IMPLIED VOLATILITY FORECASTS IN THE GRAINS COMPLEX

DAVID P. SIMON

This article finds that the implied volatilities of corn, soybean, and wheat futures options 4 weeks before option expiration have significant predictive power for the underlying futures contract return volatilities through option expiration from January 1988 through September 1999. These implied volatilities also encompass the information in out-of-sample seasonal Glosten, Jagannathan, and Runkle (GJR; 1993) volatility forecasts. Evidence also demonstrates that when corn-implied volatility rises relative to out-of-sample seasonal GJR volatility forecasts, implied volatility substantially overpredicts realized volatility. However, simulations of trading rules that involve selling corn option straddles when corn-implied volatility is high relative to out-of-sample GJR volatility forecasts indicate that none of the trading rules would have been significantly profitable. This finding suggests that these options are not necessarily overpriced. © 2002 Wiley Periodicals, Inc. Jrl Fut Mark 22:959–981, 2002

The author thanks Joel Sternberg, Jay Sultan, Arthur Wilson, and seminar participants at Brandeis University, the Northeast Business Economics Association, and Financial Management Association Meetings for helpful comments, and James Zeitler and Irina Anganina for excellent programming assistance.
For correspondence, David P. Simon, Bentley College, 175 Forest Street, Waltham, Massachusetts 02452-4705; e-mail: dsimon@bentley.edu

Received April 2001; Accepted March 2002

■ *David P. Simon is an Associate Professor of Finance at Bentley College in Waltham, Massachusetts.*

The Journal of Futures Markets, Vol. 22, No. 10, 959–981 (2002)

Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fut.10042

INTRODUCTION

Market participants commonly view implied volatility as a consensus forecast of realized volatility, despite evidence suggesting that the predictive power of implied volatility is mixed. Day and Lewis (1992) examine S&P 100 index options from 1983 to 1989, and find that, although implied volatility has information content for realized volatility over the following week, implied volatility forecasts do not outperform GARCH volatility forecasts.[1] Canina and Figlewski (1993) examine the forecasting power of the implied volatility of S&P 100 index options for the volatility of index returns through option expiration from 1983 to 1987, and find that these forecasts are biased, and, more strikingly, typically do not have statistically significant forecasting power.[2] However, Christensen and Prabhala (1998) find significant predictive power in S&P 100 index options over a longer time span from 1983–1995, while Fleming (1998) finds significant predictive power in S&P 100 index options from 1985 to 1992. In a related article, Jorion (1995) examines currency futures options traded at the International Money Market at the Chicago Mercantile Exchange, and finds that implied volatility has predictive power for the subsequent volatility of the underlying currency futures returns through option expiration but is biased.

The present article examines the forecasting power of the implied volatility of corn, soybean, and wheat futures options at the Chicago Board of Trade from January 1988 through September 1999. Evidence on the predictive power of implied volatility in the grains markets sheds further light on the general forecasting power of implied volatility.[3] The predictive power of implied volatility in the grains markets is also of interest owing to pronounced volatility swings, which should increase the signal-to-noise ratio of implied volatility and enhance its predictive power.[4] Grains volatility tends to increase during the summer when crops reach critical periods where drought or excessive heat can sharply

[1]Canina and Figlewski (1993) point out that implied volatility should reflect the market's assessment of realized volatility through option expiration and not necessarily over the following week.

[2]In each of their subsamples, implied volatility fails the unbiasedness tests, and in only 6 of 32 subsamples are the coefficients on implied volatility significantly greater than zero at the 5% level.

[3]These markets collectively are referred to as the grains markets for expositional ease, even though soybeans are an oilseed rather than a grain. Wilson and Fung (1990) examine the forecasting power of implied volatilities over the first 2 years of grains option trading from 1985–1986 and find mixed results.

[4]Christensen and Prabhala (1998) argue that one of the reasons that S&P 100 index implied volatility has predictive power from 1983–1995 is that the signal-to-noise ratio of implied volatility rose after the stock market crash because realized volatility rose sharply relative to implied volatility.

curtail yields. Following the harvest, volatility typically moves lower until the following spring. To the extent that the seasonal pattern of implied volatility is well known, implied volatility for options expiring during the growing season should reflect the seasonal pickup in realized volatility.

The present article assesses the predictive power of the implied volatilities of nearby corn, soybean, and wheat options for the volatility of the underlying futures returns from 4 weeks before option expiration through option expiration. The predictive power of implied volatility is compared to the predictive power of out-of-sample volatility forecasts from seasonal GARCH models, specified along the lines of Glosten, Jagannathan, and Runkle (GJR, 1993). Tests also are conducted to determine whether implied volatility encompasses information in GJR volatility forecasts. The article then examines the forecast errors of implied volatility when implied volatility is high relative to GJR volatility forecasts and whether more accurate volatility forecasts could have been obtained by shading implied volatility forecasts toward GJR volatility forecasts. Finally, trading rule simulations are performed in which corn option traders sell an equal dollar value of at-the-money calls and puts (straddles) when implied volatility is high relative to out-of-sample GJR volatility forecasts, combined with other entry and risk management rules.

The results indicate that grains implied volatilities have substantial predictive power for realized volatilities, and that out-of-sample seasonal GJR volatility forecasts are encompassed by implied volatility and do not have significant predictive power when implied volatility is included in the models. However, when corn implied volatility is unusually high relative to out-of-sample GJR volatility forecasts, implied volatility tends to overpredict realized volatility by a substantial extent. Finally, simulations demonstrate that selling corn straddles when implied volatility is high relative to GJR volatility forecasts would not have been significantly profitable for a variety of trading rules, which suggests that these options are not necessarily overpriced.

This article proceeds as follows. Section 1 presents the methodologies used to calculate implied volatility and the seasonal GJR volatility forecasts and discusses and provides background information on the data. Section 2 describes the tests examining the predictive power of implied volatility and presents the empirical results. Section 3 examines the forecast errors of implied volatility when implied volatility is high relative to out-of-sample seasonal GJR volatility forecasts. Section 4 provides trading simulations, and the final section summarizes the findings.

