should tend to occur following increases (decreases) in volatility. However, under the first hypothesis, the increase (decrease) in margin requirements should reduce (increase) volatility in the period subsequent to the change in margin requirements, while under the second hypothesis the volatility of the return-generating process may be independent of the prevailing margin requirement.

To distinguish between these alternative hypotheses, we examine the changes in the forward volatility of the crude oil futures market during an event window beginning 10 days prior to a change in margin requirements and ending 10 days following the change in margin requirements. If increases in initial margin requirements cause the volatility of the futures market to decline, we should expect to see the forward volatility fall to a lower level during the 10-day period following the implementation of an increase in the initial margin requirement. However, since the volatility of the crude oil futures market follows a mean-reverting stochastic process, a decrease in volatility is consistent with the hypothesis that the volatility of the futures market is independent of margin requirements so long as increases in initial margin requirements tend to follow increases in volatility. Similar logic applies to decreases in margin requirements.

To test the hypothesis that changes in margin requirements have a causative effect on volatility against the null hypothesis that the futures clearinghouse changes margin requirements to control for increases in counterparty risk, we compare futures market volatility for each day in the event window to the volatility 10 days prior to the change in margin requirements. These comparisons allow us to establish whether changes in margin requirements represent a response to changes in expected volatility. They also provide a statistical comparison of volatility before and after changes in margin requirements. To distinguish between margin effects and the impact of mean-reversion in the period subsequent to a change in margin requirements, we then compare the spot volatility for each day following a change in margin requirements with the spot volatility that would have been expected given the volatility on the day of the change in margin requirements.

Note that whereas previous research examines the (ex post) variance of returns in the time periods preceding and subsequent to changes in margin requirements, our estimates of the forward volatility of the crude oil futures market represent ex ante forecasts of the future volatility implicit in the prices of futures call options. The use of an ex ante estimate of the volatility allows the forward volatilities to be interpreted in terms of changes in investors' perceptions concerning the future volatility of the crude oil futures market. While identical results are obtained from comparisons based solely on spot volatilities, the forward volatilities illustrate the impact of the volatility shocks during the event window on the anticipated future volatility of the crude oil futures market.

5.1 Formulation of the null hypothesis

To test the significance of changes in volatility during the periods both prior to and subsequent to the change in margin requirements, we examine changes in the forward volatility of the crude oil futures market during a 21-day event window beginning 10 days prior to a change in margin requirements and ending 10 days following the change. Each change in the forward volatility is measured relative to the forward volatility at the beginning of the event window (day −10). For consistency, the forward volatilities for each day within the event window represent a future 3-month period that is fixed in relation to the first day in the event window.[9]

[9] Since the beginning of the time period over which the forward volatility is defined can be held fixed during the (event) window of time surrounding changes in margin requirements, the forward volatility provides a useful benchmark against which to measure the impact of both exogenous

Under the null hypothesis that there is no change in volatility during the event window surrounding a change in margin requirements, we should be unable to find any systematic differences in the forward volatilities during the period prior to or subsequent to changes in initial margin requirements. Therefore, if we measure changes in forward volatility relative to the average forward volatility at the beginning of the event window (day t^*), the null hypothesis can be stated as

$$H_0\colon \sigma^2_{ft}(t_1, t_2) = \sigma^2_{ft} \cdot (t_1, t_2),$$

where $\sigma^2_{ft} \cdot (t_1, t_2)$ is the forward volatility for some day t within the event window, and $t_2 - t_1$ is always a future time interval having a length of 3 months. For example, if t_1 is 6 months from day t^* ($t_1 = t^* + 0.50$), t_2 is 9 months from day t^* ($t_2 = t^* + 0.75$). Note that for increases in margin requirements, we would expect the forward volatility for the futures market to increase from the beginning of the event window until the date that the initial margin requirement is increased and then either decrease or remain unchanged depending on the impact of the change in initial margin requirement. Therefore, for increases in the margin requirements, the alternative hypothesis is that the difference in the forward volatilities should be significantly greater than zero. Conversely, for decreases in margin requirements, the alternative hypothesis is that the difference in the forward volatilities is significantly less than zero.

