Volatility Based Tests for Informational Efficiency on Commodity Options Markets

by

Robert J. Myers, Steven D. Hanson,

Jing-Yi Lai, and Hong Wang

Volatility Based Tests for Informational Efficiency on Commodity Options Markets

Robert J. Myers, Steven D. Hanson, Jing-Yi Lai and Hong Wang*

This paper has two objectives. The first is to develop a simple, computationally tractable procedure for estimating implied GARCH volatilities from commodity options price data. The second is to apply this procedure to elicit implied volatilities from soybean option price data and investigate how well the resulting volatility forecasts predict ex-post "realized" volatilities. We find that filtering option prices through a GARCH option pricing model provides informative forecasts of daily volatilities, but that these forecasts can generally be improved upon using additional information available at the time the options are being priced. The results have implications for forecasting volatility, as well as for the informational efficiency of soybean options markets.

Introduction

This paper has two objectives. First, we construct a simple, computationally tractable procedure for estimating implied GARCH volatilities from commodity options price data. Previous research has used market option prices to estimate an implied GARCH volatility process for the underlying asset (e.g., Engle and Mustafa; Hanson, Myers and Wang). However, existing methods involve complex numerical simulation and optimization routines which are somewhat costly because they are difficult to implement and can take a very long time to converge. Complexity and delay limit the usefulness of these numerical simulation routines in a number of important applications, including using implied volatility estimates from past option price data to get a timely estimate of the fundamental value of current options. Thus, it would be useful to have a simple, computationally tractable procedure for estimating GARCH volatilities implied by observed commodity option prices.

The most well known procedure for estimating implied volatilities is to back them out from the Black-Scholes option pricing formula, treating market option prices as known and average volatility to maturity as the unknown (Black and Scholes; Black; and Chiras and Manaster). While this approach has the advantages of simplicity and tractability, it is inconsistent in the following sense. In order to compute the current implied volatility from the Black-Scholes model the volatility is implicitly assumed deterministic over time. But the resulting implied volatility estimates, together with time-series models using historical data, suggest that volatilities are in fact time-varying and stochastic (Bollerslev; Lamoureux and Lastrapes; Baillie and Myers; Yang and Brorsen). Some researchers have tried to overcome

---

* Robert Myers and Steven Hanson are Associate Professors, Jing-Yi Lai is a Research Assistant and Hong Wang is a Visiting Assistant Professor, all at Michigan State University, East Lansing, MI.

this inconsistency by using implied binomial trees to estimate the underlying binomial probability model assumed to be driving the asset price (e.g. Rubinstein, 1994; Jackwerth and Rubinstein); or by specifying a deterministic volatility function which allows volatility to vary deterministically with the asset price and time (e.g., Dumas, Fleming and Whaley). Others have essentially ignored the evolution of daily volatilities over time and concentrated on using option prices to obtain an estimate of the distribution of the underlying asset price at maturity, conditional on information available when the option price is observed (Fackler and King).

Here, however, we take a different approach. Because GARCH processes have been shown to do a good job of modeling stochastic volatility in commodity futures prices we assume that the market knows the underlying futures price is GARCH. Under this assumption it should be possible to recover estimates of the implied GARCH volatility process from a sequence of option prices. This approach is consistent in the sense that the implied GARCH volatilities are being derived from an option pricing model that assumes the underlying stochastic process is GARCH.

The second objective of the paper is to use the implied GARCH volatility estimates to examine the market efficiency hypothesis for U.S. soybean options markets. By the "market efficiency hypothesis" we mean the proposition that there is no currently available information which could be used to forecast the volatility of an underlying asset better than the forecasts implied by current market option premiums. If this condition fails then the options market is not informationally efficient and, ignoring transaction costs and risk considerations, it should be possible to devise a profitable options trading strategy. As usual in this type of test, options market efficiency can only be tested under the maintained hypothesis that the underlying option pricing model is correct. Thus, the test is really a joint test of options market efficiency and the GARCH option pricing model. Results of the test will indicate how well soybean option market participants are able to use and assimilate information to predict future volatility paths for the underlying soybean futures prices.