THE VOLATILITY FORECASTS

Implied Volatility

Implied volatilities are calculated from Black's (1976) model for European call and put options on futures contracts,

$$c = [FN(d_1) - KN(d_2)]e^{-rt} \tag{1}$$

$$p = [KN(-d_2) - FN(-d_1)]e^{-rt} \tag{2}$$

where c and p are European call and put option daily settlement prices, F is the settlement price of the futures contract underlying the option, K is the strike price, $N(\cdot)$ is the cumulative normal distribution, $d_1 = \ln(F/K)/\sigma\sqrt{t} + \sigma\sqrt{t}/2$, $d_2 = d_1 - \sigma\sqrt{t}$, t is the time to option expiration, r is the risk-free rate, and σ is the implied standard deviation or the implied volatility of the underlying futures contract. Because option prices and all of the variables in the above formulas are observed except for implied volatility, Equations 1 and 2 each can be solved for implied volatility. Implied volatility is calculated under the assumption that corn, soybean, and wheat options are European options, despite the fact that they are American options, because the value of early exercise is very small for the near term, close to the money options that are examined in this article.[5]

The predictive power of implied volatility is examined 4 weeks prior to the expiration of each option contract to avoid potential biases that could arise from the extreme degree of overlapping forecasts if daily data were used instead.[6] Christensen and Prabhala (1998) point out that the extreme degree of autocorrelation in daily overlapping implied volatility forecast residuals raises the possibility of a spurious regression problem, which causes regression coefficient estimates to be inconsistent.[7] Corn and wheat futures options and futures expire in March, May, July, September, and December, and soybean futures options and futures expire in January, March, May, July, August, September, and November.

[5]The value of early exercise is roughly 1/10 of a percentage point for an at the money corn call with time to expiration and implied volatilities equal to 4 weeks and 22%, respectively.

[6]The analyses were also performed separately at horizons 5 and 6 weeks before option expiration. These results were qualitatively similar to those reported for 4-week horizons.

[7]Christensen and Prabhala (1998) show that in regressions of realized S&P 100 index volatility on implied volatility, the slope coefficients are considerably higher with nonoverlapping data compared to daily overlapping data. These issues are explored further in Hansen, Christensen, and Prabhala (2001).

The options contracts expire on the third Friday of the month prior to futures settlement months, whereas futures trade until the last week of contract months.[8]

Implied volatilities are calculated from the settlement quotes on four options—the closest in-the-money and out-of-the-money calls and the closest in-the-money and out-of-the-money puts on the nearby contract.[9] At-the-money implied volatilities for calls and puts are constructed separately from a weighted linear interpolation of the implied volatilities of the two closest to-the-money options, where the weights are the proximity of the strike price to the underlying futures price. The at-the-money implied volatilities used in this artricle are then formed by averaging these interpolated at the money call and put implied volatilities. The reason for following this procedure is that although these options are actively traded, using four options guards against implied volatility fluctuations owing to the possibility of inaccurate quotes. Also, by interpolating across the options whose strike prices bracket the underlying futures price, the potential distorting effects of the skewness in grains-implied volatilities is mitigated. The skewness reflects a tendency for grains-implied volatilities for both calls and puts to be higher at higher strike prices.[10]

One explanation for the skewness in grains implied volatilities is that the volatility of the underlying futures returns rises more during large rallies than during large selloffs. This phenomenon could reflect the likelihood that factors that cause large rallies, such as droughts and excessive heat, raise uncertainty, and the possibility of sharply lower grains supplies and the need for much higher prices to ration limited supplies. By contrast, selloffs often are associated with the passing of weather scares and a resolution of uncertainty, which causes both realized and implied volatilities to fall. This article later presents evidence consistent with this hypothesis.

[8]For example, the July corn futures contract trades until the last week of July and the July corn futures option contract stops trading on the third Friday in June.

[9]The analyses performed in this article were also conducted with each of the options used to construct implied volatilities and the results were qualitatively similar.

[10]For the whole sample period, the average deviation between the implied volatilities of the closest out-of-the-money calls and the closest in-the-money calls for corn, soybeans, and wheat are 1.09, 1.56, and 0.68 percentage points, respectively. Similarly, the mean difference of the implied volatilities of the closest in-the-money puts and the closest out-of-the-money puts for corn, soybeans, and wheat are 0.95, 1.65, and .53 percentage points, respectively. These differences are all statistically significant at the one percent level using t-tests. The direction of skewness is the opposite of what is found in U.S. stock index options, where both call and put options at lower strike prices have higher implied volatilities.

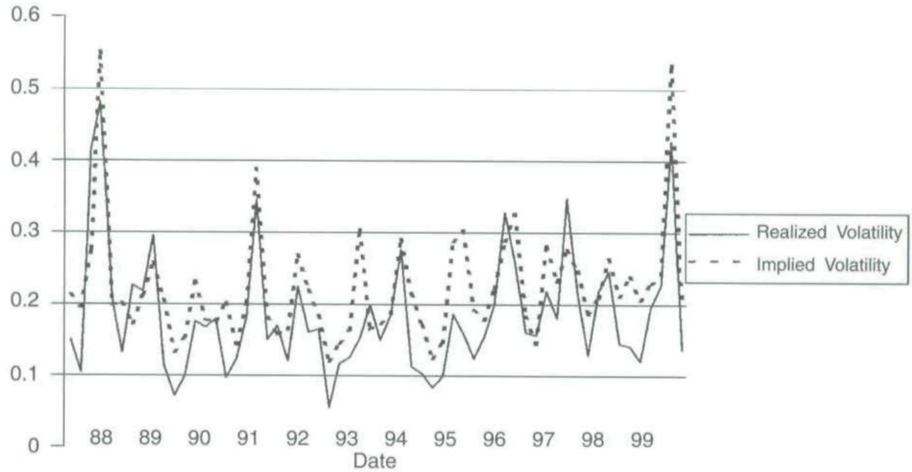

FIGURE 1
Corn volatilities.

The realized daily volatilities of the returns on the underlying futures contract through option expiration are calculated as

$$\sigma_{t,T} = \left(\frac{1}{T-t}\right)\sqrt{\sum_{i=1}^{T-t}(R_{t+i} - \overline{R}_{t,T})^2} \tag{3}$$

where $T - t$ is the number of trading days until option expiration, which is typically 28 days, R_{t+i} is the daily returns calculated as the log of the futures price at $t + i$ minus the log of the futures price at t, and $\overline{R}_{t,T}$ is the average daily return through option expiration from day t. The realized daily volatilities are annualized by multiplying the above volatility by the square root of 252, the typical number of trading days in a year. Corn implied volatility and realized volatility through option expiration over the sample period from January 1988 through September 1999 are shown in Figure 1. The spikes in the chart represent the seasonal pickups in volatility and weather scares.[11]

Seasonal GJR Volatility Forecasts

The next section compares the predictive power of implied volatility to the predictive power of out-of-sample GARCH volatility forecasts. The GARCH forecasts for the volatility of the returns of the futures contracts underlying the nearby option contract are obtained by estimating the

[11]The general patterns of implied volatility and realized volatility for soybeans and wheat are similar to those of corn.

