

**Intermediate Volatility Forecasts Using Implied Forward Volatility:
The Performance of Selected Agricultural Commodity Options**

Thorsten M. Egelkraut^{a,*}

Philip Garcia^b

*Paper presented at the NCR-134 Conference on Applied Commodity Price
Analysis, Forecasting, and Market Risk Management
St. Louis, Missouri, April 18-19, 2005*

Copyright 2005 by Thorsten M. Egelkraut and Philip Garcia. All rights reserved.
Readers may make verbatim copies of this document for non-commercial purposes by any
means, provided that this copyright notice appears on all such copies.

^{a,b}Thorsten M. Egelkraut and Philip Garcia are Graduate Fellow and Professor and T. A. Hieronymus Chair in the Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign. The authors are also in the Office of Futures and Options Research.

*Corresponding author: egelkrau@uiuc.edu; Tel.: +1-217-333-0644; Fax +1-217-333-5538; Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 326 Mumford Hall, MC-710, 1301 West Gregory Drive, Urbana, Illinois 61801, USA.

Intermediate Volatility Forecasts Using Implied Forward Volatility: The Performance of Selected Agricultural Commodity Options

Options with different maturities can be used to generate an implied forward volatility, a volatility forecast for non-overlapping future time intervals. Using five commodities with varying characteristics, we find that the implied forward volatility dominates forecasts based on historical volatility information, but that the predictive accuracy is affected by the commodity's characteristics. Unbiased and efficient corn and soybeans market forecasts are attributable to the well-established volatility during crucial growing periods. For soybean meal, wheat, and hogs volatility is less predictable, and investors appear to demand a risk premium for bearing volatility risk.

Keywords: Implied forward volatility, Options, Forecasts, Agricultural commodity, Efficiency

Introduction

Options markets are markets in future volatility – each option implies a particular volatility forecast. This forecast, obtained from the observed premium by inverting a theoretical pricing model, is referred to as the implied volatility and commonly interpreted as the expected average volatility until expiration. The implied volatility however is not the only information about future volatility contained in option premiums. The premiums also hold information about the implied forward volatility. The implied forward volatility is generated from two options with consecutive maturities and represents the expected average volatility for the non-overlapping future time interval between their expiration dates. Figure 1 illustrates this concept for a pair of options maturing at T_1 and T_2 , $T_1 < T_2$. At t_0 , implied volatilities for two different intervals can be recovered, $\sigma_{IV(t_0, T_1)}$ and $\sigma_{IV(t_0, T_2)}$. In addition, the option premiums also contain the implied forward volatility, $\sigma_{IFV(T_1, T_2)}$, over the interval T_1 to T_2 .

Options are generally considered to provide the most accurate predictions of future volatility because investors have the ability to incorporate all publicly available information into prices. A large empirical literature has examined this hypothesis for volatility forecasts over nearby time horizons using options with short maturities (for an overview see Poon and Granger, 2003). In contrast, volatility forecasts for more distant horizons have received almost no attention. The lack of research is somewhat surprising since an understanding of the behavior of volatility in the longer term and the forecast accuracy of implied forward volatility are equally important for efficient derivative pricing and effective hedging decisions.

This study evaluates the implied forward volatility as a forecast of subsequent realized volatility for intermediate future time intervals. Specifically, we investigate the predictive performance of implied forward volatility for several important agricultural commodities with different degrees of seasonality in price and production behavior. The choice of these commodities permits us to assess the forecast accuracy of the implied forward volatility in a more comprehensive manner as the nature of production that affects these markets is known.

Moreover, by focusing on just one sector/commodity type we minimize the impact of external factors on the analysis.

The implied forward volatility constitutes an unconventional method to recover volatility forecasts for more distant time intervals from the options market. Our assessment extends previous research on predicting volatility in several important dimensions. First, the limited ability of the traditional time series models to provide accurate volatility forecasts beyond the short term requires exploring alternative forecasting methods (Tomek, 1997; Poon and Granger, 2003). Second, the possibility of using the information contained in options with multiple maturities simultaneously to generate an implied forward volatility has been largely ignored. Yet, Egelkraut et al. (2003) report that the implied forward volatility performs well in predicting the volatility of corn futures prices over various time horizons. Finally, volatility forecasts are typically evaluated relative to their immediate historical volatility as an alternative forecast. For commodities however, this approach may favor the option-based forecasts because it does not adequately account for the potential of commodity-specific patterns of uncertainty resolution. Therefore, we assess the predictive performance of the implied forward volatility against three alternative predictors of volatility – the traditional historical volatility, the volatility realized during the same time interval in the previous year, and a composite forecast that incorporates both recent information and seasonal effects.

Literature

The financial literature has proposed a wide range of statistical forecasting techniques to predict an asset's future volatility. Spurred by the introduction of ARCH (Engle, 1982) and GARCH (Bollerslev, 1986) models, the number of studies on the subject has exploded over the past decade. Numerous empirical studies, primarily of financial markets, tend to confirm that these models can provide powerful predictions of short-term volatility when estimated correctly (Anderson and Bollerslev, 1998; Poon and Granger, 2003). For long-term volatility forecasts however, the traditional time series models and the members of the ARCH and GARCH family may not be appropriate as their long-term volatility forecast will generally revert to the unconditional mean. Day and Lewis (1993), for example, report little explanatory power of GARCH and E-GARCH models in predicting long-term volatility of crude oil futures, and Holt and Moschini (1992) find that ARCH and GARCH models provide poor forecasts of long-term variances in real hog prices.

