AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Combining Time-Varying and Dynamic Multi-Period Optimal Hedging Models

by

Michael S. Haigh and Matthew T. Holt

WP 02-08

Department of Agricultural and Resource Economics
The University of Maryland, College Park

Title: **Combining Time-Varying and Dynamic Multi-Period Optimal Hedging Models**

Authors: **Michael S. Haigh, University of Maryland ^a and Matthew T. Holt, North Carolina State University^b**

Abstract: This paper presents an effective way of combining two popular, yet distinct approaches used in the hedging literature – dynamic programming (DP) and time-series (GARCH) econometrics. Theoretically consistent yet realistic and tractable models are developed for traders interested in hedging a portfolio. Results from a bootstrapping experiment used to construct confidence bands around the competing portfolios suggest that while DP-GARCH outperforms the GARCH approach they are statistically equivalent to the OLS approach when the markets are stable. Significant gains may be achieved by a trader, however, by adopting the DP-GARCH model over the OLS approach when markets exhibit excessive volatility.

Key Words: bivariate GARCH; dynamic programming; multi-period hedging

JEL Classification: C32, D81, D92, G11

Manuscript status: European Review of Agricultural Economics, in press.

Michael S. Haigh (corresponding author)^a
Assistant Professor
Department of Agricultural and Resource Economics
University of Maryland
College Park, MD 20742
(301) 405-7205 (Tel)
(301) 314-9091 (Fax)
mhaigh@arec.umd.edu (e-mail)

Matthew T. Holt ^b
Professor
Department of Economics
North Carolina State University
Raleigh, NC 27695-8109
(919) 513-4605 (Tel)
(919) 515-5613 (Fax)
matt_holt@ncsu.edu (e-mail)

We thank seminar participants at the 2000 American Agricultural Economic Association for their helpful advice and comments on an earlier version of this paper. We also thank three anonymous reviewers and the editor for helpful comments that substantially improved the paper.

Copyright © 2002 by Michael S. Haigh and Matthew T. Holt
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.

Combining Time-Varying and Dynamic Multi-Period Optimal Hedging Models

Abstract: This paper presents an effective way of combining two popular, yet distinct approaches used in the hedging literature – dynamic programming (DP) and time-series (GARCH) econometrics. Theoretically consistent yet realistic and tractable models are developed for traders interested in hedging a portfolio. Results from a bootstrapping experiment used to construct confidence bands around the competing portfolios suggest that while DP-GARCH outperforms the GARCH approach they are statistically equivalent to the OLS approach when the markets are stable. Significant gains may be achieved by a trader, however, by adopting the DP–GARCH model over the OLS approach when markets exhibit excessive volatility.

Key Words: bivariate GARCH; dynamic programming; multi-period hedging

JEL Classification: C32, D81, D92, G11

1. Introduction

Recent empirical research on optimal hedging has focused on two distinct areas. One branch of the literature has investigated models wherein there is a multi-period hedging horizon and the optimal hedge may be updated each period. This research has relied largely on applications of dynamic programming (DP), with examples including Anderson and Danthine (1983); Karp (1983); Martinez and Zering (1992); Mathews and Holthausen (1991) and Vukina and Anderson (1993). Overall, these studies find that, while intuitively appealing (but operationally complicated), multi-period hedging models have not resulted in significant improvements over static models from a risk management perspective.

The other thrust in the literature has been to use time-series econometrics to model conditional variance and covariance dynamics for commodity cash and futures prices. To this end various versions of the ARCH/GARCH framework advanced by Engle (1982) and Bollerslev (1986) have been utilized, the result being that time-varying one-step-ahead hedge portfolios are estimated (e.g., Baillie and Myers, 1991).^{1,2} Gagnon et al. (1998), Haigh and Holt (2000), Kroner and Sultan (1993), Lin et al. (1994), Myers (1991), Park and Switzer (1995), Sephton (1993), Tong (1996), and Yeh and Gannon (2000) have employed ARCH/GARCH methods, and have reported significant gains in hedge performance relative to more traditional OLS techniques include. Like the DP approach, GARCH methods allow hedge updating. Unlike the DP approach, the time-series approach thus far has been limited to determining only sequentially updated one-period-ahead optimal hedge ratios (OHRs).

In this paper we first present a method of utilizing GARCH time series models within a DP framework to construct dynamic optimal hedging portfolios. Dynamic hedging strategies are developed for a hypothetical firm (merchandiser) interested in purchasing in advance but on a

weekly basis commodities (specifically, cocoa and sugar) used in food manufacturing. For each commodity we begin by specifying theoretically consistent yet realistic and tractable risk management models. We allow the representative firm to adjust the optimal hedge several times between initial hedge placement and eventual commodity purchase several weeks into the future. In each case the hedger is assumed to minimize the variability of total costs. By employing DP recursion relations, OHRs are thereby derived.