Glosten, Jaganathan, and Runkle (GJR) model, which nests the GARCH model of Bollerslev (1986), and allows conditional volatility to respond differently to equal-magnitude lagged positive and negative return innovations. In the present article, these models are augmented by an intercept indicator variable that captures the seasonal pattern of volatility in the grains markets.[12] The seasonal GJR model is specified as,

$$\begin{gathered} R_t = \mu + \eta_t, \quad \eta_t = N(0, h_t^2), \\ h_t^2 = \alpha_0 + \alpha_1 D^{\mathrm{SEAS}} + \beta h_{t-1}^2 + \delta_1 \eta_{t-1}^2 + \delta_2 D \eta_{t-1}^2 \end{gathered} \tag{4}$$

where R_t is the daily return expressed as the difference from $t - 1$ to t of the natural logs of the futures price, μ is the mean daily return, and η_t is the innovation of the mean daily return. The innovation is assumed to be distributed normally with a mean equal to zero and a conditional variance equal to h_t^2. The variable, D^{SEAS}, is a seasonal indicator variable that takes on a value of 1 if the underlying futures contract settles during the growing season, defined as July through September, and zero otherwise. If the coefficient on the seasonal indicator variable is significantly positive, the conditional volatility of futures return innovations is greater during the growing season, ceterus paribus.[13] The variable D is an indicator variable that takes on the value 1 if the lagged innovation is positive, and zero, otherwise. This specification allows equal magnitude positive and negative squared lagged innovations to have different effects on the conditional variance, which is captured by the δ_2 coefficient. If this coefficient is significantly positive, positive return innovations have a greater effect on conditional volatility than equal magnitude negative returns.

Table I presents estimates of the seasonal GJR models of the daily returns of corn, soybean, and wheat futures contracts underlying the nearby option contracts for the whole sample period from January 1988 through September 1999 and for subperiods from January 1988 through December 1993 and from January 1994 through September 1999. The autoregressive and the moving average terms in the conditional variance equations are highly statistically significant for each commodity for each period, consistent with the presence of conditional heteroskedasticity. The seasonal indicator variable enters the corn model with a significantly

[12]Models were also estimated that allow the coefficients on the lagged squared innovation in the conditional volatility equation to be different during the growing and nongrowing seasons. However, these terms did not enter with significant coefficients and are not reported.

[13]Volatility estimates and realized volatility do not reflect potential discontinuities owing to contract rollovers because futures returns on the first and last day that a contract is followed are calculated with the same contract.

TABLE I
Seasonal GJR Model Estimates
Daily Data from January 1988–September 1999

$$R_t = \mu + \eta_t, \quad \eta_t = N(0, h_t^2)$$
$$h_t^2 = \alpha_0 + \alpha_1 D^{\text{SEAS}} + \beta h_{t-1}^2 + \delta_1 \eta_{t-1}^2 + \delta_2 D \eta_{t-1}^2$$

Commodity	$\mu \times 10^{-3}$	$\alpha_0 \times 10^{-4}$	$\alpha_1 \times 10^{-4}$	β	δ_1	δ_2	*Likelihood Function*
Corn							
1988–1999	−0.108	0.029**	0.0027**	0.9029**	0.0616**	0.0230*	11,929
	(0.184)	(0.005)	(0.009)	(0.0090)	(0.0078)	(0.0102)	
1988–1993	0.026	0.019**	0.030**	0.9150**	0.0532**	0.0226	6142
	(0.244)	(0.005)	(0.012)	(0.0103)	(0.0091)	(0.0123)	
1994–1999	−0.29	0.049**	0.021	0.8862**	0.0642**	0.0315	1471
	(0.29)	(0.012)	(0.021)	(0.0182)	(0.0131)	(0.0185)	
Soybeans							
1988–1999	−0.05	0.024**	0.029**	0.9097**	0.0350**	0.0653**	11,746
	(0.20)	(0.004)	(0.005)	(0.0074)	(0.0064)	(0.0104)	
1988–1993	−0.22	0.019**	0.024**	0.9191**	0.0409**	0.0443*	4962
	(0.32)	(0.007)	(0.008)	(0.0092)	(0.0116)	(0.0183)	
1994–1999	−0.05	0.043**	0.044**	0.8725**	0.0375**	0.0992**	6784
	(0.27)	(0.009)	(0.011)	(0.0166)	(0.0096)	(0.0168)	
Wheat							
1988–1999	−0.37	0.055**	−0.018	0.9043**	0.0535**	0.0330**	2985
	(0.22)	(0.011)	(0.015)	(0.0113)	(0.0084)	(0.0112)	
1988–1993	−0.16	0.059**	−0.019	0.8854**	0.0591**	0.0425*	6054
	(0.28)	(0.013)	(0.019)	(0.0169)	(0.0114)	(0.0112)	
1994–1999	−0.67	0.13**	−0.035	0.8845**	0.0326**	0.0525*	5567
	(0.35)	(0.04)	(0.038)	(0.0297)	(0.0131)	(0.0218)	

Note. R_t is the daily return on the commodity futures, μ is the average daily return, η_t is the innovation from the mean daily return, h_t^2 is the conditional variance of the innovation of daily returns, D^{SEAS} is an indicator variable that equals 1 when the futures contract settles during the growing season between June and September, and zero otherwise, and D is a indicator variable that equals one when the lagged innovation of futures returns is greater than zero, and zero otherwise. Standard errors are in parentheses and one and two asterisks denote statistical significance at the 5 and 1% levels, respectively.

positive coefficient for the whole sample and for the first subsample, indicating a tendency for conditional volatility to be higher during the growing season. The seasonal indicator variable also enters the soybean model with a significantly positive coefficient for the entire period and for both subsamples. By contrast, the seasonal indicator variable does not enter the wheat model with significantly positive coefficients.[14]

[14]Nevertheless, Table II shows that wheat realized volatility is about 3% higher during the growing season than during the nongrowing season. The type of wheat traded at the Chicago Board of Trade is soft red wheat, which is planted in the fall and harvested in mid-summer. The seasonal indicator variable does not enter significantly for wheat when the growing season is assumed to be earlier in the year such as from May through July.

The economic and statistical significance of the conditional volatility asymmetries vary both across commodities and across subsamples. Conditional volatility asymmetries for corn are significantly positive at the 5% level for the whole period but just miss being significant for both subperiods. The conditional volatility asymmetry terms for soybeans and wheat enter with statistically significant positive coefficients for the whole period and for both subperiods. These asymmetries are economically pronounced. The estimates for soybeans for the whole sample period indicate that conditional volatility increases by about three times as much following positive versus negative return innovations. Conditional volatility asymmetries for corn and wheat are present but less pronounced—for the entire period corn and wheat conditional volatilities are about 1/3 and 2/3 higher following positive versus negative return innovations of a given magnitude, respectively.