GARCH Option Pricing Models

Suppose that commodity futures prices follow a GARCH(1, 1) process:

$$\Delta f_t = \mu + \epsilon_t$$

$$\epsilon_t \,|\, \Omega_{t-1} \sim t(0, h_t, \nu) \qquad (1)$$

$$h_t = \omega + \alpha \epsilon_{t-1}^2 + \beta h_{t-1}$$

where $\Delta f_t = \ln F_t - \ln F_{t-1}$ is the change in the logarithm of futures prices, Ω_{t-1} is a set of information available at $t-1$, and ϵ_t is a random shock which follows a student t distribution with conditional variance h_t and degrees of freedom ν. There is now considerable evidence that (1), or slight variations of (1), do a good job of modeling time-varying volatility in many commodity futures price series (Baillie and Myers; Yang and Brorsen).

The GARCH option pricing model assumes that futures prices follow (1) and options on futures are priced at their discounted expected value:[1]

$$P_t = e^{r_t(t-T)} \int_k^{\infty} [e^{f_t} - K] g(f_T) df_T \tag{2}$$

where r_t is the risk free interest rate at date t, $k = \ln K$ is the logarithm of the strike price K, $f_T = \ln F_T$ is the logarithm of the futures price at maturity, and $g(\cdot)$ is the density of f_T conditional on Ω_t (which includes f_t). The standard Black-Scholes formula assumes that $g(\cdot)$ is normal. Under the GARCH model (1), however, $g(\cdot)$ has no closed form solution and so the expectation in (2) is evaluated by simulating repeated realizations of the GARCH process using Monte Carlo methods. The sample mean of the simulated option values over all realizations is then discounted back to the present at r_t to get an estimate of the current option price.[2]

Now suppose we want to use observed market option prices to estimate an implied GARCH volatility process. Let $\hat{P}_{it}(\omega,\alpha,\beta,h_{t+1})$ denote the price of option i estimated at time t using the GARCH option pricing model conditional on parameters $(\omega,\alpha,\beta,h_{t+1})$. Then if the GARCH option pricing model is correct we would expect $\hat{P}_{it}(\omega,\alpha,\beta,h_{t+1})$ to be close to the observed market option price P_{it}. This suggests that the GARCH volatility process implied by options can be estimated by solving:

$$\min_{\omega,\alpha,\beta,h_{t+1}} \sum_{i=1}^{m} \sum_{j=t-n+1}^{t} \theta_{ij}[P_{ij} - \hat{P}_{ij}(\omega,\alpha,\beta,h_{t+1})]^2 \tag{3}$$

which is a weighted sum of squared deviations between realized market option prices and those estimated via the GARCH option pricing model. The summation i is over put and call options with identical maturities but different strike prices. The summation over j determines the number of previous days included in the analysis (since multiple days will be required to identify information about how volatility evolves over time). The weights θ_{ij} may be set to one (no weights) or each option might be weighted by some measure of its liquidity.

---

[1]See Myers and Hanson for an explanation and justification for using a risk-neutral valuation formula. Equation (2) is for call options but an analogous equation holds for puts.

[2]See Myers and Hanson, and Engle and Mustafa, for more details on this procedure.

While the approach just outlined is clearly feasible in principle there are a number of practical difficulties. First, even with a quadratic loss function the gradients of (3) are highly nonlinear which makes optimization somewhat difficult. Second, any search procedure over the parameter space of (3) requires re-simulating the GARCH model at every iteration because each $\hat{P}_{ij}(\omega, \alpha, \beta, h_{t+1})$ has to be evaluated at each prospective value for $(\omega, \alpha, \beta, h_{t+1})$. Third, to get a good estimate of the implied GARCH volatility process we may have to sum over a fairly large number of days (large n) because it is the variation in option prices over time that provides most of the information about the time-varying volatility process. These three factors combined make it extremely difficult, time-consuming, and costly, to estimate implied volatility processes in this way. Thus, attempts to estimate implied GARCH volatilities from market option prices have, to date, been somewhat limited in scope (see, for example, Hanson, Myers and Wang, and Engle and Mustafa).