We test for the significance of changes in our forecasts of future volatility using two different approaches. First, we perform t-tests for the significance of the average changes in volatilities using an approach similar to Day and Lewis (1988). Since it is possible that the measurement errors associated with the implied spot volatility are heteroskedastic, we also test the null hypothesis using two nonparametric procedures: the Wilcoxon signed-rank test[10] and the Mann–Whitney rank test.[11] Because the forward volatilities are monotone transfor-

changes in the spot volatility and any corresponding changes in margin requirements.

[10] For increases in margin requirements, we test the null hypothesis that there is no significant increase in volatility in either the 10 days prior to or the 10 days following an increase in the margin requirements. Therefore, the Wilcoxon signed-rank test statistic is calculated using the sum of the ranks of the negative differences in the forward volatilities form day t^* to t, which should be less than the sum of the ranks of the positive differences under the alternative hypothesis. The smaller the sum of the negative ranks, the smaller is the probability of rejecting a null hypothesis that is true. This logic is reversed in testing the null hypothesis for the subset of decreases in initial margin requirements.

[11] Whereas the Wilcoxon signed-rank statistic is used to test the null hypothesis that the average difference in the forward volatilities is equal to zero, the Mann–Whitney rank test considers the null hypothesis that the forward volatility on day t comes from the same distribution as the forward volatility on day t^*. For increases in initial margin requirements, the forward volatility on day t would be greater than the forward volatility on day t^* under the alternative hypothesis. This rank test statistic is determined by ranking the forward volatilities and then computing the number of

mations of the spot volatility, the test statistics and related p-values are the same across different forecasting horizons. Therefore we only report one set of values.

We also examine the impact of changes in margin requirements by comparing the average spot volatility for each of the 10 event days following the change in margin requirements with the average expected volatility for that day conditional on the spot volatility at day 0 [see Equation (8)]. Note that the comparison described above explicitly distinguishes between decreases in volatility (relative to day 0) that are attributable to the normal expected decay of above average levels of volatility and decreases that are attributable to the dampening effect of increases in margin requirements.

5.2 Empirical results

The extent and magnitude of changes in the volatility of the crude oil futures market during the periods prior to and subsequent to changes in initial margin requirements provide important evidence concerning the impact of margin requirements on futures market volatility. In this section we present evidence that shows that although increases and decreases in margin requirements are preceded by significant changes in futures market volatility, margin policy has no permanent impact on the volatility of the futures market.

The average changes in forward volatilities relative to the forward volatility at the beginning of the 21-day event window surrounding changes in margin requirements are presented in Table 3. These forward volatilities represent forecasts of future volatility for 3-month periods that respectively begin 3, 6, 9, and 12 months from the first day of the event window (day -10). In panel A we report the average changes in forward volatilities in the event window surrounding the 22 increases in initial margin requirements included in the sample. Note that each of the four forward volatility series reflects a noticeable increase in volatility on the day prior to the increase in initial margin requirements (day -1). For example, the average increase in the 3-month forward volatility in the period prior to the margin change is 0.0307, which represents an increase of approximately 23% relative to the forward volatility at the beginning of the event window. This increase is statistically significant for all forecasting periods. Note that the average increase in forward volatility decreases as the forecasting horizon increases. This reflects both the mean-reverting nature of the square-root diffusion and the tendency for increases in initial margin

forward volatilities for day t that are exceeded by each of the forward volatilities on day t^*. This statistic should be small under the alternative hypothesis that the forward volatility at day t is greater than the forward volatility at day t^*.