Instead of predicting future volatility based on a series of past price observations, forecasts based on the volatility implied in options premiums take a different approach. In an efficient options market, the implied volatility is the best available volatility forecast because options premiums impound all information about past volatility as well as expectations about future volatility.¹ If the options-based volatility forecast is obtained using an options pricing model that is linear in volatility and if there is no premium for bearing volatility risk, the implied volatility equals the expected average volatility until expiration. Under this assumption, the difference between two implied volatilities from options maturing in T_1 and T_2 , $T_1 < T_2$, reflects

¹ Poon and Granger (2003) report that the results from 44 out of 53 studies evaluated support the notion that market forecasts contain the most information.

average volatility that market participants expect to prevail during the non-overlapping time interval T_1 to T_2 (Figure 1). This expected average volatility for the non-overlapping time interval is referred to as the implied forward volatility. Since options trade with various maturities, implied forward volatilities can be obtained for various time horizons. Decomposing the expected average volatilities implied in options with different maturities therefore represents a novel approach for obtaining volatility forecasts for intermediate and distant time intervals, where ARCH- and GARCH-type time series models have displayed poor predictive power.

The forecasting performance of the implied forward volatility was examined by Gwilym and Buckle (1997) for one- and two-month maturity American options on the FTSE 100 index from June 1993 to September 1995. Comparing the implied forward volatility between the two expiration dates with the realized volatility over this period, they find that the implied forward volatility consistently overstates realized volatility as evaluated by mean absolute and mean squared errors, and to have poor forecasting ability. Egelkraut et al. (2003) report contrasting results for the forward volatilities implied in corn futures options. Using a substantially larger data set, they examine the market's ability to predict the level of future volatility for intermediate time intervals and to forecast the direction and magnitude of future volatility changes for distant time intervals. Their results indicate that the implied forward volatility predicts future volatility well. For intermediate time intervals, the implied forward volatility provides unbiased forecasts and captures a larger portion of the systematic variability in the realized volatility than forecasts based on historical volatilities. Egelkraut et al. (2003) attribute the difference between the informational content of FTSE 100 index options and corn options to the characteristics of the underlying assets. In contrast to the FTSE 100 index, the volatility of corn futures prices displays strong seasonality across years which is reflected in the implied forward volatilities.

Methods

Implied Forward Volatility

An option's present value is its expected future payoff at maturity discounted at the risk free rate. Hence, the current premiums of European call and put futures options, V_c and V_p , can be written as

$$V_c(x) = b(T) \int_0^{\infty} \max(0, F_T - x) dG(F_T) \quad [1]$$

$$V_p(x) = b(T) \int_0^{\infty} \max(0, x - F_T) dG(F_T) \quad [2]$$

where $b(T)$ is the discount factor, x is the options' strike price, T is the time to expiration, F_T is the price of the underlying futures at maturity, and $G(F_T)$ is the risk neutral valuation measure, i.e. the futures' cumulative distribution function. If $G(F_T)$ is log-normal, these relationships represent Black's (1976) standard formula for European futures options. Estimates of the implied volatility can then be obtained by inverting this pricing model and solving for the standard deviation.

At any moment, there are commonly several implied volatilities for a given maturity because options trade with different strike prices and as calls and puts. Multiple weighing

schemes have been developed to attain a single best implied volatility from the various estimates, but differences in the resulting composite implied volatilities are small. Scott and Tucker (1989), for example, argue that as long as greater weight is placed on at-the-money options, the choice of the weighting scheme is secondary. Because at-the-money options are approximately linear in volatility, and hence most sensitive to changes, all implied volatilities used in this study are obtained from options nearest to being at-the-money. Moreover, these options are the most actively traded and therefore least impacted by noise resulting from wide bid-ask spreads and non-synchronous trading. Possible measurement errors are further reduced by averaging the volatility estimates of the nearest-to-the-money call and put.

Black's (1976) model has been repeatedly questioned. In fact, the formula's underlying assumptions do not hold for most financial markets. Commodity futures, for example, may have return distributions that are not log-normal, and their associated options can typically be exercised any time before expiration rather than only at maturity. If the options are American type rather than European, Black's (1976) implied volatility is upward biased because it does not implicitly embed a premium for the right of early exercise in the options price. This error, however, is small and at a minimum for at-the-money options so that the European pricing formula serves as a good approximation (Ramaswamy and Sundaresan, 1985; Barone-Adesi and Whaley, 1987). There is also some empirical evidence that the distribution of the logarithmic futures returns is not normal but skewed and has leptokurtic tails. The thick-tailed and sometimes non-symmetric return distribution is frequently attributed to be a result of a stochastic volatility process requiring a stochastic volatility model. Despite their less restrictive nature, stochastic volatility models reveal only small biases of Black's (1976) formula, which essentially disappear when at-the-money options are used (Hull and White, 1987; Heston, 1993; Heynen et al., 1994). On the whole, the bias introduced by Black's (1976) formula has been shown to be at most marginal for at- or near-the-money options. When used appropriately the model provides reasonably accurate estimates of the implied volatilities.

Denoting $\sigma_{IV(t_0, T_1)}$ and $\sigma_{IV(t_0, T_2)}$ as the implied volatility estimates expressed in annual terms for the time intervals t_0 to T_1 and t_0 to T_2 , and denoting $D_{(t_0, T_1)}$ and $D_{(t_0, T_2)}$ as the number of trading days between t_0 and T_1 as well as between t_0 and T_2 , the implied forward volatility between the two expiration dates is defined as

$$\sigma_{IFV(T_1, T_2)} = \sqrt{\frac{D_{(t_0, T_2)} \times \sigma_{IV(t_0, T_2)}^2 - D_{(t_0, T_1)} \times \sigma_{IV(t_0, T_1)}^2}{D_{(T_1, T_2)}}} \quad T_2 > T_1, \quad [3]$$

where $D_{(T_1, T_2)}$ refers to the number of trading days between T_1 and T_2 . The implied forward volatility $\sigma_{IFV(T_1, T_2)}$ represents the market's expectation of the average volatility that will occur during this future interval (Figure 1).

