Journal of Forecasting
J. Forecast. **24**, 155–172 (2005)
Published online in Wiley InterScience (www.interscience.wiley.com). **DOI**: 10.1002/for.950

Conditional Volatility Forecasting in a Dynamic Hedging Model

MICHAEL S. HAIGH*
US Commodity Futures Trading Commission and University of Maryland, USA

ABSTRACT
This paper addresses several questions surrounding volatility forecasting and its use in the estimation of optimal hedging ratios. Specifically: Are there economic gains by nesting time-series econometric models (GARCH) and dynamic programming models (therefore forecasting volatility several periods out) in the estimation of hedging ratios whilst accounting for volatility in the futures bid–ask spread? Are the forecasted hedging ratios (and wealth generated) from the nested bid–ask model statistically and economically different than standard approaches? Are there times when a trader following a basic model that does not forecast outperforms a trader using the nested bid–ask model? On all counts the results are encouraging—a trader that accounts for the bid–ask spread and forecasts volatility several periods in the nested model will incur lower transactions costs and gain significantly when the market suddenly and abruptly turns.

KEY WORDS bid–ask spread, multi-period hedging, dynamic programming, forecasting volatility, multivariate GARCH

INTRODUCTION

Volatility forecasting permeates the world of economics and finance. Indeed, a voluminous literature has emerged in the area of risk management literature simply because volatility, usually measured as the standard deviation of portfolio returns, is often taken to represent portfolio risk. However, much of the literature on volatility forecasting tends to be univariate in nature. In a hedging application this makes little sense since forecasting hedging ratios (which are mainly functions of variances and covariances of asset returns) requires multivariate forecasting. Therefore, since Engle (1982) and Bollerslev (1986) first introduced ARCH/GARCH models, time-series econometrics have been applied to model conditional variance and covariance dynamics for both cash and futures prices jointly—used in the forecasting of optimal hedging ratios. Significant gains have been reported from a risk management standpoint relative to more traditional, basic econometric techniques.

The underlying concept in the hedging literature is the notion that traders may optimally select combinations of cash and futures positions to minimize risk. However, as highlighted by Campbell *et al*. (1997), rather than there simply being one futures price for an asset there are in fact three rel-

* Correspondence to: Michael S. Haigh, US Commodity Futures Trading Commission, 1155 21st Street, NW, Washington, DC 20581, USA. E-mail: mhaigh@cftc.gov

evant to a trader: bid price, an asking price and the transaction price. Given that neither the bid nor the ask are usually reported by most open outcry markets, only a limited number of research papers have evaluated the importance of the bid–ask spread in trading activities, simply because bid–ask data is not usually available. Results from studies that have employed available bid–ask data have suggested that ignoring the spread can be costly. For instance, Bae *et al*. (1998) show that failing to consider bid–ask spreads would lead to false conclusions regarding the profitability of stock index futures arbitrage.

It is therefore the principle objective of this paper to introduce the role of the bid–ask spread on optimal hedging models that require multi-step time-varying volatility forecasts, by uniting the GARCH time-series hedging approach (which tends to rely on one-period-ahead forecasts) with another competing, yet very different approach which tends to rely on multi-step-ahead forecasts of volatility—dynamic programming (DP). An assessment is also made as to whether the hedging ratios, calculated from competing (traditional) models, are in fact statistically *and* economically different from this nested DP–GARCH model. Indeed, as pointed out by Granger (2001) and Granger and Pesaran (2000), having statistical measures may not be enough and so a quantifiable measure of economic significance should be required and a closer link between economic decision theory and forecasts should be made.

The remainder of this paper is as follows. First a brief overview of optimal hedging is presented, followed by the introduction of the nested DP–GARCH that ignores the role of the bid–ask spread (model I). Next the role of the bid–ask spread and the role of forecasting in the nested DP–GARCH approach is discussed and a modified DP–GARCH model (model II) is presented. The data used in the empirical analysis is discussed and the econometric estimation results are then presented, followed by a presentation of hedging results, complete with a description on the development of time-varying confidence bands around the optimal hedging ratios. Finally, the results and their implications are summarized in the conclusion.

HEDGE RATIO ESTIMATION

A popular method of determining an optimal hedging strategy is to employ what is commonly known as the minimum variance (MV) framework, where it is assumed that a trader minimizes the variability of wealth associated with an expected purchase (or sale). The MV methodology has been utilized in many studies because the components of the MV hedge ratio may be retrieved from variance and covariance forecasts of the underlying cash and futures prices. These combinations of cash and futures positions (typically expressed as proportions to one another) are referred to as hedge ratios. To understand the MV hedge ratio framework, consider a simple one-period wealth function

$$W_t = \left(C_t - \overline{C_{t-1}}\right) + b_{t-1}\left(\overline{F_{t-1}} - F_t\right) \tag{1}$$

where W_t is comprised of both the return on purchasing the asset at time $t - 1$ and selling it at time t and the returns from hedging the cost. Here C_t denotes the future cash price associated with commodity sale, $\overline{C_{t-1}}$ is the known cash price associated with the initial purchase, $\overline{F_{t-1}}$ is the known (short) futures price locked in at time $t - 1$, and F_t is the (long) futures price obtained to offset the original futures transaction at time t. b_{t-1} represents the OHR to be determined. The concept of the MV hedge method is to minimize the variance of the wealth function. For the simple case illustrated in (1), the variance of wealth, Var_{t-1}, may be written as

$$\mathrm{Var}_{t-1}(W_t) = \mathrm{Var}(C_t) + b_{t-1}^2 \mathrm{Var}(F_t) - 2b_{t-1}\mathrm{Cov}(C_t, F_t) \tag{2}$$

After taking the first-order condition for an extremum associated with the OHR and solving, we obtain

$$b_{t-1} = \frac{\mathrm{Cov}(C_t, F_t)}{\mathrm{Var}(F_t)} \tag{3}$$

which is comprised of the one-step-ahead forecasts of the variance and covariance estimates for underlying cash and futures prices.[1]

Traditionally the calculation of OHRs rested on the erroneous assumption that the covariance matrix, and hence the OHR, was constant through time. Consequently a large body of recent research has focused on utilizing the GARCH framework to compute time-varying (conditional) hedge portfolios. To estimate time-varying hedge ratios it is necessary to model jointly the first two moments of the cash and futures settlement prices. Focusing on the constant correlation model (one of the more basic approaches to modelling the second moments of asset prices), we model the prices as follows:

$$\Delta P = \mu + \varepsilon_t \qquad \varepsilon_t | \Omega_{t-1} \sim N(0, H_t) \tag{4}$$

where $\Delta P_t = (C_t, F_t)^{\mathrm{T}}$ is a (2×1) vector containing cash and futures prices (T is a transpose operator); μ is a (2×1) mean vector of cash and futures prices (the intercept or drift terms), respectively; ε_t is a (2×1) vector of mean zero, bivariate normally distributed cash and futures price innovations; Ω_{t-1} is the information set available at time $t - 1$; and H_t, where $\mathrm{vech}(H_t) = (h_{11,t}, h_{12,t}, h_{22,t})^{\mathrm{T}}$, is a (2×2) conditional covariance matrix. The constant correlation parameterization implies that the H_t matrix may be specified according to

$$\begin{aligned} h_{iit} &= \omega_i + \alpha_{i1}\varepsilon_{it-1}^2 + \beta_{i1}h_{iit-1} \\ h_{ijt} &= \rho_{ijt}(h_{iit}h_{jjt})^{\frac{1}{2}} \qquad i, j = 1(C_t), 2(F_t), i \neq j \end{aligned} \tag{5}$$

where ρ_{ijt} denotes the ijth constant conditional correlation. In general ρ_{ijt} can be time-varying (hence the t subscript), but consideration simplifications arise in estimation and inference if it is assumed that ρ_{ij} is constant for all t. Returning to the OHR problem, it follows that, given the (time-varying) nature of the variance–covariance matrix H_t, the time-varying OHR may be expressed as

$$b_{t-1} = \frac{\mathrm{Cov}(F_t, C_t | \Omega_{t-1})}{\mathrm{Var}(F_t | \Omega_{t-1})} = \frac{h_{12,t}}{h_{22,t}} \tag{6}$$

where b_{t-1} is the OHR conditional on all available information at time $t - 1$, represented by Ω_{t-1}.