Table 3
Average forecasts of volatility during the event window surrounding changes in margin requirements

Average forward volatility estimate	Days relative to initial margin change						
	−10	−5	−1	0	1	5	10
Panel A: Initial margin increases ($n = 22$):							
3-month forward volatility:							
Average forward volatility	0.1331	0.1437	0.1639	0.1711	0.1606	0.1735	0.1397
Volatility difference		0.0106	0.0307	0.0379	0.0274	0.0404	0.0066
6-month forward volatility:							
Average forward volatility	0.1132	0.1197	0.1322	0.1366	0.1302	0.1381	0.1173
Volatility difference		0.0065	0.0190	0.0234	0.0169	0.0249	0.0041
9-month forward volatility:							
Average forward volatility	0.1009	0.1049	0.1126	0.1154	0.1114	0.1163	0.1034
Volatility difference		0.0042	0.0117	0.0144	0.0105	0.0154	0.0025
12-month forward volatility:							
Average forward volatility	0.0933	0.0958	0.1005	0.1022	0.0998	0.1028	0.0948
Volatility difference		0.0025	0.0072	0.0089	0.0065	0.0095	0.0016
t-statistic		1.07	2.05	1.61	1.69	1.26	0.33
Wilcoxon signed-rank test		0.153	0.055	0.143	0.245	0.207	0.863
Mann–Whitney rank test		0.789	0.658	0.450	0.545	0.600	0.816
Panel B: Initial margin decreases ($n = 8$):							
3-month forward volatility:							
Average forward volatility	0.4292	0.2999	0.2824	0.2714	0.2795	0.2747	0.2727
Volatility difference		−0.1293	−0.1468	−0.1580	−0.1497	−0.1545	−0.1566
6-month forward volatility:							
Average forward volatility	0.2959	0.2161	0.2053	0.1985	0.2036	0.2006	0.1993
Volatility difference		−0.0798	−0.0906	−0.0974	−0.0924	−0.0953	−0.0966
9-month forward volatility:							
Average forward volatility	0.2137	0.1644	0.1578	0.1536	0.1567	0.1548	0.1540
Volatility difference		−0.0492	−0.0559	−0.0601	−0.0570	−0.0588	−0.0596
12-month forward volatility:							
Average forward volatility	0.1629	0.1325	0.1284	0.1258	0.1277	0.1266	0.1261
Volatility difference		−0.0304	−0.0345	−0.0371	−0.0352	−0.0363	−0.0368
t-statistic		−1.00	−0.94	−0.99	−0.96	−0.94	−0.88
Wilcoxon signed-rank test		0.547	0.461	0.461	0.461	0.547	0.641
Mann–Whitney rank test		0.796	0.680	0.538	0.643	0.505	0.505

The average forward volatility for a 3-month period starting n months from the first day (−10) of a 21-day event window is reported for selected days within the event window. These forward volatilities, along with the changes relative to the forward volatility at the beginning of the event window, are reported for 3-month periods beginning 3, 6, 9, and 12 months from day −10. Heteroskedasticty-consistent t-statistics are for the average difference between the forward volatilities for day t and day −10. The Wilcoxon signed-rank test reports the p-value for the hypothesis that the day t forward volatility is unaffected by changes in the initial margin requirements. The Wilcoxon–Mann–Whitney test indicates that the null hypothesis that the forward volatilities from day t and day −10 come from the same distribution should not be rejected for any significance level below this value.

requirements to occur when the spot volatility exceeds the long-term volatility of the futures market.

The increases in forward volatilities observed from day −10 to day 0 seem to reverse during the 10-day period following increases in margin requirements. In fact, the results presented in panel A show that we are unable to reject the null hypothesis that the volatility at the end of the event window is on average equal to the volatility at the beginning of the event period. Note that the data fail to reject this null hypothesis under both the parametric and nonparametric tests that we conduct.

The results for the eight decreases in the initial margin requirements are reported in panel B. Note that the average levels for the forward volatilities are uniformly greater than the corresponding forward volatilities reported in panel A. This reflects the fact that several of the decreases in margin requirements occurred in January and February 1991, when the volatility of the oil futures market was declining from unprecedented levels. While it is difficult to make general claims about eight observations, it is of interest to note that the forward volatilities decline in the week prior to the announcement week (days -10 to -6), but remain essentially unchanged during the remainder of the event period (days -5 to 10).

The results presented in Table 3 are consistent with the argument that the clearinghouse increases and decreases initial margin requirements in response to increases and decreases in the volatility of the futures market. Panel A shows that increases in forward volatility during the period prior to the increase in margin requirements seem to reverse themselves in the subsequent 10-day period. A related result is observed in panel B, where we see that the drop in futures market volatility that precedes the reduction in margin requirements does not continue into the second half of the event period.