An Alternative Approach to Estimating Implied GARCH Volatilities

GARCH models can be used to forecast future volatility paths. Assuming a GARCH(1, 1) model, and using prediction formulas derived by Baillie and Bollerslev, the variance of the log futures price at maturity, conditional on information available at date t is given by:

$$Var(f_T \mid \Omega_t) = \sum_{i=1}^{T-t} E_t(\epsilon_{t+i}^2) \tag{4}$$

$$= \frac{\omega}{1-\rho}\left[T - t - \frac{1 - \rho^{T-t}}{1 - \rho}\right] + \frac{h_{t+1}(1 - \rho^{T-t})}{1 - \rho}$$

where $\rho = \alpha + \beta$ and h_{t+1} is the initial daily volatility at t (the variance of Δf_{t+1} conditional on information available at t). Equation (4) assumes $0 < \rho < 1$ and is a sum of daily volatility forecasts starting at h_{t+1} and assuming future volatility movements will be driven by a GARCH(1, 1) model.

An approximation to the GARCH option pricing model can be obtained by assuming that the distribution of f_T conditional on f_t is approximately normal with variance given by (4). In this case it can be shown that the option pricing equation (2) reverts to Black's model:[3]

[3]See Rubinstein (1976). This formula is for calls but an analogous formula holds for puts.

$$P_t = e^{r_t(t-T)}\left\{F_t C\left(\frac{\ln(F_t/K) + 0.5\,var(f_T\,|\,\Omega_t)}{\sqrt{var(f_T\,|\,\Omega_t)}}\right) - KC\left(\frac{\ln(F_t/K) - 0.5\,var(f_T\,|\,\Omega_t)}{\sqrt{var(f_T\,|\,\Omega_t)}}\right)\right. \tag{5}$$

where $C(\cdot)$ is the cumulative distribution function of the standard normal. Options can then be priced by estimating a GARCH model from historical data and using the resulting estimates to evaluate (4) and (5). This approximation allows for stochastic time-varying volatility but assumes the distribution of f_T conditional of Ω_t is approximately normal, an assumption which is generally violated under the GARCH model. Nevertheless, Myers and Hanson have shown that this approximation to the GARCH option pricing model predicts soybean option prices very well, outperforming Blacks formula using historical 30-day volatilities and, in many cases, the full GARCH pricing model estimated via numerical simulation.

Equations (4) and (5) suggest a very simple way of generating implied GARCH volatilities from market option prices. First, use the standard Black formula (5) to back out an implied estimate for $var(f_T\,|\,\Omega_t)$ from option prices P_{it} observed at t. Second, obtain estimates of the ω and ρ values perceived by the market at t (some alternative ways of doing this will be discussed in a moment). Third, take the implied estimates of $var(f_T\,|\,\Omega_t)$, ω, and ρ and substitute them into (4) to solve for an implied estimate of h_{t+1}. This implied estimate of h_{t+1} represents the markets perception of the current daily volatility of futures price movements, conditional on information available at t. Repeating this procedure over a range of dates provides a sequence of implied daily volatility estimates, h_{t+1}, which are consistent with the GARCH option pricing model and conditional on information available at t. All that remains is to discuss how we might obtain estimates of the market's estimates of ω and ρ at each t. Here we consider two alternative ways of doing this.

Method 1: Estimates Based on GARCH Models of Past Futures Prices

The first method is to use historical data on past futures price to estimate (1), making sure to only include information available at date t when the option is being priced. Then the historical estimates of ω and $\rho = \alpha + \beta$ are used as the market's current perception of these parameters at t. At the next date the latest observation is added, the estimates are updated, and the process is repeated. This requires estimating a new GARCH model at every t but this is fairly straightforward.