This research departs from prior studies by relaxing the assumption that hedgers will not revise their estimates of underlying variances/covariances over time and that they will not update futures positions. The representative merchandiser's variance expressions are manageable functions of non-contemporaneous, time-varying variances and covariances, requiring multi-step ahead forecasts that are incorporated directly into the DP framework. Because GARCH specifications are effective in modeling time-varying volatility, the model developed here illustrates the gains from combining dynamic hedging strategies in a DP framework with modern time-series techniques. Indeed, the research not only expands on both the time-series and dynamic programming hedging literature, thus being of interest to an academic audience, but also, because of its practicality and ease of use, should be of interest to risk-management practitioners responsible for developing optimal hedging strategies with varying time horizons.

Another unique aspect of this study is the development of confidence bands around competing portfolios. Specifically, by employing a parametric bootstrap we are able to evaluate whether portfolio variance reductions from each method are statistically different. We are therefore able to shed light on circumstances under which one method outperforms another. This analysis will offer a more complete understanding of the conditions that are likely to offer similar portfolio payoffs and, from a practical standpoint, will enable a trader to decide when to

employ a simple strategy (like OLS techniques) and when to embark on a more sophisticated strategy (like a GARCH or DP-GARCH approach).

The remainder of the paper is organized as follows. First, we present a brief overview of hedge ratio estimation, and then introduce the DP–GARCH model. This is followed by a description of the data used in the empirical analysis and econometric estimation results. We then present hedging results and the bootstrap analysis. The final section concludes.

2. Hedge Ratio Estimation

A basic concept in the hedging literature is the notion that traders optimally select combinations of cash and futures positions to minimize portfolio risk. These combinations, typically expressed in terms of proportion of cash to futures positions for an asset, are commonly referred to as Optimal Hedge Ratios (OHRs). One popular method of determining OHRs is to employ a minimum-variance (MV) framework, wherein an agent (e.g., a merchandiser) is assumed to minimize variability of outlays (costs) associated with an expected purchase. For several reasons the MV framework has become the benchmark in the hedging literature. First, MV hedge ratio is optimal for exceptionally risk averse traders (Ederington, 1979; Kahl, 1983). As well, MV hedge ratio is also optimal when futures markets are unbiased. This result is important as such a phenomenon has been verified in several empirical studies (Baillie and Myers, 1991; Martin and Garcia, 1981).³ As such, MV methodology has been widely applied, in part because of the theoretical justification of finding unbiased markets and in part because components of the MV hedge ratio may be retrieved from variance and covariance estimates of underlying cash and futures prices (see, e.g., Baillie and Myers, 1991; Kroner and Sultan, 1993).

Many studies calculate OHRs from historical data by simply regressing changes in cash prices on changes in futures prices (see Appendix equations (A1) – (A3)). The resulting slope

coefficient, b_{t-1} , is interpreted as an estimate of the OHR (Ederington, 1979; Kahl, 1983).⁴ This result holds because in the simple least squares model the estimated slope coefficient equals the term shown in (A3). This form of the hedge ratio is commonly referred to as the MV hedge ratio, and has been employed in many studies including that by Mathews and Holthousen (1991).

Implicit in traditional hedge ratio estimation methodology is the assumption that the covariance matrix of cash and futures prices—and hence the hedge ratio—is constant through time. But as Fama (1965) observed, variances and covariances of asset returns are often *not* constant over time—large changes tend to be followed by other large changes and small changes tend to be followed by other small changes. This phenomenon, known as volatility clustering, is effectively captured by Engle’s (1982) ARCH model; in an early application Cecchetti et al. (1988) employed a bivariate ARCH process to infer time-varying OHRs. Bollerslev (1986) subsequently proposed the GARCH model to circumvent the long lags often needed to specify correctly an ARCH model. A large body of research has therefore focused on utilizing the GARCH framework to construct time-varying (conditional) hedge portfolios.

Here we ignore the potential need for a time series structure for the mean of cash and futures prices. The relevant price equations may then be specified as

$$\begin{aligned}\Delta \mathbf{p}_t &= \mathbf{m} + \mathbf{e}_t, \\ \mathbf{e}_t | \Omega_{t-1} &\sim N(\mathbf{0}, H_t),\end{aligned}\tag{1}$$

where $\Delta \mathbf{p}_t = (c_t, f_t)^T$ is a (2 x 1) vector containing the first difference of cash and futures prices (T is a transpose operator); $\mathbf{m}$ is a (2 x 1) mean vector (i.e., a (2 x 1) vector of intercept or drift terms) for, respectively, cash and futures prices; $\mathbf{e}_t$ is a (2 x 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 $\text{vech}(H_t) = (h_{11,t}, h_{12,t}, h_{22,t})^T$, is a (2×2) conditional covariance matrix.⁵