The interpretation of these results is unclear. Since the average volatility in the period surrounding an increase in margins is at its highest on the announcement date, the decrease in average volatility following the change in margin requirements can be attributed either to the effect of the increase in margins or the tendency for transitory shocks to volatility to decay over time. For decreases in margin requirements, we see that the average volatility on the announcement date is well above the long-run mean of volatility. Therefore, whereas the volatility of the futures market is relatively constant in the period following the announcement date, under the null hypothesis that decreases in margin requirements have no effect on volatility, the spot volatility would be expected to decay in the period following decreases in margin requirements.

To separate the impact of changes in margin requirements from mean-reversion in volatility, we compare the average spot volatility with the corresponding average of the expected future spot volatility (conditional on the day 0 spot volatility) for each day in the event window following changes in margin requirements. These results are

Table 4
Comparison of actual and conditional expected spot volatilities

	Increases in margin requirements			Decreases in margin requirements		
Day	Average spot volatility	Average conditional volatility	*t*-statistic for significance of difference	Average spot volatility	Average conditional volatility	*t*-statistic for significance of difference
1	0.2345	0.2635	−0.27	0.4865	0.4669	0.12
2	0.3324	0.2621	0.49	0.4680	0.4640	0.03
3	0.3310	0.2621	0.37	0.4673	0.4610	0.04
4	0.2792	0.2593	0.17	0.4608	0.4581	0.02
5	0.2699	0.2580	0.10	0.4767	0.4552	0.12
6	0.2645	0.2566	0.08	0.4923	0.4523	0.21
7	0.1842	0.2552	−1.64	0.4955	0.4495	0.24
8	0.1940	0.2539	−1.25	0.4743	0.4466	0.15
9	0.1974	0.2526	−1.14	0.4736	0.4438	0.15
10	0.2009	0.2513	−1.07	0.4724	0.4410	0.16

The averages of the spot volatility and the expectation of the spot volatility conditional on volatility at day 0 are reported for each of the 10 days following a change in margin requirements. To test for significance of the differences between the actual spot volatility and the conditional expectation, we report *t*-statistics computed using the standard deviation of the empirical distribution of actual volatilities for a given event day. Similar nonparametric results for the Wilcoxon rank-sum test are reported in Section 5.

reported in Table 4. For increases in margin requirements, we find that although the average spot volatility appears to decline by more than would be expected given the expected mean-reversion in the announcement date spot volatility, the difference is not statistically significant for any of the 10 days following the change in margin requirements. This result is confirmed by performing a two-sample Wilcoxon rank-sum test for the difference between the actual and the conditional expected spot volatility on day 10. The standard normal approximation for this test statistic has a t-value of -0.49, indicating that the actual volatility at the end of the event window is not statistically different than the volatility that would be expected based on the mean reversion in the spot volatility observed on the announcement date. For decreases in margin requirements, the Wilcoxon rank-sum statistic rejects the null hypothesis that the average spot volatility on day 10 is less than the conditional expected spot volatility. Although the rejection of the null hypothesis implies that decreases in margin requirements may have slowed the forces causing volatility to decay to its long-run mean, the limited size of the sample suggests that this result should be interpreted with caution.

In summary, there is some evidence that ex ante volatility increases temporarily in the period immediately prior to increases in margin requirements. However, we are unable to reject the null hypothesis that the forward volatilities 10 days prior to changes in margin requirements are equal to forward volatilities 10 days after changes in margins. And while decreases in initial margin requirements are associated with decreases in ex ante volatility, the results suggest that volatility decreases lead (in a statistical sense) decreases in initial margin requirements. Taken together, these results indicate that changes in initial margin requirements have very little short-term impact on the forecasts of future volatility embedded in the prices of crude oil futures call options.