In the next section, the forecasting performance of implied volatility is compared to the forecasting performance of out-of-sample seasonal GJR volatility forecasts. Out-of-sample GJR forecasts of the volatility of the underlying futures returns from 4 weeks before through the expiration of the nearby option contract are obtained as follows: the seasonal GJR models described above are estimated for each commodity over the presample period from the beginning of option trading for each commodity through the end of 1987 to obtain initial parameter estimates.[15] The estimated parameters from the seasonal GJR models are used to obtain 1-day ahead volatility forecasts in 1988—the first year of the sample—using the following formula,

$$E(h_{t+1}^2) = \alpha_0 + \alpha_1 D^{\text{SEAS}} + \beta h_{t-1}^2 + \delta_1 \eta_{t-1}^2 + \delta_2 D \eta_{t-1}^2 \tag{5}$$

The forecasts of volatility more than 1 day ahead are formed using the following equation,

$$E(h_{t+s}^2) = \alpha_0 + \alpha_1 D^{\text{SEAS}} + (\beta + \delta_1 + 0.5^* \delta_2 \alpha_1) E_t(h_{t+s-1}^2) \quad \text{for } s > 1 \tag{6}$$

Equations 5 and 6 reflect the assumption that positive and negative return innovations have an asymmetric effect on conditional volatility only 1 day into the future. Beyond 1 day in the future, positive and negative return innovations are assumed to be equally likely, and thus are

[15]Corn, soybean, and wheat futures options began trading at the Chicago Board of Trade on February 2, 1985, January 2, 1985, and November 17, 1986, respectively.

invariant to the sign of the most recently observed innovation.[16] The forecast of the standard deviation of returns over the remaining life of the nearby option contract from 4 weeks before expiration is equal to the average of the annualized daily standard deviation forecasts through option expiration. For each subsequent year through 1999, out-of-sample seasonal GJR volatility forecasts are obtained in the same manner by reestimating the models with the prior year of data added to the estimation period.

Preliminary Volatility Analysis

Table II provides background information on the corn, soybean, and wheat volatility data over the sample period from January 1988 through September 1999. The table shows the realized annualized return volatilities (standard deviations) of the underlying futures contract, the implied volatilities and the out-of-sample seasonal GJR volatility forecasts. Implied volatilities are measured 4 weeks before nearby option expiration, while realized volatility and seasonal GJR volatility forecasts are measured from 4 weeks before option expiration through option expiration—the same horizons examined in the next section. The volatilities shown in the table are for the whole period and, separately, for both the growing season from July through September and the nongrowing season. The table shows the means, the standard errors, and the top and bottom deciles of the volatilities. Overall, the realized volatilities tend to be considerably higher during the growing season than during the nongrowing season, and implied volatility tends to be substantially higher than realized volatility. The realized volatilities for corn average 18.5% over the whole period, while the realized volatilities are 24.24% during the growing season and 14.50% during the nongrowing season. The corresponding implied volatilities tend to be about 4% higher. The realized volatilities for soybeans average 20.25% for the whole period and 26.75 and 15.25% during the growing and nongrowing seasons, while the corresponding implied volatilities tend to be between 1 and 3% higher.

[16]If the asymmetry indicator variable is equal to 1 because the most recent return innovation is positive, forecasting more than one step ahead using $\alpha_0 + \alpha_1 D^{SEAS} + (\beta + \delta_1 + \delta_2\alpha_1) E_t(h^2_{t+s-1})$ would be incorrect because doing so would assume a zero probability that future innovations will be negative. By the same reasoning, forecasting more than one step ahead following a negative innovation with the same formula absent the $\delta_2\alpha_1$ term would be incorrect because of the possibility that future innovations are positive. As a result, the asymmetry term is allowed to affect the forecast of the next day's volatility, but it is assumed that positive and negative innovations have the same mean reversion patterns and revert to the same unconditional mean. See Potter (1999) for a discussion of the complexities of forecasting multiple periods ahead with asymmetric GARCH models.

TABLE II

Volatilities Measured 4 Weeks Before Expiration of the Front Option Contract for Corn, Soybeans, and Wheat from January 1988 Through September 1999

	Whole Period			*Growing Season*			*Nongrowing Season*		
	Realized Volatility	*Implied Volatility*	*GJR Volatility*	*Realized Volatility*	*Implied Volatility*	*GJR Volatility*	*Realized Volatility*	*Implied Volatility*	*GJR Volatility*
Corn									
Mean	0.1851	0.2234	0.2028	0.2432	0.2793	0.2514	0.1464	0.1862	0.1703
Std. error	0.0866	0.0822	0.0684	0.1016	0.0963	0.0754	0.0449	0.0416	0.0380
Top decile	0.2999	0.3002	0.2934	0.3926	0.3812	0.3290	0.2109	0.2296	0.2208
Bot decile	0.1026	0.1464	0.1383	0.1477	0.1985	0.1931	0.0978	0.1410	0.1303
Soybeans									
Mean	0.2023	0.2232	0.1983	0.2672	0.2788	0.2383	0.1525	0.1806	0.1676
Std. error	0.0979	0.0870	0.0576	0.1103	0.100	0.0515	0.0447	0.0405	0.0409
Top decile	0.3399	0.3305	0.2640	0.4072	0.3966	0.2794	0.2218	0.2354	0.2206
Bot decile	0.1091	0.1427	0.1286	0.1463	0.1747	0.1845	0.1016	0.1360	0.1230
Wheat									
Mean	0.2113	0.2225	0.2010	0.2339	0.2448	0.2237	0.1963	0.2076	0.1858
Std. error	0.0668	0.0577	0.0316	0.0797	0.0730	0.0306	0.0528	0.0391	0.0219
Top decile	0.3129	0.2909	0.2539	0.3367	0.3175	0.2618	0.2736	0.2503	0.2015
Bot decile	0.1414	0.1746	0.1692	0.1421	0.1774	0.1874	0.1457	0.1712	0.1665

Note. Realized volatilities are the annualized standard deviations of the daily returns of the underlying futures contract from 4 weeks before through option expiration. Implied volatilities are at the money, and are calculated using Black's futures option pricing model for the two closest to the money calls and puts on the front contract 4 weeks before expiration. At the money implied volatilities are constructed by interpolating between the implied volatilities of the calls and the puts and then averaging these interpolated at the money call and put implied volatilities. The GJR volatility forecasts are out of sample forecasts of volatility from 4 weeks before through option expiration from seasonal GJR models. The growing season is assumed to run from July through September.