One method for specifying the structure generating conditional second moments is to utilize the positive semi-definite parameterization (PSD) explained in Engle and Kroner (1995). In their setup the form of H_t , the conditional covariance matrix, is

$$H_t = W^T W + A^T \mathbf{e}_{t-1} \mathbf{e}_{t-1}^T A + B^T H_{t-1} B, \quad (2)$$

where W is a symmetric (2×2) parameter matrix and A and B are (2×2) parameter matrices. Alternative GARCH specifications include the constant correlation parameterization employed by, among others, Cecchetti et al. (1988), Garcia et al. (1995), and Kroner and Sultan (1993). The set up, while parsimonious, does not allow the cash-futures covariance (and, therefore, OHRs) to switch signs in the short run as cash and futures prices move in opposite directions (see, e.g., Haigh and Holt, 2000). The implication is this model may be overly restrictive. The linear diagonal model used by Baillie and Myers (1991) is less restrictive in this regard, but does not ensure the conditional covariance matrix is positive definite for all t .

The bivariate GARCH process defining H_t in (2) is specified so that each $h_{ij,t}$ term in H_t is a linear function of lagged values of H_{t-1} , as well as lagged values of products and cross-products of innovations $\mathbf{e}_{t-1} \mathbf{e}_{t-1}^T$. A notable feature of this specification is that all elements in H_{t-1} and $\mathbf{e}_{t-1} \mathbf{e}_{t-1}^T$ are permitted to influence the ij_{th} component of H_t . As such equation (2) is the analogue to a VARMA model in the system's conditional covariance process, and therefore provides a general way of modeling time dependencies in the system's conditional second moments. By construction the specification in (2) ensures H_t is positive semi-definite at all data points and so is used in this analysis. As well, (2) allows inferences about the extent to which

lagged variances (innovations) in, say, the cash (futures) market influences the current variance in the futures (cash) market.

Returning to the optimal hedging problem, it follows that, given the time-varying nature of variance-covariance matrix H_t , the time-varying hedge ratio may be expressed as

$$b_{t-1} = \frac{Cov(f_t, c_t | \Omega_{t-1})}{Var(f_t | \Omega_{t-1})} = \frac{h_{12,t}}{h_{22,t}},$$

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

Numerous studies have focused on time-varying hedging by using variants of the ARCH/GARCH approach to model the cash-futures price distribution. OHRs are then inferred by relaxing the assumption that conditional variances (covariances) are time independent. For instance, Cecchetti et al. (1988) applied a bivariate ARCH model to financial futures prices. Myers and Baillie (1991), Myers (1991), and Sephton (1993) applied bivariate GARCH models to commodity prices. Park and Switzer (1995), Tong (1996), and Yeh and Gannon (2000) compared GARCH-generated OHRs to OLS hedging for stocks. Finally, Kroner and Sultan (1993) and Lin et al. (1994) used a bivariate GARCH framework in foreign currency hedging.

3. Combining DP Hedging Models with GARCH Time-Series Techniques

In this section we outline the theoretical multi-period hedging models for commodity merchandisers that undertake weekly *purchases* of cocoa and sugar. Our representative firm is assumed to hedge cash price uncertainty associated with anticipated commodity purchase several weeks in future. That is, the merchandiser maintains the required amount of cocoa or sugar for processing on a weekly basis. To minimize uncertainty associated with each week's anticipated purchase, the firm uses futures market as a hedge.⁶

The hedging framework developed here is multi-period, and builds upon the somewhat restrictive dynamic models presented by Mathews and Holthausen (1991). We focus on long hedging in part because extensive research in agricultural economics has assessed risk attitudes and the use of derivative instruments from a *short hedgers* perspective (e.g., Binswanger, 1982; Pennings and Smidt, 2000). This focus has occurred in spite of Berck (1981) and Shapiro and Brorsen's (1988) findings that only 5 per cent and 11.4 per cent, respectively, of farmers' actually use futures markets to hedge. The need for more sophisticated hedging models, at least from a producer's perspective (i.e., short hedging), is therefore called into question. Alternatively, Nance et al. (1993) studied the determinants of corporate hedging, which is more typically associated with long hedging. They reported that 61 per cent of corporate firms in their sample used hedging instruments. Collins (1997) reached a similar conclusion, observing that many manufacturers do employ hedging techniques. Corporate buyers are therefore more likely to employ hedging, and accordingly are more apt to utilize sophisticated hedging techniques, than are primary producers. Empirical evidence supporting actual futures market utilization by corporate hedgers (i.e., merchandisers or processors) provides the first of two reasons why we focus on long hedging.

The second reason we focus on long hedging is there are no overriding issues regarding storage. See, for example, Lence et al. (1993). This assumption is not overly restrictive, however, as the gains from hedging effectiveness from using a model that incorporates speculative storage decisions are, according to these authors, negligible. Importantly, we assume the merchandiser does not purchase the commodity and store it (at least beyond one week). In this way our framework resembles more closely that of Mathews and Holthausen (1991).