This ex-ante volatility forecast can be compared to the ex-post return volatility for the corresponding interval. The realized volatility is based on the futures contract, F , underlying the call and put with the longer time to maturity and is calculated on daily log returns around an assumed mean of zero. Two reasons warrant this approach. First, in an efficient futures market, no arbitrage requires that the mean return from holding futures contracts is zero. Second, Figlewski (1997) cautions that when dealing with short sample periods as is the case in this

study, noisy price movements can result in deviations from the true mean and make its estimate very inaccurate. Expressed in annual terms, the realized volatility during the interval T_1 and T_2 is obtained as

$$\sigma_{REAL(T_1, T_2)} = \sqrt{\frac{\sum_{t=1}^{D_{(T_1, T_2)}} (\ln F_t - \ln F_{t-1})^2}{D_{(T_1, T_2)}}} \times 252. \quad [4]$$

Alternative Volatility Forecasts

The predictive performance of the implied forward volatility is evaluated with respect to alternative predictors of future volatility in order to assess whether market participants incorporate new information into their volatility forecasts. Three alternative forecasts are considered²

- (1) the immediate historical volatility (IHV) defined as the realized volatility during the period immediately preceding the date of the forecast, where – following standard practice - the length of this period is chosen to equal the length of the forecasted interval,
- (2) the one-year lagged historical volatility (LHV) defined as the realized volatility during the same time period as the forward interval in the previous year, and
- (3) the composite historical volatility (CHV) defined as the rolling out-of-sample forecast that is generated by regressing realized volatility on the moving average of the realized volatilities during the same time period as the forward interval in the three previous years and the realized volatility during the time period immediately preceding the forecast date, where the length of this period equals the length of the forecasted interval.³

The immediate historical volatility is the conventional alternative forecast used in most research. Despite its popularity, this approach is not always the most appropriate. When volatility contains seasonal components, as is the case for many agricultural commodities, the immediate historical volatility may provide poor predictions of subsequent realized volatility. Consequently the performance of the proposed forecast that is compared to the immediate historical volatility will be positively biased. To eliminate this potential bias, we offer two additional alternatives. The first which is much in the spirit of a partial adjustment introduces the one-year lagged historical volatility as a second alternative forecast. The second is a composite forecast that allows recent information and seasonal patterns to be incorporated simultaneously by combining information from the immediate historical volatility with a three-year moving average of historical volatility. Here, a three-year time frame is chosen to model

² ARCH- and GARCH-type models were not used in the analysis. ARCH- or GARCH-type models have been found to be successful predictors of short-term volatility when estimated using high frequency (e.g. daily) data. However, these models have little if any predictive power one month or more into the future as the forecasts revert to the unconditional mean. Averaging data over longer time periods or sampling at lower frequency (e.g. monthly) reduces the number of observations and ARCH effect dramatically, making estimation problematic and forecasts unreliable. Early experimentation confirmed these concerns and supports research by Day and Lewis (1993) and Holt and Moschini (1992) who were among the first to report difficulties in forecasting long-term volatility with these models.

³ The composite historical volatility provided moderately better volatility forecasts than the three-year moving average of historical volatility alone. The results of the three-year moving average as a fourth alternative forecast are therefore not presented here.

seasonal effects because it has been shown to be an effective forecast horizon for agricultural crops as it reduces the impact of non-systematic deviations and yet remains rather flexible in adjusting to structural changes in the underlying commodity market.

Forecast Evaluation

Consistent with the literature, the predictive ability of the implied forward volatility is assessed according to three criteria – forecast unbiasedness, superior predictive power, and informational efficiency relative to alternative forecasts. Each of the criteria is stated as testable hypothesis and then explained.

H1: The implied forward volatility is an unbiased forecast of future realized volatility.

Unbiasedness of the implied forward volatility is examined within the following regression framework

$$\sigma_{REAL} = \alpha_0 + \alpha_{IFV} \sigma_{IFV} + \varepsilon \quad [5]$$

where σ_{REAL} and σ_{IFV} refer to the annualized realized and implied forward volatilities. A significant coefficient α_{IFV} indicates that the implied forward volatility contains information about future realized volatility, and a significant constant term α_0 indicates an average level of stochastic volatility that the market is unable to predict. In this context, an unbiased forecast is characterized by $\alpha_0=0$ and $\alpha_{IFV}=1$ which can be tested using a standard F -test. Moreover, if the residuals ε are white noise and independent, the implied forward volatility is efficient.

H2: The implied forward volatility has more predictive power than alternative forecasts of future realized volatility.

The predictive power of the implied forward volatility relative to alternative forecasts of future realized volatility is evaluated by comparing the results from Equation 5 with those obtained for alternative volatility forecasts using

$$\sigma_{REAL} = \alpha_0 + \alpha_{AF} \sigma_{AF} + \varepsilon \quad [6]$$

where σ_{AF} refers to the annualized volatility of a particular alternative forecast (IHV, LHV, or CHV). Greater explanatory power will be reflected in α_0 closer to zero, α_{IFV} closer to one, and a larger adjusted R^2 for the implied forward volatility in Equation 5 than for the alternative forecasts in Equation 6.

The differences in accuracy of the volatility forecasts are further evaluated based on relative forecast errors using mean absolute percentage errors (MAPEs) and mean squared percentage errors (MSPEs)

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{(\sigma_{FORECAST,i} - \sigma_{REAL,i})}{\sigma_{REAL,i}} \times 100 \right| \quad [7]$$

$$MSPE = \frac{1}{n} \sum_{i=1}^n \left(\frac{(\sigma_{FORECAST,i} - \sigma_{REAL,i})}{\sigma_{REAL,i}} \times 100 \right)^2, \quad [8]$$

where $\sigma_{FORECAST}$ refers to the annualized volatility of a particular forecast (IFV, IHV, LHV, or CHV) and where n , the total number of forward intervals, depends on the commodity examined. These error measures are then compared for different forecasts using the Modified Diebold Mariano (*MDM*) test proposed by Harvey, Leybourne, and Newbold, HLN (1997). The procedure involves specifying a cost-of-error function, $g(e)$, of the forecast errors e and testing

pair-wise the null hypothesis of equality of expected forecast performance. The test statistic, which HLN (1997) indicate should be compared with the critical values from the Student's t distribution with $(T - 1)$ degrees of freedom, is computed for one-step ahead forecasts as

$$MDM = \sqrt{\frac{T-1}{\sum_{t=1}^T (d_t - \bar{d})^2}} \bar{d}, \quad [9]$$

where $d_t = g(e_{t,1}) - g(e_{t,2})$, \bar{d} is the average difference across all years, and the null hypothesis is $E(d_t) = 0$. For example, when testing for significant differences of the MAPEs of two forecasts, $g(e_{t,1}) = |e_{t,1}|$ is the absolute percent forecast error of method 1, $g(e_{t,2}) = |e_{t,2}|$ is the absolute percent forecast error of method 2, and $d_t = e_{t,1} - e_{t,2}$ is the difference between the respective absolute percent forecast errors at time t .