[1] Several studies have used the rule in (3) to calculate the hedge ratio using OLS regression techniques because the slope coefficient equals the term shown in (3).

COMBINING DP HEDGING MODELS WITH GARCH TIME-SERIES TECHNIQUES

Assume that a trader starts with an initial amount of wealth, which is invested in the commodity for resale at a later date. Also the hedging decision for each purchase is made several periods prior to the terminal date and the hedged portfolio can then be updated (modified) at each period up to that time using GARCH time-series techniques. For simplicity, consider a three-period problem, $t-2$, $t-1$, t which is the delivery date. Three periods before the sale of the commodity, after the cash position has been established, the trader decides on the initial futures position that covers the trading period between period three and period two, $b_{(S)t-2}$ (where (S) indicates the use of just the futures settlement prices). The futures position then evolves in the sense that the quantity hedged in the next period, $b_{(S)t-1}$, may be different from the quantity hedged in $t-2$. The following period, t, the trader closes out all outstanding futures positions, sells the cash commodity and collects the proceeds. Hence, from the perspective of the current period, $t-2$, we can define wealth at the terminal date t, W_t, as

$$W_t = (1+r)\left(-\overline{C_{t-2}}\right) + (1+r)\left(\overline{F_{t-2}} - F_{t-1}\right)b_{(S)t-2} + (F_{t-1} - F_t)b_{(S)t-1} + C_t \tag{7}$$

Variable $\overline{C_{t-2}}$ is the initial (known) price at which the exogenously determined cash commodity is purchased; C_t is the stochastic cash price at which the commodity must be sold at the end of the three periods. $\overline{F_{t-2}}$ is the (known) futures price available at the initial time period that the decision is made; F_{t-1} and F_t are the stochastic futures prices in the respective periods; r is the one-period risk-free interest rate; $b_{(S)t-2}$ and $b_{(S)t-1}$ are the hedging ratios that capture the quantity of futures sold (bought if negative). The terminal monetary wealth W_t reflects the fact that the trader's futures account is marked to market, meaning that all profits and losses related to the futures positions are realized each period.

To remain consistent with the OHR literature, at each decision date the trader first decides the quantity to be hedged in order to minimize the variance of terminal wealth, given the cash position. The trader's objective at the initial time period, $t-2$, is to calculate the hedge ratio, $b_{(S)t-2}$, that minimizes the variability of terminal wealth. Therefore, to find the hedge ratio that would be used at time $t-2$, $b_{(S)t-2}$, the trader must forecast the hedge ratio that would be employed the following week, $t-1$. Following Mathews and Holthausen (1991), and hence working backwards, the trader estimates the hedge ratio that would be used the week prior to the cash sale, $b_{(S)t-1}$, in order to minimize the variability of wealth. The conditional variance of the wealth (and suppressing conditioning information notation for ease of reading) associated with that week is therefore

$$\begin{aligned}\mathrm{Var}_{t-1}(W_t) &= \mathrm{Var}_{t-1}\left[(1+r)\left(-\overline{C_{t-2}}\right) + (1+r)\left(\overline{F_{t-2}} - \overline{F_{t-1}}\right)\overline{b_{(S)t-2}} + \left(\overline{F_{t-1}} - F_t\right)b_{(S)t-1} + C_t\right] \\ &= b^2_{(S)t-1}\mathrm{Var}_{t-1}(F_t) + \mathrm{Var}_{t-1}(C_t) - 2b_{(S)t-1}\mathrm{Cov}_{t-1}(C_t, F_t)\end{aligned} \tag{8}$$

After obtaining the first-order condition for an extremum, and then solving for the optimal hedging ratio, we are left with precisely the same hedge ratio presented in equation (6). Therefore, the hedge ratio is simply comprised of the one-step-ahead forecast of the cash and futures settlement price covariance divided by the one-step-ahead forecast of the futures settlement price variance. Substituting the expression for $b_{(S)t-1}$ into the wealth expression (8), we can find the variance of wealth at the initial trade date.[2] Therefore, in the initial period (two periods prior to the eventual cash sale) the trader minimizes the variability of terminal wealth relevant at that date:

$$\begin{aligned}\mathrm{Var}_{t-2}(W_2) &= \mathrm{Var}_{t-2}(1+r)\left(-\overline{C_{t-2}}\right)+(1+r)\left(\overline{F_{t-2}}-F_{t-1}\right)b_{(S)t-2}+(F_{t-1}-F_t)b_{(S)t-1}+C_t\\ &=(1+r)^2 b^2_{(S)t-2}\mathrm{Var}_{t-2}(F_{t-1})+b^2_{(S)t-1}\mathrm{Var}_{t-2}(F_{t-1})+b^2_{(S)t-1}\mathrm{Var}_{t-2}(F_t)+\mathrm{Var}_{t-2}(F_t)\\ &\quad-2(1+r)b_{(S)t-2}b_{(S)t-1}\mathrm{Var}_{t-2}(F_{t-1})+2(1+r)b_{(S)t-2}b_{(S)t-1}\mathrm{Cov}_{t-2}(F_{t-1},F_t)\\ &\quad-2(1+r)b_{(S)t-2}\mathrm{Cor}_{t-2}(F_{t-1},C_t)-2b^2_{(S)t-1}\mathrm{Cov}_{t-2}(F_{t-1},F_t)+2b_{(S)t-1}\mathrm{Cov}_{t-2}(F_{t-1},C_t)\\ &\quad-2b_{(S)t-1}\mathrm{Cov}_{t-2}(F_t,C_t)\end{aligned} \tag{9}$$

The variance at time period $t-2$ is a function of several variances and covariances; the hedge ratio that is used at time $t-2$, $b_{(S)t-2}$ (the operational hedge ratio) and the expected hedge ratio to be used the next time period, $b_{(S)t-1}$ (the forecasted ratio). Taking the first-order condition of equation (9) and solving for the optimal first-period futures position (the OHR) gives

$$b_{(S)t-2}=\frac{\mathrm{Cov}_{t-2}(F_{t-1},C_t)}{(1+r)\mathrm{Var}_{t-2}(F_{t-1})}+\frac{b_{(S)t-1}}{(1+r)}\left[1-\frac{\mathrm{Cov}_{t-2}(F_t,F_{t-1})}{\mathrm{Var}_{t-1}(F_{t-1})}\right] \tag{10}$$

This hedge ratio might be viewed as the sum of an inter-subperiod hedge ratio plus the discounted next period hedge ratio weighted by a small positive weight if $\mathrm{Cov}_{t-2}(F_t, F_{t-1}) < \mathrm{Var}_{t-2}(F_{t-1})$. However, if the futures market in question can be shown to have little or no systematic bias then all terms remaining after the first term on the right-hand side of each OHR disappear. This is so because if a futures market can be shown to be unbiased then the hedge ratio at each time period can be shown to be independent of all other hedge ratios.[3] This implies that if we have unbiased futures markets, the OHRs developed within the DP framework differ because of the discount rate, and by the timing of the forecasts of volatility, but are independent of any future hedge ratios.