TABLE III

Root Mean Square Errors of Implied Volatility and Out-of-Sample Seasonal GJR Models for Realized Volatility from 4 Weeks Before Nearby Option Expiration Through Option Expiration from January 1988–September 1999

	Whole Period		*Growing Season*		*Nongrowing Season*	
	Implied Volatility	*GJR*	*Implied Volatility*	*GJR*	*Implied Volatility*	*GJR*
Corn	0.0531	0.0478	0.0611	0.0657	0.0477	0.0419
Soybeans	0.0417	0.0535	0.0464	0.0729	0.0380	0.0386
Wheat	0.0437	0.0482	0.0432	0.0616	0.0440	0.0394

Note. The realized volatilities are the annualized standard deviations of the daily returns of the underlying futures contract from 4 weeks before through option expiration. Implied volatilities are at the money and are calculated using Black's futures option pricing model for the two closest to the money calls and puts on the front contract 4 weeks before expiration. At the money implied volatilities are constructed by interpolating between the implied volatilities of the calls and the puts and then averaging these interpolated at the money call and put implied volatilities. The GJR volatility forecasts are out of sample forecasts of volatility from 4 weeks before through option expiration from seasonal GJR models. The growing season is assumed to run from July through September.

The realized volatilities for wheat average 21% for the whole period and 23.50 and 19.50% during the growing and nongrowing seasons, while the corresponding implied volatilities tend to be about 1% higher. The table also shows that GARCH volatility forecasts tend to be roughly in line with realized volatilities. Finally, the top and bottom deciles, and the standard error of the means demonstrate that all three volatilities are considerably more variable during the growing season than during the rest of the year.

Table III displays the root mean square errors (rmse) of the implied volatility and the out-of-sample seasonal GJR volatility forecasts at the 4-week horizons for the whole period from January 1988 through September 1999 and for the growing and nongrowing seasons. The results indicate that the root mean square forecast errors are fairly large, ranging from about 3.50 to 6 percentage points. Corn implied volatilities have higher rmse than out-of-sample seasonal GJR forecasts, while soybean and wheat implied volatilities have lower rmse. In the next section, the predictive power and the bias of the forecasts are examined.

PREDICTIVE POWER TESTS

This section examines the predictive power of implied volatility and the out-of-sample seasonal GJR volatility forecasts by regressing the realized volatility of the returns of the underlying futures contract on a constant and on implied volatility and the seasonal GJR volatility forecast separately,

$$\sigma_{t,T} = \alpha + \beta\sigma_t^{\mathrm{IV,GJR}} + \varepsilon_{t,T} \tag{7}$$

where $\sigma_{t,T}$ is the realized volatility of returns from 4 weeks before nearby option expiration through option expiration. σ_t^{IV} and σ_t^{GJR} are implied volatility and out-of-sample seasonal GJR forecasts of realized volatility through nearby option expiration, respectively, and are both measured 4 weeks before nearby option expiration.[17] If the volatility forecasts in Equation 7 are unbiased, the intercept coefficient should not be significantly different from zero, and the slope coefficient should not be significantly different from 1. If the volatility forecasts have predictive power, the parameter estimates on the volatility forecasts should be significantly greater than zero.

The models originally were estimated with separate variables for the intercept term and implied volatility during the growing and nongrowing seasons to determine whether the predictive power of implied volatility varies across the year. These results indicate no significant difference between the forecast power of implied volatility and the out-of-sample seasonal GJR forecasts during the growing and nongrowing seasons, and each is statistically significantly positive for all three commodities. As a result, the implied volatilities and the GJR volatility forecasts across the growing and nongrowing seasons can be pooled as in Equation 7.

Whether implied volatility encompasses information about future realized volatility embodied in the GJR volatility forecasts is examined by regressing realized volatility on a constant and on both implied volatility and the out-of-sample seasonal GJR volatility forecasts,

$$\sigma_{t,T} = \alpha + \beta\sigma_t^{\mathrm{IV}} + \lambda\sigma_t^{\mathrm{GJR}} + \varepsilon_{t,T} \tag{8}$$

If implied volatility encompasses the information contained in the GJR volatility forecasts, the coefficients on the GJR forecasts should not be statistically significant and the coefficient on implied volatility should remain significantly positive. All of the regressions are run with the White correction for heteroskedasticity.

The estimates of Equations 7 and 8 are shown in Table IV. The results indicate that the implied volatilities of corn, soybeans, and wheat have significant predictive power at the 1% confidence level. The results

[17]Realized volatility is an estimate of true underlying volatility and hence contains an error component.

TABLE IV

Tests of the Predictive Power of Implied Volatility and Out-of-Sample Seasonal GJR Forecasts from January 1988–September 1999

$$\sigma_{t,T} = \alpha + \beta\sigma_t^{\mathrm{IV,GJR}} + \varepsilon_{t,T}$$

$$\sigma_{t,T} = \alpha + \beta\sigma_t^{\mathrm{IV}} + \lambda\sigma_t^{\mathrm{GJR}} + \varepsilon_{t,T}$$

	Constant	*Implied Volatility*	GJR	$\bar{R}^2$	D-W
Corn	−0.003 (0.012)	0.843[a,b] (0.062)	—	0.64	1.80
	−0.003 (0.019)	—	0.929[a,b] (0.108)	0.53	2.05
	−0.021 (0.013)	0.628[a,b] (0.142)	0.323[b] (0.184)	0.65	2.06
Soybeans	−0.018 (0.014)	0.986[a] (0.069)	—	0.77	1.65
	−0.031 (0.027)	—	1.17[a] (0.157)	0.47	1.62
	−0.018 (0.015)	0.984[a] (0.105)	0.004[b] (0.119)	0.76	1.66
Wheat	0.036 (0.029)	0.786[a] (0.138)	—	0.45	1.82
	0.062 (0.051)	—	0.743[a] (0.262)	0.11	1.63
	0.072 (0.040)	0.889[a] (0.189)	−0.291[b] (0.263)	0.45	1.80

Note. The returns volatility of the underlying futures contract over the remaining 4 weeks of the nearby option contract is regressed on a constant and on either implied volatility, out-of-sample seasonal GJR volatility forecasts or both. Standard errors are estimated using the White correction for heteroskedasticity, and are shown in parentheses.
[a]Indicates significantly different from zero at the 1% level.
[b]Indicates significantly different from one at the 5% level.

also indicate that corn implied volatility is a biased forecast of realized volatility. The estimated 0.84 coefficient on corn implied volatility is significantly less than 1, and the intercept coefficient is not significantly different from zero, indicating that rising implied volatility leads to smaller increases in realized volatility. By contrast, soybean and wheat-implied volatilities are unbiased forecasts of realized volatility as the intercept estimates are not significantly different from zero and the slope estimates are not significantly different from 1.[18] Implied volatility

[18]These results could be affected by the possibility that limits on how much futures prices can change on a given day do not allow futures prices to move as much as they otherwise would move, which would tend to bias downward the GARCH forecasts of volatility relative to implied volatility. However, the above results are robust to throwing out outliers.

explains roughly 50 to 80% of the total variation of the realized volatility of the three commodities. These findings indicate that grains-implied volatilities have substantial forecast power for realized volatility.