HLN (1998) demonstrate that the size of the MDM test is insensitive to contemporaneous correlation between the forecast errors, and that its power declines only marginally with departures from normality. They argue that these characteristics are important since researchers attempting to differentiate between forecasts are often faced with correlated forecasts that possess occasional large errors. Other advantages of the MDM test include its applicability to multiple-step ahead forecast horizons, its non-reliance on an assumption of forecast unbiasedness, and its applicability to cost-of-error functions other than the conventional quadratic loss. HLN (1997) assert that the MDM test constitutes the “best available” method for determining the significance of observed differences in competing forecasts.

H3: The implied forward volatility is informationally efficient, and no alternative forecast of future realized volatility contains additional information that is not already incorporated in the implied forward volatility.

This hypothesis is assessed in the context of the alternative forecasts by incorporating the implied forward volatility and a particular alternative forecast of future realized volatility in the same regression equation

$$\sigma_{REAL} = \alpha_0 + \alpha_{IFV} \sigma_{IFV} + \alpha_{AF} \sigma_{AF} + \varepsilon. \quad [10]$$

Informational efficiency requires that $\alpha_0=0$, $\alpha_{IFV}=1$, $\alpha_{AF}=0$, which can be tested by a standard F-test, and that the residuals ε be independent and distributed as white noise. A non-significant coefficient estimate α_{AF} means that the information provided by the alternative forecast is already contained in the implied forward volatility. If however, the coefficient estimate α_{AF} is significant then the alternative forecast does provide additional information about future volatility not yet contained in the implied forward volatility. Since market participants can incorporate all publicly available information about past prices into their volatility forecasts, the latter case implies that the options market is inefficient and potentially profitable arbitrage opportunities exist.

Data and Construction of Volatility Intervals

Daily closing prices of futures and standard futures options on five agricultural commodities – corn, soybeans, soybean meal, wheat, and hogs – were obtained from the Chicago Board of Trade (CBOT) and the Chicago Mercantile Exchange (CME). The futures data extends from

November 8, 1978, to February 28, 2002, and the options data from January 02, 1992, to December 31, 2001, providing ten complete years of options observations. Since the contract months traded are different for each commodity, the length and number of forward intervals that can be generated from the options first and second in maturity varies (Table 1). The forward intervals are either one, two, or three months long, resulting in a total of 50 intervals for corn and wheat, 70 for soybeans and hogs, and 80 for soybean meal. All intervals are essentially fixed across years because the futures options always mature at approximately the same point in time. The expiration dates vary only by a few days from year to year.

The data are first filtered to exclude uninformative options observations. Such observations include (1) options that are listed but did not actually trade, i.e. zero volume observations, (2) options violating monotonic strike-price patterns, and (3) options with prices less than three times their minimum tick size. The first criterion is used because options prices with no associated trades are simply price quotes. As such, they are not the result of a (negotiation) process in which market participants reach an agreement on their value and form a common volatility expectation. The second criterion removes options that are inconsistent with monotonic strike prices. Call premiums must decrease with increasing strike price and put premiums must increase with increasing strike price. The third criterion avoids possible distortions of the implied volatility calculation introduced by the discrete nature of option prices.

All forward volatilities are derived from options that traded one month before the beginning of every interval, i.e. one month before the expiration of the options with the shorter maturity. Because the forward intervals are one, two, or three months long, this approach assures independent and non-overlapping observations. The computation occurs in two steps. First, the volatility estimates for each of the two option maturities that enter Equation 3 are computed as the arithmetic average of Black's (1976) implied volatilities of the nearest-to-the-money call and the nearest-to-the-money put. The necessary discount factors $b(T)$ are calculated by compounding the corresponding three-month T-Bill rates obtained from the Federal Reserve Board over the time to maturity of the options. Next, the resulting volatility estimates are used to recover the implied forward volatility for the interval between the expiration dates of the two option pairs (Equation 3).

The realized volatilities for the corresponding time intervals as well as the alternative volatility forecasts are computed according to Equation 4. Futures prices from 1978-1992 are used to begin generating the rolling composite forecasts (CHV), which are based on a fixed sample size of twelve years to estimate the most recent parameters. Finally, all volatility measures are expressed in annual terms to allow for comparisons across intervals and years.

Analysis and Results

The results from examining *H1-H3* are reported in Table 2-5.

H1: Informational Content and Unbiasedness of the Implied Forward Volatility

The results from estimating Equation 5 are displayed in Table 2. All slope coefficients are positive and significant indicating that the implied forward volatility contains information about

future realized volatility for each commodity. The α_{IFV} estimates are smaller than one, ranging from 0.468 for wheat to 0.841 for corn. Moreover, the constant terms are significant for soybeans, soybean meal, wheat, and hogs resulting in rejection of the joint hypothesis $\alpha_0=0$ and $\alpha_{IFV}=1$. For those commodities, small values of the implied forward volatility tend to over-predict and large values tend to under-predict future realized volatility. In contrast, α_0 is not significant ($p=0.092$) for corn, and the unbiasedness hypothesis cannot be rejected ($p=0.146$). Further, in light of the absence of autocorrelation in the residuals, the corn market is efficient.