BID–ASK SPREADS AND THE DP–GARCH MODEL

According to Campbell *et al*. (1997) there are three futures prices relevant to the trader, a bid price, an asking price and the transaction price, not just the settlement price often used in empirical research. Therefore, even though the bid and ask price represent prices related to the same commodity, they may not be perfectly correlated and should not be treated as such.[4] It is assumed, as in

[2] To simplify the model, we follow Mathews and Holthausen (1991) by assuming that the hedger knows $b_{(S)t-1}$ at the initial trade date. This is not a restrictive assumption, as estimates of the variances/covariances were based on historical relationships and are easy to forecast using GARCH models.

[3] For example, representing the futures price at time period t as $F_t = F_{t-1} + u_t$ with a variance expressed as $\mathrm{Var}_t(F_t) = E_t(u_t)^2$. The futures price at period $t-1$ can then be expressed as $F_{t-1} = F_{t-2} + u_{t-1}$, and because $F_t = F_{t-2} + u_{t-1} + u_{t-2}$ we have $\mathrm{Cov}_{t-2}(F_{t-1}, F_t) = E_{t-2}(u_{t-1}, u_{t-1} + u_t)^2 = E_{t-2}(u_{t-1})^2$. Therefore, $\mathrm{Var}_{t-2}(F_{t-1}) = \mathrm{Cov}_{t-2}(F_{t-1}, F_{t-2})$, and so the hedge ratio collapses to $b_{(S)t-2}=\dfrac{\mathrm{Cov}_{t-2}(F_{t-1},C_t)}{(1+r)\mathrm{Var}_{t-2}(F_{t-1})}$.

[4] Recent research by Gagnon *et al*. (1998) employed a *trivariate* GARCH system allowing for time-varying covariability between related prices (that were not perfectly correlated) in a portfolio. They discovered that significant gains in hedging performance may be enjoyed by modelling (forecasting) the prices (variances) jointly in a portfolio compared to individual strategies. Moreover, further gains may be achieved simply because significantly fewer futures contracts would be recommended in the portfolio approach, thus reducing commission charges on the futures contracts.

model I, that in period $t-2$ the trader decides on the initial futures position, $b_{(BA)t-1}$ (where (BA) indicates that the hedging ratios have been developed using the bid and ask prices). However, now instead of going short at the settlement price recorded by the exchange, the trader must go short at the bid price which can in some instances be much lower than the settlement price, particularly in thinly to moderately traded markets (see data section). The futures position then evolves (like the case presented for model I) in the sense that the quantity hedged in period $t-1$, $b_{(BA)t-1}$, may be different from the quantity hedged in $t-2$. The next period, t, the trader closes out all outstanding futures positions, by offsetting the futures position by going long at the asking price rather than the settlement price (model I). The trader would then sell the cash commodity, and collect the proceeds (if any) from the hedge. Hence, from the perspective of the current period, $t-2$, we can define wealth at the terminal date t, W_t, as

$$W_t = (1+r)\left(-\overline{C_{t-2}}\right) + (1+t)\left(\overline{B_{t-2}} - A_{t-1}\right)b_{(BA)t-2} + (B_{t-1} - A_t)b_{(BA)t-1} + C_t \tag{11}$$

All variables are as previously defined except $\overline{B_{t-2}}$, which is the futures bid price available at the initial time period that the decision is made. B_{t-1} and A_t are the stochastic futures bid and ask prices in the respective periods; and $b_{(BA)t-2}$ and $b_{(BA)t-1}$ are the hedging ratios that capture the quantity of futures sold (bought if negative). The solution is once again obtained through backward induction, so in order to find the hedge ratio that would be used at time $t-2$, $b_{(BA)t-2}$, the trader must forecast the hedge ratio that would be employed the following week. The conditional variance of the wealth associated with that period is

$$\begin{aligned}\mathrm{Var}_{t-1}(W_t) &= \mathrm{Var}_{t-1}\left[(1+r)\left(-\overline{C_{t-2}}\right) + (1+r)\left(\overline{B_{t-2}} - A_{t-1}\right)b_{(BA)t-2} + (B_{t-1} - A_t)b_{(BA)t-1} + C_t\right] \\ &= b^2_{(BA)t-1}\mathrm{Var}_{t-1}(A_t) + \mathrm{Var}_{t-1}(C_t) - 2b_{(BA)t-1}\mathrm{Cov}_1(C_2, A_2)\end{aligned} \tag{12}$$

After obtaining the first-order condition for an extremum, and then solving for the OHR, we are left with $b_{(Ba)t-1} = \mathrm{Cov}_1(C_2, A_2)/\mathrm{Var}_1(A_2)$, which is composed of the one-step-ahead forecast of the cash and futures asking price covariance divided by the forecasted variance of the futures asking price.

Substituting the expression for $b_{(BA)t-1}$ into the wealth expression (12), we can find the variance of wealth at the initial trade date. Taking that first-order condition and solving for the optimal first-period futures position leaves us with

$$b_{(BA)t-2} = \frac{\mathrm{Cov}_{t-2}(A_t, P_t)}{(1+r)\mathrm{Var}_{t-2}(A_{t-1})} + \frac{b_{(BA)t-1}}{(1+r)}\left[\frac{\mathrm{Cov}_{t-2}(B_{t-1}, A_{t-1})}{\mathrm{Var}_{t-2}(A_{t-1})} - 1\right] \tag{13}$$

In this case, the second term on the right-hand side of equation (13) may not disappear simply because the bid prices and ask prices exhibit different behaviour in the short run, particularly in thinly to moderately traded markets (see data section). So only if $\mathrm{Cov}_{t-2}(B_{t-1}, A_{t-2}) = \mathrm{Var}_{t-2}(A_{t-1})$ are we left with a weighted initial hedge ratio, the first term on the right-hand side (which relies on a two-step-ahead forecast of the volatilities) where the next period hedge ratio, $b_{(BA)t-1}$, does not affect the current period hedge ratio $b_{(BA)t-2}$. As this may not be the case every time period, we have reason to believe that the next period hedge ratio will affect the current period hedge ratio.

To implement the DP–GARCH framework, regardless of whether we focus on model I or model II, a specification must be chosen for the time-varying covariance matrices. Based upon residual diagnostic tests (see econometric estimation results presented below), each series is specified as a simple martingale process, thereby satisfying the assumption of unbiased markets.[5] Because the constant correlation structure (presented in equation (5)) is parsimonious in parameters, and relatively easy to estimate, it is employed here. Multi-step-ahead forecasts of relevant variances and covariances are then required to generate the optimal hedging ratios from the underlying bivariate GARCH model (model I) and trivariate GARCH model (model II) at each trade date. Therefore, the general form of a forecast, at time t, of the volatility of the time-series variable over the period t to $t + M$ is represented as

$$\sigma^2_{t,t+m} = E[\varepsilon^2_{t+1} + \cdots + \varepsilon^2_{t+M} | \Omega_t] = \sigma^2_{t+1} + E[\sigma^2_{t+2} | \Omega_t] + \cdots + E[\sigma^2_{t+m} | \Omega_t] \tag{14}$$

For the GARCH (1, 1) model, this then becomes

$$\begin{aligned}\sigma^2_{t,t+M} &= \sigma^2_{t+1} + E[\omega + \alpha\varepsilon^2_{t+1} + \beta\sigma^2_{t+1} | \Omega_t] + E[\omega + \alpha\varepsilon^2_{t+2} + \beta\sigma^2_{t+2} | \Omega_t] \\ &\quad + \cdots + E[\omega + \alpha\varepsilon^2_{t+M-1} + \beta\sigma^2_{t+M-1} | \Omega_t]\end{aligned} \tag{15}$$