Given that we focus on a long hedging scenario, how then might we determine whether in fact a DP–GARCH framework outperforms a simple GARCH framework (without recursion relations for optimal updating) and a time-invariant OLS, and, correspondingly, SUR, framework, where the differenced futures price is regressed on the differenced cash price?⁷ Such comparisons are of interest for several reasons. First, based on results reported by Mathews and Holthausen (1991) there is little evidence that dynamic (DP) hedge ratios differ in magnitude over the hedging horizon; there may be negligible improvement (measured by the reduction in portfolio variability) obtained by following a DP approach relative to static OLS models, casting doubt on the usefulness of the DP approach. Second, and more importantly, previous empirical research has assumed a trader does *not* optimally update variance and covariance terms over time in the DP framework. For this reason we also compare our advancement of the DP approach (the DP–GARCH model) with a basic GARCH approach, which, by itself, has been shown to outperform myopic alternatives (Baillie and Myers, 1991; Gagnon et al., 1998; Haigh and Holt, 2000; Park and Switzer, 1995; Yeh and Gannon, 2000). Both DP-GARCH and basic GARCH are then compared with the static OLS/SUR approach. Doing so isolates gains from using GARCH relative to OLS/SUR and, as well, gains from using DP–GARCH over GARCH and OLS/SUR approaches. We may therefore isolate the importance of introducing time-varying variances-covariances without the DP framework with that of time-varying variances-covariances within a DP framework, thereby allowing a comparison of a truly time-varying *and* dynamic GARCH approach with the more typical time varying but non-dynamic GARCH model.

In the DP–GARCH setup we assume merchandiser’s price paid for the commodity in future, that is, several weeks ahead, is uncertain. In other words, the merchandiser’s purchasing

cost for the raw commodity is stochastic. For each commodity we therefore derive, on a weekly basis, the optimal number of contracts for the merchandiser to lock into in order to minimize total cost variability. Each merchandiser is then allowed to update relevant variance-covariance estimates—used to generate the hedged position—multiple times between when the initial hedge is placed and when the commodity is ultimately purchased in the cash market.

By using a DP-GARCH approach the merchandiser decides the amount hedged at each decision date in order to minimize the variance of total terminal cost, $Cost_t$. To illustrate, with four trading dates total cost is defined as

$$Cost_t = -c_t + b_{t-1}(f_{t-1} - f_t) + rb_{t-2}(f_{t-2} - f_{t-1}) + r^2b_{t-3}(f_{t-3} - f_{t-2}) + r^3b_{t-4}(f_{t-4} - f_{t-3}), \quad (3)$$

where week t denotes the terminal period and r represents the merchandiser's opportunity cost of funds. As indicated in (3), at the end of week t the commodity is finally purchased in the cash market at price c_t . At the beginning of week $t-1$ and for two weeks prior to that (i.e., weeks $t-2$ and $t-3$) the trader may adjust the position set initially at $t-4$. $Cost_t$ is therefore total (discounted) cost paid for the commodity at the end of the four-week horizon, accounting for gains (losses) from futures market. Setup (3) is similar to that utilized in prior research employing DP hedging. The difference is we focus on a merchandiser who will ultimately *purchase* the commodity; our methodology could of course also be applied to a short hedgers problem.

In the DP-GARCH framework the merchandiser's aim is to minimize total variance of cost by determining the optimal number of futures contracts to lock into each period $t-i$, where $0 \leq t-i \leq t$. In other words, at each period $t-i$ the merchandiser chooses the optimal number of futures contracts to minimize total variance of cost, which is simply the sum of final cash price paid, c_t , and (discounted) futures market gains/losses accrued over the four-week period preceding cash purchase. As illustrated by Mathews and Holthausen (1991) and Anderson and

Danthine (1983) the general solution to this problem is obtained through backward induction. Therefore, to find cost-minimizing hedge ratio b_{t-4} at initial trade date $t-4$, the merchandiser must also determine relevant hedges to employ at $t-3$, $t-2$, and $t-1$. That is, to minimize variance of discounted purchase cost we work backwards with the merchandiser first choosing the OHR that would be used at the start of delivery week, $t-1$. The conditional variance—where we have suppressed conditioning information notation—is shown in Appendix equation (A2). The first-order condition for an extremum and the corresponding OHR are also presented in the Appendix.