The results also indicate that the out-of-sample seasonal GJR forecasts have significant predictive power for all three commodities at the 1% confidence level and are unbiased. A comparison of $\overline{R}^2$s indicates that implied volatility typically explains considerably more of realized volatility than the GJR model forecasts. Finally, when both implied volatility and the GJR model forecasts are included in the model, implied volatility for all three commodities continues to have significant predictive power at the 1% level, while the GJR forecasts no longer have significant predictive power. Thus, implied volatility encompasses the information available about realized volatility contained in out-of-sample time series forecasts.

The results indicate that implied volatility has significant forecast power for realized volatility, explains more of the variation of realized volatility than out-of-sample seasonal GJR volatility forecasts, and encompasses the information contained in the GJR volatility forecasts. However, the possibility remains that when implied volatility is unusually high relative to out-of-sample seasonal GJR volatility forecasts, market participants could forecast realized volatility more accurately by shading implied volatility forecasts toward these time series forecasts. The next section examines this issue.

A CLOSER EXAMINATION OF IMPLIED VOLATILITY FORECASTS

The results in the previous section demonstrate that out-of-sample seasonal GJR volatility forecasts have no significant forecasting power beyond that contained in implied volatility. However, these results do not rule out the possibility that market participants should reduce their volatility forecasts below implied volatility when implied volatility is considerably higher than out-of-sample GJR volatility forecasts. If implied volatility tends to overpredict realized volatility more as implied volatility rises relative to GJR volatility forecasts, GJR volatility forecasts contain useful information, and could possibly help traders to profit from overpriced options.[19] The possibility that implied volatility overpredicts realized volatility by more when implied volatility is high relative to GJR

[19]For example, when implied volatility spikes relative to GARCH volatility forecasts, options traders might profit from selling puts and calls on a delta neutral basis.

volatility forecasts is tested in the following model,

$$\sigma_t^{\mathrm{IV}} - \sigma_{t,T} = \beta_1 + \beta_2(\sigma_t^{\mathrm{IV}} - \sigma_t^{\mathrm{GJR}}) + \varepsilon_{t,T} \tag{9}$$

where the difference between implied volatility and realized volatility from 4 weeks before option expiration through option expiration—the overprediction error of implied volatility—is regressed on a constant and on the deviation of implied volatility from the out-of-sample seasonal GJR volatility forecast. The β_2 coefficient provides information on how much of the deviation of implied volatility from GJR volatility forecasts systematically shows up as overprediction error. The finding of a significantly positive β_2 coefficient would imply that when implied volatility rises relative to GARCH volatility forecasts, implied volatility tends to overpredict subsequent volatility by a greater extent.

Another possibility is that the overprediction error of implied volatility increases when extreme divergences exist between implied volatility and GJR volatility forecasts. To test this possibility, the below equation is also estimated,

$$\sigma_t^{\mathrm{IV}} - \sigma_{t,T} = \beta_1 + \beta_3 D^{\mathrm{DIVERGE}} + \varepsilon_{t,T} \tag{10}$$

where the overprediction error of implied volatility is regressed on a constant and on an indicator variable, D^{DIVERGE}, which takes on the value of 1 when the divergence between implied volatility and the out-of-sample GJR volatility forecasts is in its highest decile over the sample period, and zero otherwise. If the coefficient on this indicator variable is significantly positive, implied volatility tends to overpredict realized volatility by more when implied volatility is high relative to GJR volatility forecasts.[20] For corn, soybeans, and wheat, the upper decile cutoffs are equal to 7.3, 8.0, and 5.2 percentage points, respectively. Both models are estimated with standard errors adjusted for heteroskedasticity using White's method.

The results shown in Table V indicate that only in the case of corn options does the divergence between implied volatility and out-of-sample seasonal GJR forecasts have predictive power for the subsequent overprediction error of implied volatility. The results demonstrate that when implied volatility is equal to GJR volatility forecasts, implied volatility tends to overpredict realized volatility by 3 percentage points, and for each percentage point that implied volatility rises relative to GJR volatility

[20]The equation was also estimated with an indicator variable for large negative divergences between implied volatility and GARCH volatility forecasts. However, these coefficients were not statistically significant.

TABLE V

Tests of Whether Implied Volatility Overprediction Errors Are Driven by the Divergence Between Implied Volatility and Out-of-Sample Seasonal GJR Forecasts and a Dummy Variable for the Top Decile of Divergences (Daily Data from January 1988–September 1999)

$$\sigma_t^{IV} - \sigma_{t,T} = \beta_1 + \beta_2(\sigma_t^{IV} - \sigma_t^{GJR}) + \varepsilon_{t,T}$$

$$\sigma_t^{IV} - \sigma_{t,T} = \beta_1 + \beta_3 D^{DIVERGE} + \varepsilon_{t,T}$$

	β_1	β_2	β_3	$\overline{R}^2$	D-W
Corn	0.031**	0.369**	—	0.10	2.13
	(0.007)	(0.144)			
	0.032**	—	0.061*	0.10	2.17
	(0.007)		(0.024)		
Soybeans	0.020**	0.018	—	−0.01	1.56
	(0.005)	(0.110)			
	0.020**	—	0.009	−0.01	1.58
	(0.005)		(0.021)		
Wheat	0.008	0.144	—	0.00	1.84
	(0.007)	(0.218)			
	0.013*	—	−0.018	0.00	1.80
	(0.006)		(0.031)		

Note. The first model regresses the overprediction error of implied volatility on a constant and on the divergence between implied volatility and out-of-sample seasonal GJR volatility forecasts. The second model regresses the overprediction error of implied volatility on a constant and on an indicator variable that takes on the value one when the divergence between implied volatility and the seasonal GJR forecasts are in their top deciles. For corn, soybeans, and wheat, these cutoffs are equal to 7.3, 8.0, and 5.2%, respectively. The standard errors shown in parentheses are corrected for heteroskedasticity using the White procedure. One and two asterisks denote statistical significance at the 5 and 1% levels, respectively.

forecasts, the overprediction error rises by a statistically significant 4/10 percentage point. The estimated β_2 coefficient of 0.37 implies that optimal forecasts would have 0.63 weight on implied volatility and 0.37 weight on the GJR volatility forecasts (after subtracting 3 percentage points because of the constant).[21]

When the model is specified with an indicator variable that takes on the value of 1 when the divergence between implied volatility and the GJR volatility forecast is in its upper decile, the results demonstrate that corn implied volatility overpredicts realized volatility by an additional statistically significant 6 percentage points. These findings suggest that better forecasts of realized corn volatility would result from moving volatility forecasts toward the GJR forecasts when implied volatility is high relative to those forecasts.