Method 2: Estimates Based on IGARCH and Past Option Prices

The second method uses data from past option prices to get the market's implied estimate of ω and ρ. To do this we assume that $\rho = \alpha + \beta = 1$ and concentrate an

Forecast 1: Historical 30-Day Sample Variance

A simple sample variance of the previous 30 days of changes in log futures prices (historical 30-day volatility).

Forecast 2: Implied GARCH Volatility Based on Past Futures Prices

The implied GARCH volatility calculated using ω and ρ values estimated from a GARCH model using historical futures price data available at the time the forecast is being made (method 1 above).

Forecast 3: Implied GARCH Volatility Based on Past Option Prices

The implied GARCH volatility calculated using $\rho = 1$ and an ω estimated from past option prices as discussed previously (method 2 above).

One obvious difficulty in undertaking volatility based tests for informational efficiency is that "true" volatility is unobservable. Thus, there is some question as to how to best measure the actual realized volatility that we are trying to estimate based on option premiums. The measure of realized volatility used in this study is the ex-post, in-sample estimates of h_t obtained from a GARCH model on the underlying futures price using all data up to the maturity date on the option. This represents our best ex-post estimate of what daily volatility was at each date prior to maturity.

The tests are based on three different types of regression. The first is of the form:

(9) $$h_t = \gamma_0 + \gamma_1 \tilde{h}_{it} + \eta_t$$

where h_t is our "realized" GARCH volatility estimated using the full sample, and $\tilde{h}_{it}$ is a forecast of h_t conditional on information available at date $t-1$. The i subscript refers to the three different forecasts outlined above ($i = 1,2,3$). If $\tilde{h}_{it}$ is a good forecast of h_t then we would expect $\hat{\gamma}_0 = 0$, $\hat{\gamma}_1 = 1$, and a high R^2 indicating a large amount of the variation in h_t can be explained by $\tilde{h}_{it}$. Thus, results from estimating (9) will provide a general idea of how well the different volatility forecasts are performing.

The second set of regressions take the form:

(10) $$h_t - \tilde{h}_{it} = X_{t-1}\delta + \xi_t$$

where X_{t-1} is a set of information available at $t-1$. Since the volatility forecast error $h_t - \tilde{h}_{it}$ should be orthogonal to all information available at $t-1$ then informational efficiency would imply $\delta = 0$. Again, we use the alternative volatility forecasts as $\tilde{h}_{it}$. For the information set X_{t-1} we could use alternative forecasts of h_t conditional on information available at t and/or other relevant information, such as futures price realizations at $t-1$ and prior dates.

The final regression is an encompassing equation that includes all three of our volatility forecasts:

$$h_t = \lambda_0 + \lambda_1 \tilde{h}_{1t} + \lambda_2 \tilde{h}_{2t} + \lambda_3 \tilde{h}_{3t} + \zeta_t \tag{11}$$

This essentially provides optimal forecast weights on each of the available forecasts. If the implied volatility forecasts have minimal and insignificant weight then this would imply market inefficiency.

An Application to Soybean Options

To investigate the efficiency of daily volatility forecasts implied by soybean options prices we used daily data on the futures contract maturing in July 1990 obtained from the Chicago Board of Trade data base. To calculate the "realized" ex-post daily volatility at each date over the contract life we estimated a GARCH(1, 1) model using futures price data beginning one year prior to maturity and ending at the maturity date. Results from this full in-sample estimation are reported in table 1 and indicate that time-varying volatility is present, and that it is captured well by a simple GARCH(1, 1) model. The estimated degrees of freedom parameter for the t distribution was found to be very large, indicating that the errors are very close to being conditionally normal. Thus, we assumed conditional normality when estimating the GARCH model. Notice that the α and β estimates in table 1 have relatively high asymptotic t-ratios, and that they sum to 0.998, which implies the process is approximately IGARCH. In-sample estimates of h_t from the model in table 1 are used as our ex-post "realized" volatility for evaluating all volatility forecasts.