At time $t-2$ (two weeks prior to the eventual cash purchase) the merchandiser minimizes variance of total cost relevant at that date, illustrated in Appendix equation (A4). The variance at $t-2$ is a function of several variances and covariances; the hedge ratio used at $t-2$, b_{t-2} , referred to as the operational hedge ratio; and the expected hedge ratio to be used at $t-1$ (b_{t-1}).⁸ The first-order condition with respect to b_{t-2} —Appendix equation (A5)—is a function of both b_{t-1} and b_{t-2} , a system of two equations in two unknowns. But given unbiased commodity markets, the OHR at $t-2$ reduces to

$$b_{t-2} = -\frac{Cov_{t-2}(c_t, f_{t-1})}{rVar_{t-2}(f_{t-1})}. \quad (4)$$

See Appendix for details. Importantly, with unbiased markets hedge ratio (4) is simply the covariance-to-variance ratio weighted by discount factor r .

Corresponding cost functions, variance expressions, and first order conditions at $t-3$ and $t-4$ are also functions of the operational hedge ratios for that week and, as well, the hedge ratios expected to be used during subsequent weeks prior to cash purchase. In addition to discount factor r , the resulting hedge ratios are also functions of variances and covariances over the following weeks. Again, assuming unbiased commodity markets, these hedge ratios collapse to

simple expressions similar in form to those used in later weeks. Specifically, it may be shown that the optimal hedge ratios at $t - 3$ and $t - 4$ are simply

$$b_{t-3} = -\frac{Cov_{t-3}(c_t, f_{t-2})}{r^2 Var_{t-3}(f_{t-2})}, \text{ and } b_{t-4} = -\frac{Cov_{t-4}(c_t, f_{t-3})}{r^3 Var_{t-4}(f_{t-3})}, \quad (5)$$

where the influence of discount factor r increases the further away is the expected purchase date.

Based on the foregoing results for the four-week hedging horizon example, it is clear the optimal dynamic hedging rule (given unbiased markets) is easily generalized. Determining the OHR therefore involves picking an appropriate discount rate and estimating the relevant variance-covariance forecasts at each date. In the ensuing analysis we set discount value r to 10 per cent and continue with the four-week hedging scenario as in the example.

For the basic GARCH portfolio it is assumed that once a hedge has been placed it is not updated. Of course a weekly sampling frequency enables the trader relying on a basic GARCH model to update the hedge ratio for each subsequent week's purchase; however, for comparison sake this hedge is left in place for the entire four-week period. The implication is the average GARCH hedge ratio, b_{GARCH} , is identical to the DP-GARCH hedge ratio b_{t-4} developed at the beginning of the trading period. But as already indicated this hedge is left in place over the four-week horizon so that b_{GARCH} is unaffected by changes in volatility over the remainder of the month and, unlike DP-GARCH, is unaffected by future hedging ratios. As a result, because there is no provision for optimally updating the hedge basic GARCH hedging will likely be inferior to DP-GARCH in periods of short-run market volatility.

Of course once determined the OLS/SUR hedge ratio $b_{\text{OLS/SUR}}$ never changes, as it is simply obtained from an OLS/SUR regression of the futures price on the respective cash price. *A priori*, this approach will likely be inferior to DP-GARCH and basic GARCH, especially in

volatile markets, because it in no way adjusts to short-run volatility. As well, $b_{OLS/SUR}$ does not account for future hedging decisions and changes in volatility like the DP-GARCH approach.⁹

To summarize differences between models, a typical four-week time line for the basic GARCH and DP-GARCH hedging approaches would look as follows:

The OLS/SUR time-line would be similar to the GARCH time-line as only one hedge ratio is employed over the four-week period. The main difference is the OLS/SUR hedging ratio would be the same for *every* four-week period. Once determined the GARCH hedge is left in place for each four-week horizon; there is no updating over that month but updating can and does occur the following month.

In some instances the hedging ratios derived from these approaches might be similar; hence, the portfolio values could well be indistinguishable from each other. Specifically, if H_t in (2) is reasonably stable, changing little over the sample period, then OLS and basic GARCH models should yield similar results. Similarly, if H_t remains stable over the sample period, DP-GARCH hedge ratios would effectively differ from basic GARCH hedges (and, as well, the OLS/SUR hedges) only by the discount factor r .¹⁰

Unlike prior studies employing DP methods to obtain OHRs, we assume traders revise their estimates of relevant variances-covariances through time by using a GARCH framework. Using the PSD framework (presented in (2) above) multi-step-ahead forecasts of relevant variances and covariances, derived from the underlying bivariate GARCH process, may then be made.

To estimate time-varying hedge ratios it is necessary to model jointly the first two moments of cash and futures prices relevant to each merchandiser. 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. That is, for each cash price c_t , and for each futures price f_t , the model outlined in (1) and (2) is estimated.¹¹

With the econometric methodology for estimating time-varying variances-covariances in hand, we may now ask whether or not a DP-GARCH model that utilizes backward induction to generate and update OHRs outperforms a basic GARCH model without optimal updating and, as well, a static OLS hedging approach in practice. Obtaining information on the relative performance of these various hedging approaches from both an economic and statistical standpoint is the issue to which we now turn.