Using the definition of $\sigma^2_{t+1} = E[\varepsilon^2_{t+1}|\Omega_t]$, this can be expressed equivalently as

$$\sigma^2_{t,t+M} = \sigma^2_{t+1} + \omega + (\alpha + \beta)\sigma^2_{t+1} + \omega + (\alpha + \beta)E[\sigma^2_{t+2} | \Omega_t] + \cdots + \omega + (\alpha + \beta)E[\sigma^2_{t+M-1} | \Omega_t] \tag{16}$$

Recursive substitution yields

$$\sigma^2_{t,t+M} = \sigma^2_{t+1} + \omega + (\alpha + \beta)\sigma^2_{t+1} + \omega + (\alpha + \beta)E[\omega + (\alpha + \beta)\sigma^2_{t+1}] + \cdots \tag{17}$$

Collecting terms yields the following expression for the M-step-ahead conditional variance forecast:

$$\sigma^2_{t,t+M} = \sigma^2_{t+1} \sum_{i=0}^{M-1} (\alpha + \beta)^i + \omega \sum_{i=1}^{M-1} \sum_{j=1}^{i} (\alpha + \beta)^{j-1} \tag{18}$$

Using the facts that

$$\sum_{i=0}^{M-1} \lambda^i = \frac{1 - \lambda^M}{1 - \lambda} \quad \text{and} \quad \sum_{i=1}^{M-1} \sum_{j=1}^{i} \lambda^{j-1} = \frac{M}{1 - \lambda} - \frac{1 - \lambda^k}{(1 - \lambda)^2} \tag{19}$$

we yield the following expression for the M-step-ahead conditional variance for the GARCH (1, 1) model:

$$\sigma^2_{t,t+M} = M\sigma^{-2} + [\sigma^2_{t+1} - \sigma^{-2}] \frac{1 - (\alpha + \beta)^M}{1 - (\alpha + \beta)} \tag{20}$$

where $\sigma^{-2} = \omega/(1 - \alpha - \beta)$ is the steady state volatility of the GARCH process.

---

[5] Time-series processes are based on residual diagnostic results (Ljung–Box Q and Q^2 tests for white noise).

Estimating the hedge ratios also involves choosing an appropriate discount rate. In the ensuing analysis this discount rate is set throughout at 10%, and we continue with the three-period hedging scenario as in the example.

DATA

Daily closing futures prices (bid, ask and settlement) for white sugar traded at LIFFE were collected from Bloomberg International covering the period 13th December 1995–12th January 2000.[6] While bid–ask quotes are collected and recorded throughout the day (and are made publicly available), the final bid, ask and settlement price of the day were collected based on the assumption that our representative trader makes updates and hedging decisions towards the close of trading once a week. Therefore weekly price data (214 observations) were constructed using Wednesday prices. However, if closing bid–ask quotes based on Wednesday prices were not available, then a Thursday or a Friday price was used.[7] The futures prices are for the nearby contract month which forms the first value for the continuous series, and runs until the last day of trading of the contract.

In addition to the closing bid, ask and settlement prices weekly London cash prices for sugar covering the same time period were collected from Datastream International. Figure 1 illustrates the four price series (closing bid, ask, settlement and cash prices over the five-year horizon). As can be seen from the main graph it is extremely difficult to distinguish between the related futures price series as the average bid–ask spread as a percentage of the settlement price is very small. However, as can be seen from the smaller graphical insert, the bid, ask and settlement are quite different, and tend to move together over time, albeit by non-constant amounts. Such a phenomenon is not uncommon in moderately traded markets like the sugar market at LIFFE.[8] Figure 1 (lower panel) further illustrates this point by simply presenting a time-series plot of the bid–ask spread over the time horizon. As can be seen from the chart, the bid–ask spread varies by uniform amounts (the minimum tick size) and can on some occasions be as high as \$2 per tonne. The mean value of the spread over the time period is \$0.4557, while the modal value is \$0.2 or twice the size of the minimum tick value.

ECONOMETRIC ESTIMATION RESULTS

Each price series was first examined for the existence of a unit root using augmented Dickey–Fuller (ADF) tests. Results indicated that all four series (futures bid, ask, settlement and cash) are I(1). Correspondingly, each series was first differenced in the econometric estimation. Quasi-maximum likelihood estimates of model parameters were obtained using the BFGS (Broyden, Fletcher, Goldfarb and Shanno) algorithm so the estimates are consistent even if the conditional distribution of the residuals is non-normal (Bollerslev and Wooldridge, 1992). Residual diagnostics (Ljung–Box

[6] Unlike other commodities, sugar has always been traded on the electronic trading system, so all closing bid/asks and transaction volumes are available. In November 2000 all commodities switched from open outcry to electronic trading.
[7] A total of seven Thursday prices and five Friday prices were used.
[8] Total volume for the year 2000 for the sugar contract was 907,399, which ranked third in terms of total volume behind cocoa (1,636,322) and coffee (1,470,980).