To construct volatility forecasts using data available at some date t (prior to maturity) we used the three forecasting methods outlined above. That is, we computed historical 30-day sample variances of past changes in the logarithm of futures prices; implied GARCH volatilities using futures data available at t to estimate the market's perception of α and β; and implied GARCH volatilities assuming IGARCH and using option prices available at t to estimate the market's perception of ω. The implied volatility to maturity at each date (used to estimate the implied GARCH volatilities) was backed out of Black's model (5) using daily trading volume to weight current option prices with different strike prices.

We began making forecasts six months prior to the option's maturity date and stopped making them five days prior to expiration. This is to avoid any problem that might be caused by low liquidity for options traded a long time from maturity, or by maturity effects as the maturity date is approached. Forecasting ability of the three alternative volatility forecasts was investigated by running each of the regression models discussed above.

Results from the first set of regressions on general forecasting performance are provided in table 2. All three forecasts contain useful information about realized volatility but forecast 2 (the implied GARCH volatility assuming IGARCH and using past option prices to evaluate ω) performs best, explaining 73% of the variation in realized daily volatilities. In contrast, forecast 1 (historical 30-day sample variances of past changes in the logarithm of futures prices) explains 69% of the variation, while forecast 3 (the implied GARCH volatility using past futures prices to estimate α and β) explains only 66%. Standard errors and *t*-ratios in table 1 must be interpreted with caution because, under IGARCH, the volatility forecasts and realized volatilities may be cointegrated, and it is well known that standard errors from cointegrating regressions are biased.

For the efficiency tests we examined the orthogonality condition for each forecast error using the other forecasts and past futures price changes as the relevant information set. Futures price changes are used, rather than price levels, because we would expect the forecast errors to be stationary and futures prices to have a unit root. Thus, we want to take first differences of the futures price variable to induce stationary. Under IGARCH we might also expect that the volatility forecasts themselves have unit roots and are cointegrated. Thus, the efficiency test regressions may have a stationary dependent variable (forecast errors) and nonstationary explanatory variables that are cointegrated (the alternative forecasts). In this case the coefficients of interest can be written so that they only depend on stationary variables and Sims, Stock and Watson have shown that conventional estimation and inference methods can be applied.

Results from the efficiency tests are reported in table 3. As expected, forecast 1 (historical 30-day sample variances of past changes in the logarithm of futures prices) is not an efficient forecast. Errors in these forecasts are significantly correlated with the two implied GARCH volatility forecasts, suggesting that the information in the implied GARCH forecasts can be used to improve over forecasts based on historical 30-day sample variances. Furthermore, the Durbin-Watson statistic suggests significant serial correlation in forecast errors which is inconsistent with optimal prediction. These results are expected since we already have considerable evidence suggesting that historical variances do not do a very good job of predicting future volatility.

On the other hand, the implied GARCH volatility forecast 2, which uses GARCH models based on past futures prices to compute the implied volatility, performs much better. In this case none of the information variables are significant at the 5% level, although the historical sample variance (forecast 1) has a *t*-ratio of 1.861 and the Durbin-Watson test still shows evidence of serial correlation. These results suggest that forecast 2 performs much better than forecast 1 but that there still may be information available which could improve forecasting performance. Results for implied volatility forecast 3, which assumes IGARCH and uses past option prices to compute the implied volatility, clearly suggest that it is an inefficient forecast. Both of the other forecasts 1 and 2 contain information which is significantly correlated with the forecast errors from forecast 3, suggesting that using such information would lead to improved volatility forecasts. Furthermore, the Durbin-Watson statistic also suggests serial correlation which is inconsistent with optimal forecasting.

To further investigate the issue of forecasting performance we estimated optimal forecast weights for the three forecasts and report the results in table 4. The estimates suggest that the three forecasts combined can explain about 80% of the total variation in realized volatility, and that most of the weight is given to forecast 3. This should have been expected given that forecast 3 had the highest R^2 in table 1. Nevertheless, the other forecasts do contribute to an improved forecasting performance (R^2 increases by an additional 7 percentage points once forecasts 1 and 2 are added to the model). This is consistent with the efficiency tests suggesting that forecast 3 is not efficient (information from the other forecasts improves forecasting performance). The t values in table 4 should be interpreted with care because, if the underlying volatility process is IGARCH, then the estimation is a cointegrating regression and will generate biased standard errors.