4. Data and Econometric Estimation Results

Weekly cash and futures prices for cocoa and sugar were collected for the period January 4th 1985–August 4th 2000, yielding a total of 814 weekly observations for the in-sample analysis. In addition cash and futures prices were also collected for the period August 11th 2000–September 9th 2002, yielding 57 observations for an out-of-sample analysis. Cocoa (Ivory Coast) and sugar (#11 World Raw) cash prices were provided by Reuters (formerly Bridge CRB). Futures settlement prices for cocoa and sugar are from the NYBOT, and were also provided by

Reuters. As cash prices are only reported on Friday, the futures price series was also constructed by using Friday prices. Futures prices are for the nearby contract month, which forms the first value for the continuous series and runs until the last trading day of the contract.

By using Augmented Dickey-Fuller (ADF) (tests), each price series was first examined for existence of a unit root. Results indicate all four series are nonstationary. Applying the same tests to the differenced series, the null hypothesis of a unit root is rejected. Therefore, we conclude each series is $I(1)$, that is, integrated of order one, in the sense of Engle and Granger (1987). Correspondingly, each series was first differenced in econometric estimation.

Maximum likelihood estimates of model parameters are obtained by using the Davidson-Fletcher-Powell (DFP) algorithm in conjunction with symmetric numerical derivatives. Parameter estimates for each bivariate PSD GARCH model are presented in Table 1. Point estimates of GARCH parameters indicate substantial evidence of conditional variance dynamics for each commodity. As well, a number of off-diagonal elements in A and B matrices for each PSD model are large in absolute terms relative to their asymptotic standard errors. Estimates of skewness and kurtosis parameters for the standardized residuals—defined by m_3 and m_4 , respectively—are reported for each equation in each portfolio model in the lower panel of Table 1. Results indicate no serious departures from the conditional normality assumption.

To more formally assess the relevance of the structure implied by a bivariate PSD GARCH specification, several sets of maintained restrictions were imposed on parameter matrices A and B and the models re-estimated. One set of models was estimated wherein all parameters in A and B were set to zero. These restrictions result in a model with a constant conditional covariance structure, that is, a model without GARCH effects, and is consistent with a constant (non time-varying) hedge ratio as provided by an SUR model. Alternatively, each

PSD model is re-estimated by restricting all off-diagonal elements of A and B to zero. As reported in Table 2, for each model both sets of restrictions were overwhelmingly rejected by the data. In addition, intercept (drift) terms were placed in each equation. In both instances likelihood ratio test results revealed these terms could be omitted without harming model fit. On balance there is substantial empirical support for the PSD parameterizations reported in Table 1.

The $Q(20)$ and $Q^2(20)$ test statistics for, respectively, standardized residuals and squares and cross products of (standardized) residuals for each model are reported in the lower panel of Table 1. Results of these diagnostics tests indicate the models do a reasonable job of explaining conditional mean and variance dynamics of cash and futures prices. Overall, the PSD–GARCH(1,1) models do a good job of characterizing essential features of the data, and are therefore potentially useful tools for examining dynamic hedging strategies.

5. Hedging Results

Given the empirical support for the PSD-GARCH(1,1) specification, the natural question is how well does DP-GARCH hedging portfolio perform relative to either a basic GARCH or an OLS/SUR approach for a representative cocoa/sugar merchandiser? To examine this issue in-sample average hedge ratios, along with standard deviations and minimum and maximum values, for each of the DP–GARCH, GARCH, and OLS models and for each portfolio are reported in the upper panel of Table 3. Comparable summary statistics for the out-of-sample analysis are recorded in the lower panel of Table 3. As well, in sample plots of DP–GARCH, basic GARCH, and OLS hedge ratios are reported in figures 1-2 for sugar and cocoa, respectively. As indicated previously, hedge ratios in Table 3 were derived by assuming merchandisers minimizes total variance of cost and use an annualized discount rate r of 10 per cent.¹²

As illustrated in the top panel of Table 3 and figures 1-2, on average each model calls for long hedging, as indicated by the negative signs associated with OHRs. This outcome is as expected for a risk averse merchandiser anticipating making cash purchases up to four weeks hence. Furthermore, regardless of the method used more hedging is recommended in the cocoa portfolio than in the sugar portfolio. While results for the DP-GARCH portfolio show substantial variation in OHRs at each hedge horizon through time, there is relatively little variation among OHRs across hedge horizons. This general result holds for both commodities. To illustrate, during the initial period, $t-4$, the average hedge ratio b_{t-4} for a cocoa merchandiser is -0.918 ; conversely, during the last period in which the portfolio may be adjusted, $t-1$, the average hedge ratio, b_{t-1} , is -0.923 . These results indicate that on average the cocoa hedge ratio increases (in absolute terms) by only 0.50 per cent as actual purchase date approaches. Results for sugar portfolio show a similar pattern, with, on average, $b_{t-4} = -0.778$ and $b_{t-1} = -0.783$ for the DP-GARCH. If no variation occurred across hedge horizon, then DP-GARCH hedge ratios would equal GARCH hedge ratios (ignoring the discount rate), and there would be no incentive to combining approaches. Results reported in Table 3 also reveal that for DP-GARCH and GARCH portfolios, at most, a cocoa merchandiser would have hedged 212 per cent of the cash position; correspondingly, at most 112 per cent of the sugar merchandiser's expected cash position would have been hedged.