[21]This comes from rewriting Equation 9 as $\sigma_{t,T} = -\beta_1 + (1 - \beta_2)\sigma_t^{IV} + \beta_2\,\sigma_t^{GJR} + \varepsilon_{t,T}$.

For soybeans and wheat, the slope coefficients in Equations 9 and 10 are statistically insignificant, and thus do not support the hypothesis that implied volatility overpredicts realized volatility by substantial amounts when implied volatility rises relative to GJR volatility forecasts. The next section examines the possibility that profitable option trading strategies could have been implemented to exploit the tendency of corn-implied volatility to overpredict realized volatility when corn implied volatility is high relative to out-of-sample GJR volatility forecasts.

TRADING SIMULATIONS

The results in the previous section indicate a tendency for corn implied volatility to overpredict realized volatility by a substantial margin when corn-implied volatility is unusually high relative to out-of-sample seasonal GJR volatility forecasts. Although this finding suggests that corn options are overpriced during these periods, a metric for determining the practical significance of this finding for market participants is whether this tendency leads to profitable trading opportunities. This section simulates trading rules where market participants sell options that appear to be overpriced on a delta-neutral basis.[22] Finding that these trades would not have provided attractive risk-adjusted returns would undermine the view that these options are overpriced.[23]

The simulations involve selling $50 worth of both the nearby, closest to-the-money call and put (straddle) at the close.[24] This assumption keeps the ex-ante risks of each trade similar, and is consistent with a common practice of risking a given fraction of equity on an individual trade. The maximum profit on a short straddle trade is limited to the option premiums received plus the interest earned on the premiums. The seller profits when the call and put prices drop and the short

[22]Noh, Engle, and Kane (1994) examine trading strategies in S&P 500 futures options, in which traders either buy or sell at the money straddles depending on deviations of realized option prices from theoretical option prices based on GARCH volatility forecasts. These authors do not find statistically significant profits after accounting for transactions costs. Harvey and Whaley (1992) examine trading rules in S&P 100 index options based on deviations of realized option prices from option prices forecast on the basis of the mean reverting properties of implied volatility. These authors also do not find statistically significant profits after transactions costs.

[23]Figlewski (1997) discusses the view that implied volatility is determined in light of the expected risks and the transactions costs involved in delta hedging short options positions.

[24]Typically, straddles are equal quantities rather than equal dollar amounts of the closest to-the-money calls and puts. However, given that the closest strike price for calls and puts can be almost 5 cents in or out of the money, equal dollar amounts rather than equal quantities of the closest to-the-money calls are closer to delta neutral. Simulations were also performed, which involve buying straddles when implied volatility is low relative to GARCH volatility forecasts. These trading rules were not significantly profitable.

positions can be bought back at lower prices or the straddle expires worth less than the selling price. Because the simulations assume that straddle sellers are short an equal dollar amount of at the money calls and puts, the deltas of the calls and puts are roughly equal. Consequently, small underlying futures price changes typically have little impact on the value of the straddle. However, if the underlying futures price moves sharply in either direction, the value of the straddle will increase. For example, if corn futures rally sharply, the call price will rise more than the put price falls, and the straddle value will rise, causing the seller to have a loss on the position.[25] The straddle seller typically profits because the underlying futures contract does not move sharply in either direction or because implied volatility falls or because of time decay.

The main issue for options traders is whether the premium collected for selling straddles adequately compensates for the risk of sharp movements in the underlying futures price and for the risk of implied volatility spikes. Not surprisingly, the underlying futures contract typically is volatile when implied volatility is high both relative to GARCH volatility forecasts and in an absolute sense. During such periods, selling straddles might be more profitable when, in addition to implied volatility being high relative to GARCH volatility forecasts, indicators suggest that implied volatility has crested. Also, selling straddles might be unwise when implied volatility is at unusually high levels. To investigate these possibilities, the simulated entry rules incorporate other factors that might enhance the odds of profitable trades. The simulated trades are triggered at the close by combinations of the following rules:[26] (1) 5 and 10% divergences between implied volatility and GJR volatility forecasts; (2) implied volatility at the close is lower than at the previous close; and (3) implied volatility is below 40%.

Because the potential profit from selling straddles is limited to the value of the options sold plus interest earned but potential losses are unlimited, strict risk management rules are important and should be incorporated for realism. The simulations assume exit strategies that are consistent with the limited upside potential and the unlimited downside risk of selling straddles. Traders are assumed to take profits and losses quickly, exiting trades when profits or losses are 50% of the value of the initial short straddle position. Also, straddles are exited if the value of the calls sold increases or decreases 50% relative to the value of the puts

[25]This occurs because as the moneyness of an option increases (decreases), its delta increases (decreases).

[26]Although the 5 and 10% divergences are based on data from the entire sample period, the cutoff points would not be much different based on data observed by traders at a particular time.

sold because the deltas would have shifted substantially, rendering the trade largely a directional bet on futures prices. Finally, trades are exited if the options that the trader is short expire.

The profits from selling straddles are

$$\begin{aligned}\text{Profit} = {} & Q_C{}^*[C_S - C_B - \text{unit trans cost}] \\ & + Q_P{}^*[P_S - P_B - \text{unit trans cost}] \\ & + [Q_C{}^*C_S + Q_P{}^*P_S]^*r^*t/365 \qquad (11)\end{aligned}$$

where Q_C and Q_P are the quantity of calls and puts sold, C_S and P_S are the selling prices of the calls and puts when short straddles are initiated, C_B and P_B are the purchase prices of the calls and puts when straddles are bought back to close positions, unit trans cost is the transactions cost for each option sold, r is the 3-month Treasury bill rate, and t is the number of days that the position is held.

The trading simulations are performed for two types of traders—floor traders and retail traders—who face different transactions costs. The floor trader faces transactions costs owing to the bid–ask spread, whereas retail traders also pay brokerage costs. Following Noh, Engle, and Kane (1994) and Harvey and Whaley (1992), the simulations in this article assume that 1/2 tick is paid both going in and out of each call and put option because the closing option quotes used in this article should range between the bid and the ask at the end of the day. Given a tick size in corn options of 1/8 cent, the round trip transactions cost per contract associated with the bid–ask spread is assumed to be 1/8 cent. Retail traders are also assumed to pay $25 round trip commissions per contract. This corresponds to 1/2 cent per contract because a 1 cent change in the price of a corn option is worth $50.

The results are shown in Table VI. Panel A shows the results for short straddle trades for floor traders, while Panel B shows the results for short straddle trades for retail traders. The entries show for each combination of entry rules, the mean $ profit per $100 of straddles sold, the associated standard error of the mean, the largest winning and losing trade, the number of winning trades and the average number of days that winning trades are held, and the number of losing trades and the average number of days that losing trades are held.