Overall, the results suggest that daily volatility forecasts obtained by filtering option prices through a GARCH option pricing model contain a lot of useful information about future realized volatilities. Nevertheless, it would appear that the implied GARCH volatility forecasts are not fully efficient because additional information available at the time options are priced can be used to improve forecasting performance. There are several possible explanations for this result. First, it could be that the GARCH option pricing model used to filter the implied volatility forecasts is incorrect and any "inefficiency" is simply a result of having the wrong option pricing model. This problem is inherent in all efficiency tests of this type. Second, it could be that the sub-optimal forecasting performance of the implied GARCH volatilities is statistically significant but has small economic value. The question of how the market would value any potential improvement in volatility forecasts has not been addressed here and must await further research. Third, it could be that improved volatility forecasts are obtainable and do have economic value but that, because of transaction costs, imperfect information, or other constraints, are not being incorporated into the market's assessment of volatility. Unfortunately, studies such as the one conducted here shed little light on which of these explanations is correct.

Concluding Comments

This paper has developed a simple, computationally tractable procedure for estimating implied GARCH volatilities from commodity options price data. The method assumes that the distribution of the log futures price at maturity, conditional on information available today, is approximately normal. In contrast to Black's model, however, daily volatility varies stochastically and follows a GARCH process. Estimates of the market's perception of the underlying GARCH process can then be used to elicit the current volatility implied by option prices. We present two ways of getting the market's perception of the underlying GARCH process, one from a simple GARCH model estimated using data available at the time the option is priced and the other assuming IGARCH and using past option prices to estimate the drift parameter ω.

The implied GARCH volatilities were then examined to assess how well they predicted ex-post realized volatilities, as measured by an in-sample GARCH model estimated using all data up to the maturity date. In an application to soybean option prices it was found that filtering market option prices through the GARCH option pricing model provides an informative forecast of realized volatility. Nevertheless, the results also suggest that the implied GARCH volatility forecasts could generally be improved upon by using additional information available at the time the options are being priced. This is a precondition for market inefficiency but does not necessarily imply it because the underlying GARCH option pricing model may be incorrect, or the improved forecasts may have no economic value once transaction costs and/or risk premia are taken into account. Nevertheless, the analysis remains useful because it provides a simple way to elicit implied GARCH volatilities from commodity options price data, and evaluates the extent to which these forecasts are consistent with estimates based on time series analysis of the historical data.

References

Baillie, R.T., and T. Bollerslev. *Prediction in Dynamic Models with Time Dependent Conditional Variances*. Working Paper, Department of Economics, Michigan State University, November 1990.

Baillie, R.T., and R.J. Myers. "Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge." *Applied Econometrics* 6(1991): 109-24.

Black, F. "The Pricing of Commodity Contracts." *J. Finan. Econ.* 3(1976):167-79.

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *J. Pol. Econ.* 81(1973): 637-654.

Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." *J. Econometrics* 31(1986): 307-27.

Chiras, D.P., and S. Manaster. "The Information Content of Option Prices and a Test of Market Efficiency." *J. Finan. Econ.* 6(1978): 213-34.

Dumas, B., J. Fleming, and R.E. Whaley. *Implied Volatility Functions: Empirical Tests.* InResearch Symposium Proceedings. Chicago Board of Trade, 1996: 227-262.

Engle, R.F., and C. Mustafa. "Implied GARCH Models from Options Prices." *J. Econometrics* 52(1992): 289-311.

Fackler, P.L. and R.P. King. "Calibration of Option-Based Probability Assessments in Agricultural Commodity Markets." *Am. J. Agr. Econ.* 72(1990):73-83.