Results reported in the top portion of Table 3 and figures 1-2 also illustrate that during the in-sample period hedge ratios vary considerably. As suggested by figures 1 and 2, there is more variability in OHRs at every horizon for sugar than for cocoa. As well, for the sugar portfolio it is clear that during a few periods short hedging is called for by the DP-GARCH model (a hedge ratio of approximately 0.52). But as cash and futures variances (prices) may

(and do) move in opposite directions in the short run, it is not unreasonable for the optimal hedge to call for the merchandiser to go short in some instances. Indeed, it is precisely for this reason that we chose to apply unrestrictive GARCH structures (i.e., the PSD), as other specifications may not pick up on such patterns. As indicated in figure 2, short hedging is not observed for cocoa using the DP-GARCH model; the minimum hedge ratio for cocoa is near zero (-0.064).

The final row in Table 3 reveals that OLS/SUR hedge ratio for cocoa is -0.899 , suggesting somewhat less hedging on average than either the DP-GARCH or the basic GARCH model. Alternatively, OLS/SUR hedge ratio for sugar is -0.723 , which is slightly greater than those called for by DP-GARCH. By adopting an OLS/SUR approach the hedger would employ the same hedge ratio every week over the entire sample period, indicating no variability in the proportion hedged. Overall, results for cocoa and sugar portfolios provide an interesting contrast in that cocoa basis is generally more predictable than sugar basis, as suggested by the observation that cocoa hedge ratios are typically closer to one in absolute value than are sugar hedge ratios (see, e.g., figures 1-2).

Traders are perhaps less concerned about past events than they are with future performance. It is for this reason that an out-of-sample evaluation was made of the competing models. Therefore, for each approach a forecast is made of the following week's hedge ratio. For each period the required variance and covariance terms are predicted and the hedge ratio calculated. Each model is then re-estimated with the new observation included in the sample and the evaluation repeated. This prediction/re-estimation algorithm is continued until all out-of-sample data are exhausted. Summary statistics of hedging ratios are presented in the lower panel of Table 3. As with the in-sample analysis, DP-GARCH portfolio shows variation in OHRs at each hedge horizon through time, although the variation is not as great as for the in-sample

period. Alternatively, there is relatively little variation among OHRs across hedge horizons. For instance, during initial period, $t-4$, the average out-of-sample hedge ratio b_{t-4} for cocoa is -0.945 ; whereas, during the last period, $t-1$, average hedge ratio b_{t-1} is -0.950 , indicating again that, on average, proportions hedged increase as the purchase date approaches. A similar pattern emerges for sugar portfolio where, on average, $b_{t-4} = -0.797$ and $b_{t-1} = -0.801$ for DP-GARCH. Results also reveal for DP-GARCH and basic GARCH portfolios that, out-of-sample, at most 120 per cent of the cash position would have been hedged by a cocoa merchandiser; correspondingly, at most 96 per cent of a sugar merchandiser's expected cash position would have been hedged.

For both commodities out-of-sample hedging ratios exhibit less volatility than do their in-sample counterparts (the standard deviation drops from approximately 0.138 to 0.081 for cocoa and from 0.198 to 0.180 for sugar), suggesting that more complicated models that incorporate emerging volatility patterns (like DP-GARCH and basic GARCH) are less likely to perform better than OLS/SUR. Of course out-of-sample the OLS/SUR model also exhibits some variability, with standard deviations of 0.002 for both cocoa and sugar. This result occurs because data are added each week and the models re-estimated. Even so, out-of-sample OLS/SUR hedge ratios are still nearly constant, and would likely perform better in more stable markets where the relationship between cash and futures remains steady.

Overall sample and average hedge ratios are instructive, and give some guidance as to the potential usefulness of each model; they do not, however, indicate much about how the various models perform in each respective portfolio. To this end average variance of total cost is computed for each model and for each commodity. These results, along with other descriptive statistics including the standard deviation of cost variance and minimum and maximum cost

variances are reported, for each model, in Table 4. The top half summarizes the results from the in-sample analysis, the lower half for the out-of-sample analysis. First, the results are quite striking for the in-sample analysis. Importantly, there appears to be rewards to using the DP–GARCH approach relative to the basic GARCH approach in terms of average variance reductions. This result may not be particularly surprising, however, because the simple GARCH model, while allowing for time-variation in hedge ratios, does not optimally use information to update hedge ratios in a way provided for by the DP–GARCH approach.