The results indicate that the profits for the various trading strategies for floor traders are economically small—roughly $2 per $100 of straddles sold—and none are statistically significant. The frequency of profitable to unprofitable trades varies across the trading rules but is typically about even. The results also do not indicate much difference in outcomes when

TABLE VI

Simulations of Short Corn Option Straddle Trades (January 1988–September 1999) (Mean profit/std error; Range (high/low); # winners/days held; # losers/days held)

Entry Rules	*ΔIV and IV Are Not Factors*	*ΔIV Is Less Than 0*	*IV Is Less Than 40%*	*ΔIV Is Less Than 0 and IV Is Less Than 40%*
Panel A: Assuming Transactions Costs of a Floor Trader				
Divergence is greater than 5%	2.30 (1.49)	2.20 (2.01)	2.33 (1.66)	2.15 (2.31)
	50/−36	50/−55	50/−36	50/−55
	52/14	32/11	44/16	25/14
	49/12	24/10	43/13	22/14
Divergence is greater than 10%	1.96 (1.72)	2.63 (2.57)	0.29 (2.12)	2.29 (3.74)
	22/−19	22/−19	22/−19	22/−19
	17/7	11/7	10/6	7/7
	15/5	6/6	11/6	3/10
Panel B: Assuming Transactions Costs of a Retail Trader				
Divergence is greater than 5%	−3.11 (1.52)*	−2.45 (2.06)	−3.54 (1.72)	−3.01 (2.37)
	42/−45	42/−63	42/−45	42/−63
	41/16	24/12	34/18	18/17
	60/11	32/9	53/12	19/12
Divergence is greater than 10%	−0.99 (1.77)	−0.13 (2.55)	−3.11 (2.15)	−1.26 (3.68)
	20/−22	21/−24	17/−22	17/−22
	15/7	8/7	8/6	3/6
	17/5	9/6	13/6	7/9

Note. Straddle sales are triggered by combinations such as 5 or 10 percentage point divergences between implied volatility and out of sample GJR volatility forecasts, implied volatility falling from the previous close to the current close and implied volatility at the close being less than 40%. Traders are assumed to sell roughly $50 of calls and puts each. Trades are exited after straddle values have risen or fallen by 50% or the value of calls sold is 50% higher or lower than the value of puts sold. The results shown in Panel A assume the slippage costs that floor traders would face such as bid-ask spreads, while the results in Panel B also assume $25 roundtrip per contract retail brokerage fees are paid. Each entry shows the mean dollar profit along with associated standard error in parentheses, the biggest winning and losing trades, the number of winning trades, the average number of days that winning trades are held, and the number of losing trades and the average number of days that losing trades are held.

the trigger for entering short straddles is raised from a 5 to a 10 percentage point deviation between implied volatility and GJR volatility forecasts or when these rules are augmented by the conditions that implied volatility is below 40% or implied volatility has fallen from the previous close.

The results for retail traders indicate that the trading strategies show losses from $0.13 to $3.54 per $100 of straddles sold. For short straddle trades triggered only by 5 percentage point deviations between implied volatility and GJR volatility forecasts the mean loss of $3.11 is statistically significant. Overall, the simulation results indicate that trading rules aimed at taking advantage of seemingly overpriced options would not have been significantly profitable—even before adjusting for risk.

CONCLUSIONS

This article finds that nearby corn, soybean, and wheat futures implied volatilities have substantial predictive power for realized futures return volatility over the last 4 weeks of the options' lives. These results are in line with Jorion (1995), who finds that implied volatility has substantial predictive power in the foreign exchange market, and with Christensen and Prabhala (1998) and Fleming (1998), who find significant predictive power in S&P 100 index option-implied volatilities. Corn implied volatility provides biased forecasts of realized volatility, with rising implied volatility associated with smaller increases in realized volatility. By contrast, soybean and wheat implied volatility provide unbiased forecasts of realized volatility. The article also demonstrates that grains-implied volatilities encompass the information contained in out-of-sample seasonal GJR volatility forecasts.

The article also finds that corn implied volatility tends to overpredict realized volatility by substantial amounts when corn implied is unusually high relative to out-of-sample GJR volatility forecasts. Simulations of a variety of trading strategies that involve selling corn option straddles when implied volatility is high relative to out-of-sample seasonal GJR volatility forecasts along with other entry rules and risk management rules indicate that such strategies typically would not have been significantly profitable for even floor traders who face low transactions costs. Although corn implied volatility frequently is considerably higher than GJR volatility forecasts and realized volatility, simulations demonstrate that short straddle trading strategies aimed at taking advantage of the apparent overpricing of options would not have been profitable. This finding suggests that these options are not necessarily overpriced.

BIBLIOGRAPHY

Baillie, R., & Bollerslev, T. (1989). The message in daily exchange rates: A conditional-variance tale. Journal of Business and Economic Statistics, 7, 297–306.

Black, F. (1976). The pricing of commodity options. Journal of Financial Economics, 3,167–179.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31, 307–327.

Canina, L., & Figlewski, S. (1993). The information content of implied volatility. Review of Financial Studies, 6, 659–681.

Christensen, B., & Prabhala, N. (1998). The relation between implied and realized volatility. Journal of Financial Economics, 50, 125–150.

Day, T., & Lewis, C. (1992). Stock market volatility and the information content of stock index options. Journal of Econometrics, 52, 267–287.

Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50, 987–1007.

Figlewski, S. (1997). Forecasting volatility, financial markets, institutions and instruments, vol. 6, No. 1. Boston: MA: Blackwell Publishers.

Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. Journal of Empirical Finance, 5, 317–345.

French, D. (1984). The weekend effect and the distribution of stock prices: Implications for option pricing. Journal of Financial Economics, 13, 547–559.

Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance, 48, 1779–1801.

Hansen, C., Christensen, B., & Pbrabhala, N. (2001). Accounting for the overlapping data problem in the implied-realized volatility regression. Center for Analytical Finance Working Paper, University of Aarhus, Aarhus, Denmark.

Harvey, C., & Whaley, R. (1992). Market volatility prediction and the efficiency of the S&P 100 index option market. Journal of Financial Economics, 31, 43–73.

Jorion, P. (1995). Predicting volatility in the foreign exchange market. Journal of Finance, 50, 507–528.

Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. Econometrica, 59, 347–370.

Noh, J., Engle, R., & Kane, A. (1994). Forecasting volatility and option prices of the S&P 500 index. Journal of Derivatives, 2, 17–30.

Potter, S. (1999). Nonlinear time series modeling: an introduction. Journal of Economic Surveys, 13, 505–528.

Wilson, W., & Fung, H. (1990). Information content of volatilities implied by option premiums in grain futures markets. Journal of Futures Markets, 10, 13–27.