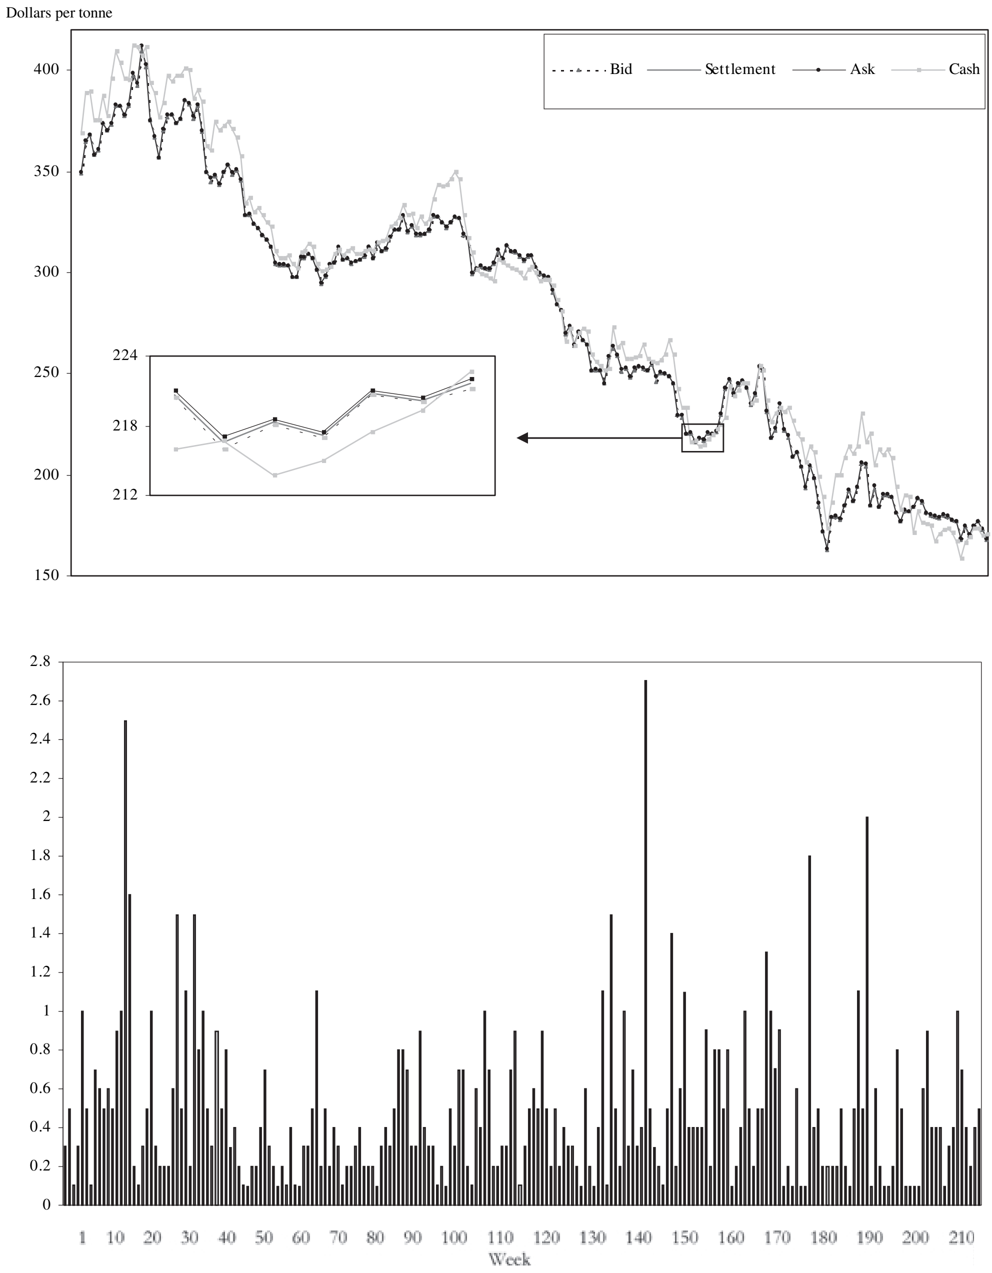

Figure 1. Weekly bid, settlement, ask and cash prices: December 1995–January 2000 (upper panel) and weekly Wednesday closing bid–ask spread (lower panel)

Q and Q^2 test statistics) for each model suggest that the GARCH (1, 1) models appear to do a reasonable job of explaining conditional mean and variance dynamics of all the prices.[9]

HEDGING RESULTS

The key question in any forecasting (hedging) evaluation study is how well does the proposed model perform relative to other models? To answer this question, we first turn our attention to the DP–GARCH model that forecasts volatility but only utilizes the cash and futures settlement prices (model I) and we then turn our attention to the DP–GARCH model that incorporates the bid–ask prices (model II) and forecasts volatilities several periods ahead. For both models an evaluation is made of their performance relative to other more standard models including a straightforward GARCH, OLS, naïve and unhedged models which do not employ any kind of recursive substitution.

Hedging results: model I

The left-hand panel of Table I presents sample average hedge ratios, along with standard errors around the average and minimum and maximum hedge ratios, for each of the DP–GARCH $b_{(S)t-2}$ and $b_{(S)t-1}$) and GARCH ($b_{(S)GARCH}$) models. Also included are the average values for the OLS, naïve

Table I. Descriptive statistics for hedge ratios for risk-minimizing static and dynamic objectives, three-week hedging horizon

Hedge model	Model I		Model II	
DP–GARCH	$b_{(S)t-2}$	$b_{(S)t-1}$	$b_{BA(t-2)}$	$b_{(BA)t-1}$
Avg.	0.761	0.763	0.748	0.749
SE	0.129	0.129	0.046	0.042
Min	0.489	0.490	0.552	0.567
Max	1.129	1.131	0.887	0.885
GARCH	$b_{(S)GARCH}$		$b_{(BA)GARCH}$	
Avg.	0.761		0.748	
SE	0.129		0.046	
Min	0.489		0.552	
Max	1.129		0.887	
OLS	$b_{(S)OLS}$		$b_{(BA)OLS}$	
	0.710		0.710	
Naïve	$b_{(S)NAIVE}$		$b_{(BA)NAIVE}$	
	1		1	
Unhedged	$b_{(S)UNHEDGED}$		$b_{(BA)UNHEDGED}$	
	0		0	

Note: The annualized discount rate, r, is 0.10. Avg. denotes sample average, SE is the corresponding standard error of the average of the hedge ratios. Min is the sample minimum and Max is the sample maximum. The GARCH hedge ratio for both model I ($b_{(S)GARCH}$) and model II ($b_{(BA)GARCH}$) represents the average hedge ratio that would be used by the trader over the entire trading period. It is equal to the hedge ratio used at $t-2$ by the DP–GARCH user as it is assumed that the simple GARCH user uses weekly data to form the hedge ratio to be applied at $t-2$ and left in place until the commodity is purchased at the end of the trading horizon. The OLS, $b_{(S)OLS}$ and $b_{(BA)OLS}$ and naïve, $b_{(S)NAIVE}$ and $b_{(BA)NAIVE}$ hedge ratios used each week are not, like the DP–GARCH and GARCH counterparts, updated each week.

[9] All these econometric results are excluded to conserve space but are available from the author upon request.

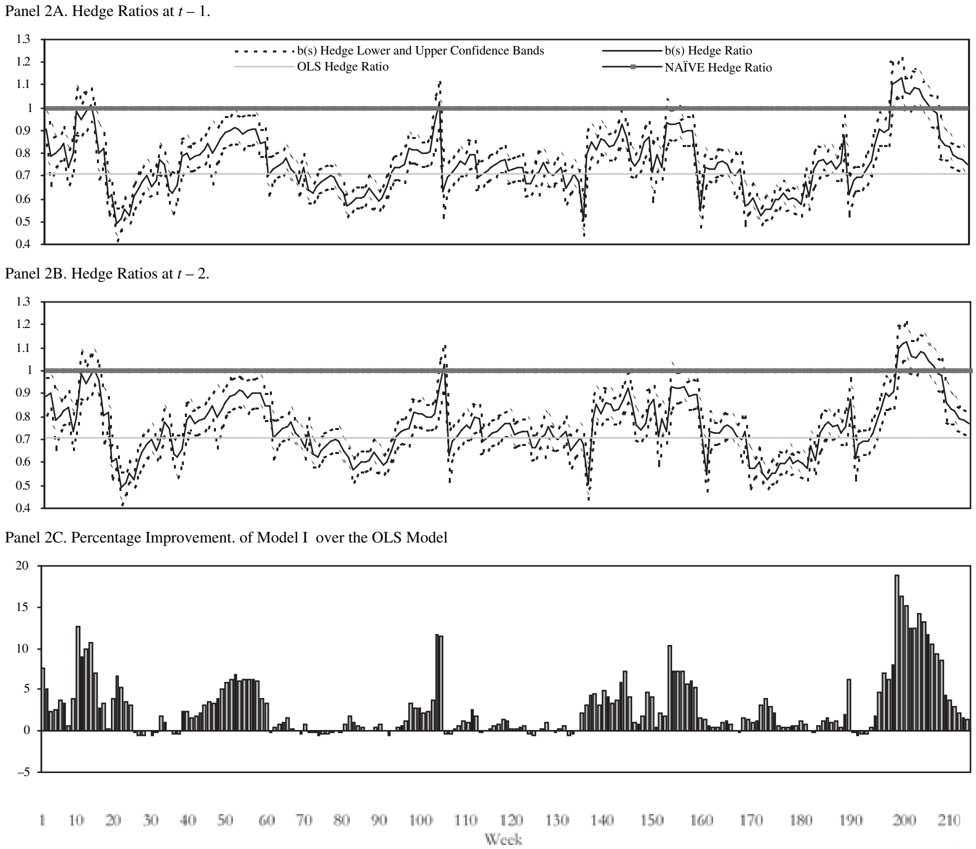

Figure 2. Model I (DP–GARCH model using $b(S)$ hedge ratios) with confidence bands at $t - 1$, $t - 2$ respectively, and percentage improvement over the OLS model
Note: The $b_{(S)}$ hedge ratios refer to the optimal hedging ratios calculated using the futures settlement price data.