6. Tests for Granger Causality

Since the event study approach used in the previous section does not represent an explicit test for a causal relation between margin requirements and futures market volatility, this section provides a formal test of whether initial margin requirements Granger-cause volatility (or vice versa). The causal relation between margin requirements and volatility is examined using monthly time series for the average initial margin requirement and the corresponding volatility of the return on the near-term futures contract. The monthly average initial margin requirement is expressed as a fraction of the value of a crude oil futures contract by dividing the average dollar amount of the initial margin requirement during the month by the average value of the futures contract based on the near-term futures price during the month. We use two proxies for monthly volatility. The first is the (historic) variance of the return on the near-term futures contract. The second is the monthly average implied spot volatility. Note that these definitions of margin requirements and volatility differ from the moving average proxies used by Hardouvelis (1988) and Hsieh and Miller (1990). Since each of the observations in the time series that we create depend only on the events within the month, our results are not affected by the residual autocorrelation created by using overlapping observations.

Table 5 presents a test of whether initial margin requirements Granger-cause futures market volatility. The test is implemented by regressing futures market volatility on lagged values of volatility and the margin requirement for the preceding period. Following Hsieh and Miller (1990), we increase the lags of the dependent variable included in our regressions until the serial correlation in the regression residuals has been eliminated. The results presented in Table 5 show very little evidence that initial margin requirements Granger-cause either measure of futures market volatility. In each case the coefficient of the initial margin requirement fails to be statistically significant, which implies that margins do not Granger-cause volatility.

Since the behavior of the implied forward volatilities is consistent with the hypothesis that margin requirements change in response to an increase in options market volatility, Table 6 presents tests of whether futures market volatility Granger-causes initial margin requirements. For the case where the lagged values of futures market

Table 5
Granger causality tests for the effect of margin requirements on volatility

	Monthly spot volatility	Monthly historic volatility
Number of observations[a]	51	51
β_1[b]	0.6495	4.6160
	(0.9111)	(1.2163)
χ^2-statistic for $\beta_1 = 0$	0.8302	1.4795
(p-value)	(0.3622)	(0.2239)
R^2	0.5321	0.2210
D.W.	1.9763	1.9268
L.M.[c]	1.3658	0.0965
(p-value)	(0.2425)	(0.7560)
m	12	12
$Q(m)$[d]	2.309	8.462
(p-value)	(0.999)	(0.748)

We test for the effect of margins on volatility by regressing the volatility for month t on lagged values for volatility and the average margin requirement in the previous month,

$$\sigma_t^2 = \alpha_0 + \sum_{k=1}^{K} \alpha_{t-k}\sigma_{t-k}^2 + \beta_1 M_{t-1} + \varepsilon_t.$$

[a]The number of observations reflect the number of lagged dependent variables needed to remove the residual autocorrelation. For example, 51 observations reflect two lags.
[b]The t-statistics in parentheses are computed using the Newey and West (1987) correction for heteroskedasticity and serial correlation.
[c]The Breusch (1978) and Godfrey (1978) LaGrange multiplier test for first-order serial correlation is robust to the bias in the Durbin Watson statistic caused by the inclusion of lagged dependent variables in the regression specification.
[d]$Q(m)$ is the Ljung–Box statistic for m lags of the residual autocorrelation function.

volatility are measured by the historic standard deviation of returns, the results show no evidence that volatility Granger-causes margin requirements. This is in contrast to the results obtained by Hsieh and Miller (1990), who find that the moving average representations of stock market volatility Granger-cause margins.

When we measure volatility using the implied spot volatility from the options market, we are unable to reject the hypothesis that monthly average spot volatility Granger-causes initial margin requirements.[12] This finding is hardly surprising in light of the evidence concerning

[12] The Granger causality tests were repeated using daily data. The daily proxies for margin requirements and volatility respectively were the dollar amount of the initial margin requirement divided by the price of the near futures contract and the implied spot volatility. As was the case for the monthly data, the daily tests show that volatility Granger-causes initial margin requirements. For the daily regressions of volatility on margin requirements, the estimated coefficient of the daily margin requirement is statistically significant, but is positive rather than negative as would be expected if margins have a causal effect on volatility in the futures market.