H2: Relative Predictive Power

Corn, soybeans, soybean meal, and wheat: As expected, the immediate historical volatility does not capture the systematic volatility changes associated with crop production, and hence provides the least accurate predictions for corn, soybeans, soybean meal, and wheat. The small adjusted R^2 s and non-significant α_{IHV} estimates in Equation 6 reflect little informational content (Table 3). Relative to the implied forward volatility (Table 2), the immediate historical volatility possesses larger constant terms, smaller slope coefficients, and greater MAPEs and MSPEs, all indicating lower predictive power. Using the *MDM* test, the MAPEs and the MSPEs of each forecast are compared more formally. The error function $g(e)$ is specified as the absolute and the squared percent forecast error and tests for statistical significance in the differences of the MAPEs and the MSPEs between the immediate historical volatility and the implied forward volatility. The p -values displayed in Table 3 show that for both specifications of the error function, these differences are significant.

In contrast to the immediate historical volatility, the one-year lagged historical volatility and the composite forecast do incorporate the volatility patterns associated with crop production. As a result, they possess greater predictive power than the immediate historical volatility (Tables 4 and 5). The composite forecast further outperforms the lagged historical volatility because it incorporates seasonal effects in volatility as well as recent available information. Despite this improvement in accuracy, the implied forward volatility continues to dominate the alternative forecasts for corn, soybeans, and soybean meal. The adjusted R^2 s in Equation 5 (Table 2) remain larger and the MAPEs and MSPEs smaller than those reported for the lagged historical volatility (Table 4) and the composite forecast (Table 5). Yet, the differences between the error measures become less significant when evaluated with the *MDM* test. For wheat, the implied forward volatility provides more accurate forecasts than the lagged historical volatility (Table 4 and Table 2) but somewhat less accurate predictions than the composite forecast (Table 5 and Table 2) as indicated by comparable coefficient estimates and adjusted R^2 s (adjusted $R^2_{IFV}=0.088$ and adjusted $R^2_{CHV}=0.088$) but greater error measures (MAPE_{IFV}=16.802 and MAPE_{CHV}=13.915; MSPE_{IFV}=4.480 and MSPE_{CHV}=3.222).

Hogs: The difference between the implied forward volatility and the immediate historical volatility is less pronounced in hogs (Table 3). Both forecasts predict about equally well (adjusted $R^2_{IFV}=0.182$ and $R^2_{IHV}=0.185$). Though the implied forward volatility has a slightly larger slope coefficient and a smaller constant term than the immediate historical volatility as well as smaller forecast errors (MAPE_{IFV}=22.185 and MAPE_{IHV}=24.614; MSPE_{IFV}=7.568 and MSPE_{IHV}=8.971), the differences in these errors measures are not significant ($p_{MAPE}=0.135$ and

$p_{\text{MSPE}}=0.149$). Since hog production is largely weather independent,⁴ the periods of greater and smaller volatility characteristic for corn, soybeans, soybean meal, and wheat are not present in hogs. Therefore, neither the lagged historical volatility nor the composite forecast possess more predictive power than the immediate historical volatility or the implied forward volatility (Tables 2-5).

H3: Informational Efficiency Relative to Alternative Forecasts

The informational efficiency of the implied forward volatility is also examined by incorporating the implied forward volatility and each of the alternative forecasts in one regression (Equation 10). Similar to results under *H1*, we find only the implied forward volatility from the corn options providing evidence of informational efficiency (joint *F*-test for $\alpha_0=0$, $\alpha_{IFV}=1$, $\alpha_{AF}=0$: IHV $p=0.283$, LHV $p=0.283$, and CHV $p=0.029$; all other commodities and forecasts $p<0.004$). The corresponding adjusted R^2 s reported in Tables 3-5 change only marginally relative to those from Equation 5 in Table 2. Furthermore, the slope coefficients for the implied forward volatility are either equal (wheat) or larger than those for the alternative forecasts. These findings suggest that the implied forward volatility already captures most of the systematic variability in the realized volatility. With the exception of soybeans ($\alpha_{IHV}=-0.303$, $p=0.025$), none of the slope coefficients for the alternative forecasts is significant confirming that the alternative forecasts contain no additional information (Tables 3-5).

Further Analysis

To assess the robustness of the findings, we follow Christensen and Prabhala (1998) and employ an instrumental variable approach. The instrumental variable approach can be useful in the presence of measurement error in the implied forward volatilities which could result in biased coefficients and inappropriate statistical inference. Focusing on the best alternative forecast for each commodity and using lagged implied forward volatility and the respective alternative forecast as instruments, the instrumental variable results do not alter the basic character of our findings; our quantitative tests and qualitative conclusions from examining hypotheses *H1-H3* remain. However, for soybeans, we find that the unbiasedness hypothesis (using lagged implied forward volatility as instrument) and the efficiency hypothesis can no longer be rejected ($p=0.342$; $p=0.366$), suggesting the presence of some measurement error. Hence, for none of the commodities do we find evidence that historical volatilities provide any significant information beyond what is already contained in the implied forward volatility.

Interpretation and Discussion of Differences

The varying degree of forecast accuracy across commodities reflects different levels of difficulty in correctly anticipating when and how much uncertainty will be resolved over time. The finance literature frequently models volatility as a stochastic process around a long-run mean. For commodities, however, the volatility process can contain seasonal components depending on the characteristics of the specific commodity. For example, the realized volatility of corn displays strong seasonality as depicted in Figure 2, whereas for hogs such seasonality is absent (Figure 3). In declining order, the systematic nature of the volatility patterns of soybeans, soybean meal, and wheat (not depicted) are between those of corn and hogs.

⁴ The shift away from traditional farm based hog production began in the 1970s and proceeded at a rapid pace. Today, almost all hogs are raised in confined operations with large, factory-like dimensions (Rhodes, 1995).

The periods of higher and lower corn futures volatility follow the growing and non-growing cycle of the crop. This cycle is particularly pronounced in corn because the plants grow according to an internal clock and cannot generate new growth to compensate for stress during key growth periods. Intervals that contain these short, but critical, periods are therefore characterized by greater volatility than periods where weather has a less profound impact on crop development and future yields. Because the critical periods repeat annually, traders know the approximate times of higher risk and uncertainty and subsequently incorporate the expected greater price volatility into the options premiums. Furthermore, the growing region for corn is geographically limited making it less likely that adverse weather conditions in one area are compensated through favorable environmental factors in another. This combination of the crop's particular temporal and spatial characteristics leads to a concentration of uncertainty resolution over narrow time periods, resulting in more accurate volatility forecasts.