and unhedged hedging strategies. Plots of the DP–GARCH, OLS ($b_{(S)OLS}$) and naïve ($b_{(S)NAIVE}$) hedge ratios for the sample period along with confidence bands are reported in the upper two panels of Figure 2. Since analytical expressions for hedging ratio standard errors may be impossible to obtain, an asymptotic approximation is applied here. In particular the delta method is employed, which amounts to a Taylor series approximation (Kendall and Stuart, 1977) for deriving standard errors around the hedge ratios. Specifically, for each of the hedging ratios: $b_{(i=S \text{ or } BA)}$ can be expressed as function $b_i(\theta)$ of a parameter vector $\theta = (\theta_1, \theta_2, \ldots, \theta_P)'$. If the covariance matrix of $\hat{\theta}$ is C and J is the gradient of $b_i(\theta)$, then approximately: $b_i(\theta) \doteq b_i(\theta_0) + (\theta - \theta_0)'J$, where $\theta_0 = E(\theta)$, so $\mathrm{Var}(b_i(\hat{\theta})) \doteq J'CJ$. The resulting time-varying confidence bands can then be calculated and can be seen surrounding the time-varying hedge ratios in Figure 2. To illustrate, in around week 200 the DP–GARCH hedge ratios at time $t - 1$ and $t - 2$ are larger than the naïve hedging ratio ($b_{(S)NAIVE}$), but as the confidence band overlaps with the naïve ratio we could infer that they are statistically indistinguishable from one another.

For the GARCH hedge ratio ($b_{(S)GARCH}$) it is assumed that once the hedge is in place it is not updated over the hedge horizon. Of course a weekly sampling frequency enables the trader relying on a myopic GARCH model to update the hedge ratio (but not use DP analysis); however, for comparison sake, this hedge is also left in place for the entire hedge period. Consequently, the average GARCH hedge ratio, $b_{(S)GARCH}$, is identical to the DP–GARCH hedge ratio developed at the start of trading. The OLS and naïve hedging ratios, $b_{(S)OLS}$ and $b_{(S)NAIVE}$, on the other hand, do not change from week to week, as they are simply obtained from an OLS regression of the change in the futures settlement price on the change in the cash price, or set equal to 1 respectively. The unhedged hedge ratio is set equal to 0 for each and every week.

As illustrated in Table I and Figure 2 (panels A and B) for the DP–GARCH portfolio we see substantial variation in OHRs at each hedge horizon through time, but there is relatively little variation among OHRs across hedge horizons. To illustrate, during the initial period, $t-2$, the average hedge ratio for the trader is 0.7613; conversely the hedging ratio in the next period is 0.7627, indicating that on average the hedge ratio increases modestly. The fact that the hedge ratio increases over time is consistent with the findings of Anderson and Danthine (1983). If no variation occurred across the hedge horizon, the DP–GARCH hedge ratios ($b_{(S)t-2}$ and $b_{(S)t-1}$) would be identical to the GARCH hedge ratios ($b_{(S)GARCH}$), and there would be no incentive in combining the DP and GARCH approaches and forecasting several periods ahead. Results reported in Table I also reveal for the DP–GARCH portfolio that, at most, about 113% of the cash position would have been hedged by the sugar trader, with the least amount hedged being about 49% of the cash position.

Results reported in Table I and Figure 2 illustrate that the DP–GARCH hedge ratios are quite erratic, sometimes recommending that about 50% of the hedged position be lifted, or locked into in just a matter of weeks. Such a recommendation could imply that a trader might incur substantial transaction costs associated with updating the portfolio. The lower part of Table II reveals that the OLS hedge ratio ($b_{(S)OLS}$) for the sugar trader is 0.7104, suggesting somewhat less hedging on average than either the DP–GARCH model or the static GARCH model. By adopting the OLS, naïve or unhedged approaches, the hedger would employ the same hedge ratio every week over the entire time frame, and so this approach shows no variability.

While sample and average hedge ratios are instructive, they tell us little about how the various models perform. Therefore, an economic significance measure is derived to complement the statistical measure calculated from the time-varying confidence bands. These results, along with other descriptive statistics, are reported for each model in Table II. Also, for illustration purposes panel 2C presents the time-varying variance percentage improvement of the DP–GARCH model over the OLS model.

According to Table II there appear to be some gains to using both the GARCH and the DP–GARCH approach relative to the more basic models in terms of average variance reductions. The performance of the GARCH model is of no surprise and appears to be very close (in terms of performance) to other papers that have evaluated its performance (e.g., Baillie and Myers, 1991). What are the advantages to using the DP–GARCH model compared to the GARCH approach? According to Table II not much. In particular, the average performance of the DP–GARCH model that forecasts ahead over the static GARCH is just 0.2203%. This number would clearly be more significant to a trader hedging a large quantity of sugar, but it might also be of importance to know the variance around this improvement. While the average hedge ratios appear quite similar and hence average improvement quite low, there are periods of time when the static model is beaten quite convincingly. For instance, the DP–GARCH outperforms the static GARCH by approximately 9% around week 104 (not shown). Indeed, the percentage improvement of the DP–GARCH model over the static GARCH verifies that there are indeed times when the market suddenly and abruptly turns and the trader following the static GARCH methodology may have lost out. It is clear that on average

Table II. Descriptive variance statistics for static and dynamic objectives for model I

DP–GARCH using $b_{(S)t-2}$ and $b_{(S)t-1}$	
Avg.	133.210
SE	50.960
Min	62.219
Max	337.382
GARCH, static using just $b_{(S)t-2}$	
Avg.	133.840
SE	51.127
Min	62.346
Max	337.490
OLS, static using $b_{(BA)OLS}$	
Avg.	137.930
SE	53.436
Min	62.565
Max	631.070
Naïve, static using $b_{(BA)NAIVE}$	
Avg.	142.410
SE	60.900
Min	65.803
Max	468.280
Unhedged, static using $b_{(S)UNHEDGED} = 0$	
Avg.	231.670
SE	84.900
Min	104.010
Max	464.700
% Variance reduction from using the DP–GARCH model relative to:	
GARCH, static using just $b_{(S)t-2}$	0.220%
OLS, static using $b_{(BA)OLS}$	2.840%
Naïve, static using $b_{(BA)NAIVE}$	5.118%
Unhedged, static using $b_{(S)UNHEDGED} = 0$	42.274%

Note: The annualized discount rate, r, is 0.10. Avg. denotes sample average, SE denotes the corresponding standard deviation. Min is the sample minimum and Max is the sample maximum. Results are based on 212 weekly hedging periods. The DP–GARCH model uses the optimal hedging ratios generated from the cash and futures settlement prices. The GARCH, static model only employs the hedging ratio developed in $t - 2$, and does not optimally update. The OLS model uses the hedging ratio developed from a simple regression of the futures settlement price on the cash price. The naïve and unhedged ratios are set equal to 1 and 0, respectively.

the DP–GARCH approach outperforms this basic alternative (in fact all alternatives), with the worst performing strategy being the unhedged approach (as one might expect). Interestingly, there are several times when the DP–GARCH model is outperformed by the OLS (and naïve and unhedged approaches). However, the negative percentage improvement figures associated with these time periods are not of large magnitude (see panel 2C), suggesting that even if the DP–GARCH is beaten by alternatives the trader would not be too heavily penalized.