Table 6
Granger causality tests for the effect of volatility on margin requirements

	Monthly spot volatility	Monthly historic volatility
Number of observations[a]	37	51
β_1[b]	0.1136	-0.0125
	(6.3752)	(-0.4004)
χ^2-statistic for $\beta_1 = 0$	40.6428	0.1603
(p-value)	(0.0000)	(0.6889)
R^2	0.7721	0.6377
D.W.	1.6153	1.9126
L.M.[c]	1.7512	0.0508
(p-value)	(0.1857)	(0.8216)
m	9	9
$Q(m)$[d]	5.256	3.411
(p-value)	(.811)	(0.999)

We test for the effect of volatility on margins by regressing the average margin requirement for month t on lagged values of the average margin requirement and the volatility from the previous month,

$$M_t = \alpha_0 + \sum_{k=1}^{K} \alpha_{t-k} M_{t-k} + \beta_1 \sigma_{t-1}^2 + \varepsilon_t.$$

[a]The number of observations reflect the number of lagged dependent variables needed to remove the residual autocorrelation. For example, 37 and 51 observations reflect 16 and 2 lags, respectively.
[b]The t-statistics in parentheses are computed using the Newey and West (1987) correction for heteroskedasticity and serial correlation.
[c]The Breusch (1978) and Godfrey (1978) LaGrange multiplier test for first-order serial correlation is robust to the bias in the Durbin Watson statistic caused by the inclusion of lagged dependent variables in the regression specification.
[d]$Q(m)$ is the Ljung–Box statistic for m lags of the residual autocorrelation function.

the relation between forward volatilities and changes in margin requirements presented in Section 5. However, it is interesting to note that averaging the short-term implied volatilities over a period of 1 month fails to attenuate the leading effect of the average increase of the forward volatilities in the 10-day period preceding increases in initial margin requirements. The difference in the causality results for the alternative proxies for futures market volatility is most likely attributable to the speed with which initial margin requirements are adjusted to reflect changes in the uncertainty surrounding the future supply and demand for crude oil. Whereas the implied spot volatilities from the prices of futures options adjust almost immediately to world events, the past series of returns tends to contain only limited information about the causes of short-run changes in market volatility. Given that officials at the New York Mercantile Exchange use the implied volatility from the options market to monitor volatility of the underlying crude oil futures contract, perhaps it should not be surprising that changes in margin requirements tend to be caused by changes in the implied volatility of the futures markets.

7. Conclusion

This article examines the relation between changes in initial margin requirements and the volatility of the crude oil futures market. Whereas previous studies have examined the relation between margin requirements and ex post volatility of the market in question (e.g., the stock market), this article examines the effect of margin requirements on ex ante measures of volatility obtained from the prices of crude oil futures options.

Our study of the behavior of implied stochastic volatilities in the event window surrounding margin changes provides a number of interesting results. As might be expected, we document a tendency for the ex ante volatility of the crude oil futures market to increase in the 10-day period prior to increases in initial margin requirements. However, we find insignificant differences between ex ante volatility 10 days prior to an increase in margin requirements and ex ante volatility 10 days after. Although these results are consistent with a number of hypotheses concerning the relation between volatility and margin requirements, we perform Granger causality tests that suggest that the chain of causation (in a Granger sense) may run from volatility of the crude oil futures markets to initial margin requirements.

References

Ball, Clifford A., and Antonio Roma, 1994, "Stochastic Volatility Option Pricing," *Journal of Financial and Quantitative Analysis*, 29, 589–608.

Bodie, Zvi, and Victor Rosansky, 1980, "Risk and Return in Commodity Futures," *Financial Analysts Journal*, 36, 27–39.

Breusch, T. S., 1978, "Testing for Autocorrelation in Dynamic Linear Models," *Australian Economics Papers*, 17, 334–355.

Cox, John C., Jonathan E. Ingersoll, Steven A. Ross, 1985, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–408.

Day, Theodore E., and Craig M. Lewis, 1988, "The Behavior of the Volatility Implicit in the Prices of Stock Index Options," *Journal of Financial Economics*, 22, 103–122.

Day, Theodore E., and Craig M. Lewis, 1993, "Forecasting Futures Market Volatility Using Alternative Models of Conditional Volatility," *Journal of Derivatives*, 1, 33–50.