In contrast to corn, soybeans can make up for lost growth during stress periods and are also geographically less concentrated. Unfavorable growing conditions during a particular time or in a certain region have therefore a smaller impact on future yields. As a consequence, soybean price uncertainty is resolved over a wider time window making it more difficult for market participants to anticipate intervals of greater volatility. The volatility of soybean meal follows that of soybeans but the pattern is even less pronounced, and thus more difficult to predict, because meal is only one of several products produced from soybeans and its volatility is impacted by additional supply and demand conditions.

Compared to corn and soybeans, wheat production extends over the largest area in North America. In addition to this spatial element, a temporal dimension exists – deliverable grades for the underlying futures contract include spring and winter wheat – that further reduces the weight of adverse environmental factors. Likewise, timing and geographic location have little influence on price volatility in hogs because production has largely moved towards confined operations (Rhodes, 1995). Because there is little or no concentration of uncertainty resolution, volatility is less predictable and differences in the forecast accuracy between the implied forward volatility and the composite forecast and the immediate historical volatility as the best alternative forecasts for wheat and hogs become less pronounced. Hence, the results in Tables 2-5 are rather consistent with the notion that the predictive performance of the implied forward volatility is influenced by the relative importance of the commodity's temporal and spatial characteristics which lead to different uncertainty resolution over time.

Summary and Conclusion

This study evaluates the implied forward volatility recovered from options with multiple maturities as a forecast of future realized volatility for intermediate time intervals. Using data on five agricultural commodities – corn, soybeans, soybean meal, wheat, and hogs -, the implied forward volatility is derived for one-, two-, and three-month intervals beginning one month into the future from the volatilities implied by options with one to four months to expiration. In addition, three alternative volatility forecasts are generated from the futures prices: the immediate historical volatility, the one-year lagged historical volatility, and a composite forecast that that incorporates both recent information and seasonal effects.

The results for the five commodities indicate that the corn and soybeans implied forward volatilities provide unbiased and efficient forecasts of subsequent realized volatility in futures prices. Soybean meal, wheat, and hogs provide information about realized volatility, but are biased such that small values of the implied forward volatility tend to over-predict and large values tend to under-predict future realized volatility. Examining the predictive performance of the implied forward volatility relative to the three alternative forecasts, we find that for corn, soybeans, soybean meal, and hogs, the option-based forecasts provide either equal or better predictions of future realized volatility than the best alternative forecast based on past volatility information. For wheat, the implied forward volatility dominates the immediate historical volatility and the one-year lagged historical volatility; yet the evidence for the composite forecast is mixed. Further, the historical forecasts do not contain significant information not already incorporated in the options prices. The relative accuracy of the implied forward volatility across commodities is influenced by the importance of each commodity's temporal and spatial characteristics which affects the uncertainty resolution over time. The implied forward volatility displays greater predictive power for commodities where the resolution of uncertainty is concentrated over narrow time periods and spatial production areas.

Although market-based volatility forecasts appear to provide substantial information, the unbiasedness hypothesis is rejected for soybean meal, wheat, and hogs. Biases have also been reported for the traditional implied volatilities of options on financial and non-financial assets (e.g. Jorion, 1995; Szakmary et al., 2003) as well as for the implied forward volatilities of FTSE 100 index options by Gwilym and Buckle (1997). The unbiased nature of the corn and soybean market forecasts may be attributable to the more well established volatility patterns in the realized futures prices. When volatility becomes less predictable, investors may demand a risk premium for bearing volatility risk which could explain the findings in the soybean meal, wheat, and hogs markets.

This study extends previous research on the implied forward volatility of corn by Egelkraut et al. (2003) to four additional agricultural commodities. While consistent with Egelkraut et al. (2003), our results are in contrast to Gwilym and Buckle (1997) who, analyzing FTSE 100 index options, conclude that the implied forward volatility lacks explanatory power. Though we also evaluate two assets (i.e., hogs and wheat) with little or no systematic volatility patterns, we find that the implied forward volatility does possess significant predictive ability regarding future realized volatility. The difference in findings may be attributable to the different nature of the underlying market as well as the shorter time period analyzed by Gwilym and Buckle (1997). Yet, more research on the implied forward volatility is needed for agricultural and financial markets. A better understanding of a potential premium for bearing volatility risk would also be very valuable for academics and practitioners alike.