Hedging results: model II

Results presented on the right-hand side of Table I verify that not only are the average hedging ratios lower in model II compared to model I, but their variability is also lower. For instance, at time t − 2 the average optimal hedging ratio associated with the DP–GARCH model that uses bid and ask prices ($b_{(BA)t-2}$) is 0.7476 compared to the average ratio recommended in model I ($b_{(S)t-2}$) of 0.7613. A similar pattern emerges in $t - 1$ whereby the average optimal hedging ratio ($b_{(BA)t-1}$) is 0.7490 compared to the hedge ratio of 0.7627 recommended by model I. These results seem to verify the findings of GLM (1998) that hedging ratios estimated in a portfolio tend to be lower than those estimated in a bivariate setting (suggesting lower transaction costs). The other important observation is the fact that the standard errors around the average hedging ratios are much lower in model II compared to model I. In particular the standard errors associated with $b_{(BA)t-2}$ and $b_{(BA)t-1}$ in model II are 0.0459 and 0.0420 respectively, whereas the corresponding standard errors in model I are 0.1288 and 0.1287—about three times as large. Comparing the OHRs for model I and model II, complete with their confidence bands, suggests that the OHRs from the DP–GARCH framework in model II are far more stable but do experience time-variation (this is not shown but is available upon request). The hedging ratios from the DP–GARCH framework in model I are much more volatile.

Table III presents the descriptive variance statistics for all the models and how they compare to the bid–ask DP–GARCH approach. The results are quite striking. Model II (DP–GARCH using $b_{(BA)t-2}$ and $b_{(BA)t-1}$) appears to outperform, in terms of reduced variability, model I (DP–GARCH using $b_{(S)t-2}$ and $b_{(S)t-1}$) by 1.833%, which on the surface does not seem like a dramatic improvement. However, as described previously, averages can be deceiving and it is found that when the model I hedge ratios are extremely volatile the trader would lose out by following that approach, if the true model was indeed the model incorporating the bid and the ask prices (model II). It is clear from Table III that the DP–GARCH model that utilizes both $b_{(BA)t-2}$ and $b_{(BA)t-1}$ outperforms all the 'basic' alternatives, beating the unhedged model by 52.105%. The OLS model performs better when evaluated against model II (compared to model I) because the OLS hedging ratios are closer to the DP–GARCH hedge ratios, simply because the DP–GARCH ratios are less volatile as they were estimated in a portfolio setting. Unlike the case of model I, while the percentage improvement over the OLS approach is lower, the DP–GARCH model is *never* beaten by the simpler alternative. The same results are obtained for the naïve and unhedged approach.

While the DP–GARCH approach using $b_{(BA)t-2}$ and $b_{(BA)t-1}$ performs the best out of the simpler alternatives, it does not seem to significantly outperform the static DP–GARCH approach (that just utilizes $b_{(BA)t-2}$). Indeed the percentage reduction from using the dynamic model over the static approach is 0.099%, suggesting that multi-period-ahead forecasting may not actually be worthwhile. Again, this result is not particularly surprising given the finding that the OHRs do not vary as much as the DP–GARCH model presented in model I. However, traders are probably more interested in the distribution of the variability of the improvement over time. To this end, Figure 3 provides more evidence on this. Firstly, the upper panel illustrates that in general the hedge ratios generated from the dynamic DP–GARCH ($b_{(BA)t-2}$ and $b_{(BA)t-1}$) are not, in general, statistically different from the hedge ratio employed both periods in the static version ($b_{(BA)t-2}$). That is, the confidence bands overlap for most of the period of time analysed. One might suspect therefore that if there is no statistical difference between the OHRs then there might not be any improvement, from an economic sense. The lower panel illustrates that when the percentage improvement is very low (e.g., week 40 to week 100) the OHRs are statistically indistinguishable. However, when the market suddenly and abruptly turns—around week 107—(see the graphical inset in the upper panel), the performance of the dynamic approach improves. This may lead us to a simple conclusion. The dynamic hedging ratios

Table III. Descriptive variance statistics for static and dynamic objectives for model II

DP–GARCH using $b_{(BA)t-2}$ and $b_{(BA)t-1}$	
Avg.	237.750
SE	42.855
Min	148.560
Max	515.460
GARCH, static using just $b_{(BA)t-2}$	
Avg.	238.013
SE	43.155
Min	148.960
Max	517.810
DP–GARCH using $b_{(S)t-2}$ and $b_{(S)t-1}$	
Avg.	242.131
SE	42.672
Min	150.823
Max	524.361
GARCH, static using just $b_{(S)t-2}$	
Avg.	242.850
SE	43.103
Min	150.621
Max	526.082
OLS, static using $b_{(BA)OLS}$	
Avg.	239.079
SE	43.144
Min	151.824
Max	527.319
Naïve, static using $b_{(BA)NAIVE}$	
Avg.	269.283
SE	57.064
Min	160.308
Max	656.380
Unhedged, static using $b_{(BA)UNHEDGED} = 0$	
Avg.	494.356
SE	58.082
Min	359.168
Max	822.060
% Variance reduction from using the DP–GARCH model relative to:	
GARCH, static using just $b_{(BA)t-2}$	0.099
GARCH, dynamic using $b_{(S)t-2}$ and $b_{(S)t-1}$	1.833
GARCH, static using just $b_{(S)t-2}$	2.095
OLS, static using $b_{(BA)OLS}$	0.567
Naïve, static using $b_{(BA)NAVIE}$	11.297
Unhedged	52.105

Note: The DP–GARCH model uses the hedging ratios generated from cash and bid and ask prices. The GARCH, static also employs the cash, bid and ask prices but only the hedging ratio developed in $t - 2$, and does not optimally update. The DP–GARCH (using the $b_{(S)}$ ratios) is the bid–ask model that employs both the hedging ratios developed from model I. The GARCH static also uses the bid–ask model, but employs the hedging ratio from model I ($b_{(S)t-2}$). The OLS model uses the hedging ratio developed from a simple regression of the futures settlement price on the cash price. The naïve and unhedged ratios are set equal to 1 and 0, respectively.

Panel 3A. Hedge Ratios and confidence bands at $t-1$ and $t-2$

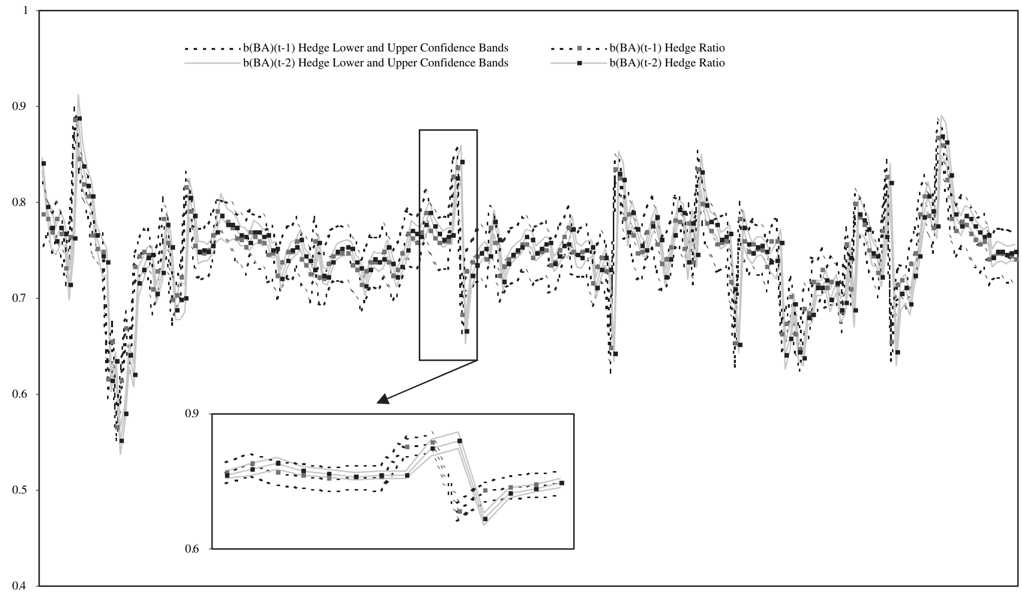

Panel 3B. Percentage Improvement of Model II over the static version of the model.