Day, Theodore E., and Craig M. Lewis, 1996, "Margin Adequacy in Futures Markets," unpublished manuscript, Vanderbilt University.

Diz, Fernando, and Thomas J. Finucane, 1992, "The Time Series Properties of Implied Volatility of the S&P 100 Index Options," unpublished manuscript, Syracuse University.

Fenn, George W., and P. Kupiec, 1993, "Prudential Margin Policy in a Futures-Style Settlement System," *Journal of Futures Markets*, 13, 389–408.

Fishe, Raymond P. H., Lawrence G. Goldberg, Thomas F. Gosnell, and Sujata Sinha, 1990, "Margin Requirements in Futures Markets: Their Relationship to Price Volatility," *Journal of Futures Markets*, 10, 541–554.

Flemming, Jeff, 1993, "The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices," unpublished manuscript, Rice University.

Godfrey, L. G., 1978, "Testing Against General Autoregressive and Moving Average Error Models when the Regressors Include Lagged Dependent Variables," *Econometrica*, 46, 1293–1302.

Gourlay, A. R., 1970, "Hopscotch: A Fast Second-Order Partial Differential Equation Solver," *Institute of Mathematics and Applications Journal*, 6, 375–390.

Gourlay, A. R., and S. McKee, 1977, "The Construction of Hopscotch Methods for Parabolic and Elliptic Equations in Two Space Dimensions with a Mixed Derivative," *Journal of Computational and Applied Mathematics*, 3, 201–206.

Granger, Clive W. J., 1969, "Investigating Causal Relations by Econometric Models and Cross-Spectral Methods," *Econometrica*, 37, 424–438.

Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

Hansen, L. P., and K. J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Non-Linear Rational Expectations Models," *Econometrica*, 50, 1269–1286; errata, 52, 267–268.

Hardouvelis, Gikas, 1988, "Margin Requirements and Stock Market Volatility," *Federal Reserve Bank of New York Quarterly Review*, Summer, 80–89.

Heston, Steven L., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–344.

Hsieh, David A., and Merton H. Miller, 1990, "Margin Regulation and Stock Market Volatility," *Journal of Finance*, 45, 3–29.

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

Johnson, Herb, and David Shanno, 1987, "Option Pricing when Variance Is Changing," *Journal of Financial and Quantitative Analysis*, 22, 143–151.

Kupiec, Paul H., 1989, "Futures Margins and Stock Price Volatility: Is There Any Link?," *Journal of Futures Markets*, 13, 677–691.

Kupiec, Paul H., and Steven A. Sharpe, 1990, "Animal Spirits, Margin Requirements, and Stock Price Volatility," *Journal of Finance*, 46, 717–731.

Kuwahara, Hiroto, and Terry A. Marsh, 1992, "The Pricing of Japanese Equity Warrants," *Management Science*, 38, 1610–1641.

Merton, Robert C., 1973, "The Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

Moore, Thomas, 1966, "Stock Market Margin Requirements," *Journal of Political Economy*, 74, 158–167.

Newey, Whitney K., and Kenneth D. West, 1987, "A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703–708.

Officer, Robert R., 1973, "The Variability of the Market Factor of the New York Stock Exchange," *Journal of Business*, 46, 434–453.

Schwert, G. William, 1988, "Business Cycles, Financial Crises and Stock Volatility," unpublished manuscript, University of Rochester.

Seguin, Paul J., 1990, "Stock Volatility and Margin Trading," *Journal of Monetary Economics*, 26, 101–121.

Seguin, Paul J., and Gregg A. Jarrell, 1993, "The Irrelevance of Margin: Evidence from the Crash of '87," *Journal of Finance*, 48, 1457–1474.

Stein, Elias M., and Jeremy C. Stein, 1991, "Stock Price Distributions with Stochastic Volatility: An Analytic Approach," *Review of Financial Studies*, 4, 727–752.

Stein, Jeremy C., 1989, "Overreactions in the Options Market," *Journal of Finance*, 44, 1011–1023.

Wiggins, James B., 1987, "Option Values under Stochastic Volatility: Theory and Empirical Estimates," *Journal of Financial Economics*, 19, 351–372.