References

- Anderson, T. G. and T. Bollerslev, 1998. Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review* 39, 885-905.
- Barone-Adesi, G., and R. E. Whaley, 1987. Efficient Analytic Approximation of American Option Values. *Journal of Finance* 42, 301-320.
- Black, F., 1976. The Pricing of Commodity Contracts. *Journal of Financial Economics* 3, 167-179.
- Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31, 307-327.
- Christensen, B. J., and N. R. Prabhala, 1998. The Relation between Implied and Realized Volatility. *Journal of Financial Economics* 50, 125-150.
- Day, T. E., and C. M. Lewis, 1993. Forecasting Futures Market Volatility. *Journal of Derivatives* 1, 33-50.
- Egelkraut, T. M., P. Garcia, and B. J. Sherrick, 2003. The Term Structure of Implied Forward Volatility: Recovery and Informational Content in the Corn Options Market. Proceedings of the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, St. Louis, Missouri, April 21-22, 2003.
http://agecon.lib.umn.edu/cgi-bin/pdf_view.pl?paperid=10915&ftype=.pdf
- Engle, R., 1982. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation. *Econometrica* 50, 987-1008.
- Figlewski, S., 1997. Forecasting Volatility. *Financial Markets. Institutions and Instruments* 6, 1-88.
- Gwilym, O. and M. Buckle, 1997. Forward/Forward Volatilities and the Term Structure of Implied Volatility. *Applied Economics Letters* 4, 325-328.
- Harvey, D., S. Leybourne, and P. Newbold, 1997. Testing the Equality of Prediction Mean Squared Errors. *International Journal of Forecasting* 13, 281-291.
- Harvey, D., S. Leybourne, and P. Newbold, 1998. Tests for Forecast Encompassing. *Journal of Business & Economic Statistics* 16, 254-259.
- Heston, S. L., 1993. A Closed Solution for Options with Stochastic Volatility with Application to Bond and Currency Options. *Review of Financial Studies* 6, 327-43.
- Heynen, R., A. Kemna, and T. Vorst, 1994. Analysis of the Term Structure of Implied Volatilities. *Journal of Financial and Quantitative Analysis* 29, 31-56.
- Holt, M. T., and G. Moschini, 1992. Alternative Measures of Risk in Commodity Supply Models: An Analysis of Sow Farrowing Decisions in the United States. *Journal of Agricultural and Resource Economics* 17, 1-12.
- Hull, J., and A. White, 1987. The Pricing of Options on Assets with Stochastic Volatilities. *Journal of Finance* 42, 281-300.
- Jorion, P., 1995. Predicting Volatility in the Foreign Exchange Market. *Journal of Finance* 50, 507-528.
- Poon, S. H., and C. W. J. Granger, 2003. Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature* 41, 478-539.
- Ramaswamy, K. and S. Sundaresan, 1985. The Valuation of Options on Futures Contracts. *Journal of Finance* 40, 1319-1340.
- Rhodes, V. J., 1995. The Industrialization of Hog Production. *Review of Agricultural Economics* 17, 107-118.

- Scott, E., and A. L. Tucker, 1989. Predicting Currency Return Volatility. *Journal of Banking and Finance* 13, 839-851.
- Szakmary, A., E. Ors., J. K. Kim, and W. N. Davidson III, 2003. The Predictive Power of Implied Volatility: Evidence from 35 Futures Markets. *Journal of Banking and Finance* 27, 2151-2175.
- Tomek, W. G., 1997. Commodity Futures Prices as Forecasts. *Review of Agricultural Economics* 19, 23-44.




Figure 1. Conceptual model of the volatilities $\sigma_{IV(t_0, T_1)}$ and $\sigma_{IV(t_0, T_2)}$ implied by two options maturing at T_1 and T_2 and the implied forward volatility $\sigma_{IFV(T_1, T_2)}$ between these expiration dates, i.e. for the interval T_1 to T_2




Figure 2. Average realized volatilities for forward intervals implied in corn futures options (solid line) and average volatility for sample period (dashed line), 1992-2001




Figure 3. Average realized volatilities for forward intervals implied in hog futures options (solid line) and average volatility for sample period (dashed line), 1992-2001

Table 1. Contracts and forward intervals

Exchange	Commodity	Contract months ^{a,b}	Number of forward intervals			
			One-month	Two-month	Three-month	Total
CBOT	Corn	Z,H,K,N,U		30	20	50
	Soybeans	U,X,F,H,K,N,Q	20	50		70
	Soybean Meal	V,Z,F,H,K,N,Q,U	40	40		80
	Wheat	N,U,Z,H,K,		30	20	50
CME	Hogs ^c	G,J,M,N,Q,V,Z ^d	20	50		70

^aIn addition to the standard contract months, a small number of serial options traded during the data period. These options were not included in the analysis because their irregular occurrence forbids the construction of independent alternative forecasts across years.

^bF=January, G=February, H=March, J=April, K=May, M=June, N=July, Q=August, U=September, V=October, X=November, Z=December

^cThe Dec 96 contract were the last live hog futures and the Feb 97 contract the first lean hog futures traded. This change in contract specification is of no consequence for this study because it focuses on volatility and not price levels.

^dMay options were introduced by the CME only in 2001. Since the data period ends before the first contract, May 02, expires these options are not part of the analysis.

Table 2. Predictive performance of the implied forward volatility (IFV) with respect to future realized volatility, 1992-2001

Commodity	Regression ^{a,b}			F-test	Errors ^c
	α_0 (<i>p</i> -value)	α_1 (<i>p</i> -value)	Adj. R ²	$\alpha_0=0$ and $\alpha_1=1$ <i>p</i> -value	MAPE MSPE
Corn	0.038	0.841	0.507	0.146	18.173
	(0.092)	(0.000)			4.819
Soybeans	0.071	0.670	0.225	0.002	25.008
	(0.031)	(0.002)			12.340
Soybean Meal	0.098	0.603	0.181	0.000	28.025
	(0.001)	(0.001)			14.489
Wheat	0.131	0.468	0.088	0.000	16.802
	(0.001)	(0.021)			4.480
Hogs	0.095	0.667	0.181	0.000	22.185
	(0.001)	(0.001)			7.568

^aFor each commodity, the results are obtained by estimating $\sigma_{REAL} = \alpha_0 + \alpha_{IFV} \sigma_{IFV} + \varepsilon$.

^bIf needed the estimates are adjusted for heteroskedasticity using the Newey-West procedure.

^cMAPE and MSPE are the mean absolute and mean squared percentage errors.