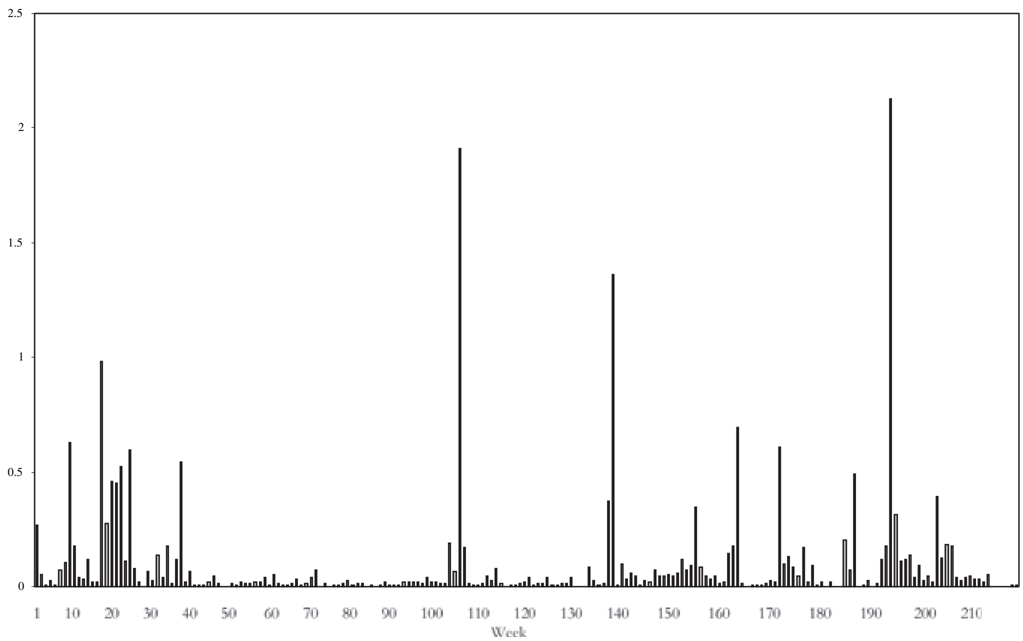

Figure 3. Model II hedge ratios at $t-1$, $t-2$ with confidence bands (panel 3A) and percentage improvement over the static version of model II (panel 3B)
Note: Model II utilizes the optimal hedges established at $t-2$ and then updates by using the optimal hedge ratio at $t-1$. The static model II utilizes the optimal hedge established at $t-2$ in both periods.

are forward-looking (they are forecasts) and so if the new hedge ratio ($b_{(BA)t-1}$) is not different from the older hedge ratio ($b_{(BA)t-2}$) from a statistical sense, then it is unlikely that any economic reward will be yielded. Consequently, the trader would save on transaction costs from updating the old hedge ratio.

CONCLUSIONS

In this paper several questions were addressed simultaneously. Are there any advantages to combining two hedging strategies that rely on forecasting—dynamic programming (DP) and time-series econometrics (GARCH) models whilst accounting for variability in the bid–ask spread? Are the generated optimal hedging ratios from this approach statistically different from more standard approaches? Are there any economic gains to be enjoyed by combining the approaches and therefore forecasting volatilities several periods ahead, and are there certain times when a trader following a more basic strategy loses out? On all counts the results here are encouraging. First, given the ability to forecast volatilities enables us to derive a rule for developing the optimal DP–GARCH hedge ratios using cash and futures settlement price data. Second, using the delta method, time-varying confidence bands were obtained such that a trader could distinguish whether or not optimal hedging ratios are statistically different from other models.

While the DP–GARCH model does slightly outperform the static GARCH approach on average, results verify that a trader that forecasts with the static GARCH approach would lose out when the market suddenly and abruptly turns. Incorporating bid–ask prices into the trader's portfolio resulted in hedge ratios lower than those recommended in the portfolio using just settlement and cash prices, suggesting lower transaction costs, with much less volatility associated with the forecasted hedging ratios. While the gains from following the DP–GARCH approach after accounting for the bid–ask spread over the static approach are small, on average there are times that the trader would lose out when the market suddenly and abruptly turns, just like in the more basic model that ignores the spread. The more sophisticated forecasting models like DP–GARCH or static GARCH that ignore the bid–ask spread are outperformed by more stable hedging strategies like the OLS, when evaluated in the bid–ask environment. This result is also consistent with previous research whereby ignoring a natural portfolio results in far more volatile hedge ratios, which may induce risk rather than reduce it. The implication here being that if the true portfolio relevant to the trader really does incorporate the bid–ask prices (which should be more readily available with the electronic platform), then the system to be estimated should involve three price series rather than two.

This research shows that if the forecasts of the DP–GARCH model incorporating the bid–ask spread suggest a statistically significantly different hedging ratio compared to the hedge ratio employed the previous period, the trader should update the portfolio. Alternatively, in more stable periods, the trader should continue forecasting with a static GARCH model that uses all three prices relevant to the trader, and enjoy potentially lower transaction costs.

ACKNOWLEDGEMENTS

Haigh is at the US Commodity Futures Trading Commission and the University of Maryland. The US CFTC disclaims any responsibility for any private publication or statement of any CFTC employee or Commissioner. This study expresses the author's views and does not necessarily reflect those of the Commission, the Commissioners, or other members of the staff.

REFERENCES

Anderson RW, Danthine JP. 1983. The time pattern of hedging and the volatility of futures prices. *Review of Economic Studies* **50**: 249–266.

Bae K, Chan K, Cheung Y. 1998. The profitability of index futures arbitrage: evidence from bid–ask quotes. *The Journal of Futures Markets* **18**: 743–763.

Baillie RT, Myers RJ. 1991. Bivariate GARCH estimation of the optimal commodity futures hedge. *Journal of Applied Econometrics* **6**: 109–124.

Bollerslev T. 1986. Generalized autoregressive conditional heteroskedacity. *Journal of Econometrics* **31**: 301–327.

Bollerslev T, Wooldridge JM. 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* **11**: 143–172.

Campbell JY, Lo AW, Mackinley AC. 1997. *The Econometrics of Financial Markets*. Princeton University Press: Princeton, NJ.

Dickey DA, Fuller W. 1981. Likelihood ratio tests for autoregressive time series with a unit root. *Econometrica* **49**: 1057–1072.

Engle RF. 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. *Econometrica* **50**: 987–1008.

Gagnon L, Lypny GJ, McCurdy TH. 1998. Hedging foreign currency portfolios. *Journal of Empirical Finance* **5**: 197–220.

Granger CWJ. 2001. Macroeconometrics—past and future. *Journal of Econometrics* **100**: 17–19.

Granger CWJ, Pesaran MH. 2000. Economic and statistical measures of forecast accuracy. *Journal of Forecasting* **19**: 537–560.

Kendall MG, Stuart A. 1977. *The Advanced Theory of Statistics*, Vol. 1, *Distribution Theory*, 4th edn. Macmillan: New York.

Mathews KH, Holthausen DM Jr. 1991. A simple multiperiod minimum risk hedge model. *American Journal of Agricultural Economics* **73**: 1020–1026.

Author's biography:

Michael S. Haigh received his PhD in Economics with a minor in Statistics from North Carolina State University. He is responsible for research and teaching related to futures and options markets and financial econometrics. Current research includes the application of economic models and experimental and quantitative methods to commodity and financial markets.

Author's address:

Michael S. Haigh, Associate Professor, University of Maryland and Senior Financial Economist, Office of the Chief Economist, United States Commodity Futures Trading Commission, 1155 21st Street, NW, Washington, DC 20581, USA.