Table 3. Predictive performance of the immediate historical volatility (IHV) with respect to future realized volatility, 1992-2001, and test of forecast encompassing by the implied forward volatility (IFV)

Commodity	Regression ^{a,b}			Adj. R ²	MDM ^c		Errors ^d	
	α_0 (p-value)	α_{IFV} (p-value)	α_{IHV} p-value		p_{MAPE} -value	p_{MSPE} -value	MAPE	MSPE
Corn	0.167 (0.000)		0.143 (0.351)	-0.002	0.003	33.693	21.766	
	0.038 (0.175)	0.840 (0.000)	0.004 (0.970)	0.497	0.000			
Soybeans	0.205 (0.000)		-0.074 (0.604)	-0.011	0.001	36.913	25.157	
	0.105 (0.003)	0.791 (0.000)	-0.303 (0.025)	0.274	0.030			
Soybean Meal	0.159 (0.000)		0.193 (0.057)	0.038	0.009	38.299	29.949	
	0.094 (0.002)	0.577 (0.003)	0.039 (0.675)	0.172	0.029			
Wheat	0.193 (0.000)		0.128 (0.153)	-0.001	0.015	21.981	6.966	
	0.132 (0.001)	0.509 (0.035)	-0.048 (0.751)	0.071	0.005			
Hogs	0.105 (0.005)		0.613 (0.009)	0.185	0.135	24.614	8.971	
	0.079 (0.026)	0.385 (0.029)	0.369 (0.165)	0.207	0.149			

^aFor each commodity, the first regression results are obtained by estimating

$$\sigma_{REAL} = \alpha_0 + \alpha_{IHV}\sigma_{IHV} + \varepsilon$$

$$\sigma_{REAL} = \alpha_0 + \alpha_{IFV}\sigma_{IFV} + \alpha_{IHV}\sigma_{IHV} + \varepsilon.$$

^bIf needed, the estimates are adjusted for heteroskedasticity using the Newey-West procedure.

^cMDM test for statistical significance in the differences of the mean absolute percentage errors and the mean squared percentage errors between the implied forward volatility and the forecast based on the realized volatility during the time period immediately preceding the date of the forecast where the length of this period equals the length of the forecasted interval.

^dMAPE and MSPE are the mean absolute and mean squared percentage errors.

Table 4. Predictive performance of the lagged historical volatility (LHV) with respect to future realized volatility, 1992-2001, and test of forecast encompassing by the implied forward volatility (IFV)

Commodity	Regression ^{a,b}			Adj. R ²	MDM ^c		Errors ^d	
	α_0 (<i>p</i> -value)	α_{IFV} <i>p</i> -value	α_{LHV} <i>p</i> -value		p_{MAPE} -value p_{MSPE} -value	MAPE MSPE		
Corn	0.112 (0.000)		0.422 (0.002)	0.172	0.000 0.000	30.474 13.786		
	0.041 (0.082)	0.892 (0.000)	-0.064 (0.625)	0.500				
Soybeans	0.158 (0.000)		0.172 (0.174)	0.020	0.001 0.011	39.838 31.712		
	0.075 (0.024)	0.714 (0.003)	-0.060 (0.610)	0.217				
Soybean Meal	0.147 (0.000)		0.261 (0.082)	0.069	0.101 0.014	35.524 32.316		
	0.089 (0.006)	0.536 (0.003)	0.100 (0.476)	0.180				
Wheat	0.157 (0.000)		0.281 (0.038)	0.068	0.432 0.066	18.791 6.410		
	0.122 (0.002)	0.337 (0.168)	0.151 (0.351)	0.086				
Hogs	0.195 (0.000)		0.093 (0.390)	-0.006	0.032 0.086	31.029 21.885		
	0.108 (0.000)	0.858 (0.001)	-0.224 (0.119)	0.204				

^aFor each commodity, the first regression results are obtained by estimating

$$\sigma_{REAL} = \alpha_0 + \alpha_{LHV} \sigma_{LHV} + \varepsilon$$

$$\sigma_{REAL} = \alpha_0 + \alpha_{IFV} \sigma_{IFV} + \alpha_{LHV} \sigma_{LHV} + \varepsilon.$$

^bIf needed the estimates are adjusted for heteroskedasticity using the Newey-West procedure.

^cMDM test for statistical significance in the differences of the mean absolute percentage errors and the mean squared percentage errors between the implied forward volatility and the forecast based on the realized volatility during the same time period as the forward interval in the previous year.

^dMAPE and MSPE are the mean absolute and mean squared percentage errors.

Table 5. Predictive performance of the composite historical volatility (CHV) with respect to future realized volatility, 1992-2001, and test of forecast encompassing by the implied forward volatility (IFV)

Commodity	Regression ^{a,b}			Adj. R ²	MDM ^c	Errors ^d
	α_0 <i>p</i> -value	α_{IFV} <i>p</i> -value	α_{CHV} <i>p</i> -value		p_{MAPE} -value p_{MSPE} -value	MAPE MSPE
Corn	0.024 (0.429)		0.919 (0.000)	0.406	0.699 0.007	19.167 6.515
	0.019 (0.525)	0.648 (0.000)	0.296 (0.179)	0.514		
Soybeans	0.094 (0.031)		0.513 (0.034)	0.107	0.496 0.107	26.576 16.042
	0.076 (0.080)	0.724 (0.006)	-0.076 (0.791)	0.214		
Soybean Meal	0.102 (0.005)		0.488 (0.010)	0.108	0.754 0.149	28.896 19.456
	0.085 (0.017)	0.515 (0.012)	0.136 (0.503)	0.176		
Wheat	0.119 (0.007)		0.479 (0.021)	0.088	0.053 0.018	13.915 3.222
	0.114 (0.010)	0.262 (0.426)	0.265 (0.431)	0.081		
Hogs	0.104 (0.002)		0.575 (0.004)	0.080	0.667 0.487	23.013 8.484
	0.088 (0.019)	0.631 (0.002)	0.070 (0.723)	0.170		

^aFor each commodity, the first regression results are obtained by estimating

$$\sigma_{REAL} = \alpha_0 + \alpha_{CHV} \sigma_{CHV} + \varepsilon$$

$$\sigma_{REAL} = \alpha_0 + \alpha_{IFV} \sigma_{IFV} + \alpha_{CHV} \sigma_{CHV} + \varepsilon.$$

^bIf needed the estimates are adjusted for heteroskedasticity using the Newey-West procedure.

^cMDM test for statistical significance in the differences of the mean absolute percentage errors and the mean squared percentage errors between the implied forward volatility and the rolling out-of-sample forecast generated by regressing realized volatility on the moving average of the realized volatilities during the same time period as the forward interval in the three previous years and the realized volatility during the time period immediately preceding the date of the forecast where the length of this period equals the length of the forecasted interval.

^dMAPE and MSPE are the mean absolute and mean squared percentage errors.