As reported in Table 4, average variance of total cost for cocoa merchandiser is 30228.77 per ton with DP–GARCH, while the corresponding value for basic GARCH is 30461.37, representing a 0.76 per cent gain by using DP–GARCH relative to the simple GARCH strategy. Also, the standard deviation associated with the average variance of the DP-GARCH model (18736) is considerably lower than for basic GARCH (20359), suggesting that following the DP-GARCH method would yield a more stable risk management strategy for a cocoa hedger. Based on estimated hedge ratios, the cocoa market basis seems more predictable than sugar as hedging ratios for cocoa are closer to one in absolute value. The implication is there may be even greater rewards to employing more sophisticated models (i.e., DP-GARCH) over standard approaches (i.e., basic GARCH) in relatively more unpredictable markets like sugar, where recommended hedging ratios are smaller in absolute terms. To illustrate, average variance for sugar total cost is 2.248 for the DP–GARCH model; the corresponding value for basic GARCH is 2.501, implying a 10.14 per cent gain associated with using DP–GARCH in lieu of the basic GARCH methodology. Not unlike the cocoa market, portfolio variance ‘volatility’ for DP-GARCH sugar model is lower than for simple GARCH (reported standard deviation of 0.051 versus 0.060). There is apparently a reward in terms of risk reduction by using the DP–GARCH framework

relative to a simple hedge-and-hold GARCH strategy, especially in markets where basis is less predictable. Rewards to following the DP–GARCH strategy are, however, highly dependent upon the total costs of buying the commodity confronted by a particular merchandiser, suggesting the DP–GARCH model would in fact benefit more greatly a merchandiser making frequent but large purchases.¹³

Turning to OLS/SUR portfolios, Table 4 reveals that while basic GARCH outperforms OLS/SUR for cocoa, the OLS/SUR approach actually dominates the GARCH approach for sugar. Interestingly, the standard GARCH model for sugar, when evaluated within a DP framework, does not perform as well as might be reported in myopic settings, where, as mentioned previously, GARCH models tend to outperform OLS/SUR approaches. Our results do show reward to following the DP–GARCH methodology over the straightforward OLS/SUR model. To illustrate, relative to OLS/SUR the DP–GARCH model yields 1.58- and 8.18 per cent reductions in total cost variability for cocoa and sugar merchandisers, respectively. Moreover, volatility of DP-GARCH variance is lower than for OLS/SUR models, as indicated by standard errors (18736 versus 20001.74 and 0.05 versus 0.07 in, respectively, the cocoa and sugar markets), suggesting a more stable risk management strategy over the OLS/SUR alternative. The implication is that using an OLS/SUR approach would result, at least in some periods, in foregone cost-reduction opportunities relative to the DP-GARCH approach, with the most compelling evidence offered in the sugar market.

Results from the out-of-sample analysis illustrates that the effectiveness of DP-GARCH and GARCH approaches is somewhat dampened. As suggested previously, sample hedging ratios seem to indicate that markets exhibited less volatility in the out-of-sample data period—particularly in the cocoa market—implying more stable (constant) models are likely to perform

better. This conjecture is confirmed by, on average, poorer performance of DP GARCH and GARCH models relative to OLS/SUR. According to the average improvement figures reported in the lower panel of Table 4, there is essentially no difference between the DP-GARCH and basic GARCH models (average percentage improvement is < -0.001 per cent). Moreover, out-of-sample DP-GARCH model is outperformed by the more stable OLS approach (average improvement of -2.32 per cent). The sugar market did exhibit more volatility out-of-sample (noting that the standard deviation of the hedging ratios was much higher at 0.18). Accordingly, DP-GARCH outperforms both GARCH and OLS/SUR approaches by 8.37 per cent and 1.80 per cent, respectively.

Like all statistics described thus far, the figures recorded in Table 4 are averages. The implication is a trader may not only be concerned with average model performance, but may also want to know whether a model could be detrimental if the market exhibits an unexpected burst of volatility or a period of prolonged stability. For this reason we also focus on period-by-period estimates to assess whether one model outperforms another under certain conditions. To assess period-by-period performance we first employ bootstrapping techniques to generate confidence intervals around both the hedging ratios and the portfolio variances. Specifically, drawing from the appropriate bivariate normal distribution, 500 parametric bootstrap estimates of both the DP-GARCH and basic GARCH portfolios are obtained for both commodities. We then assess whether one model is statistically different from another model (by determining whether or not confidence bands overlap), and if so, by how much and how many times that model outperformed the alternatives.

To summarize the results of the bootstrap evaluations, we first estimate confidence bands around the GARCH and DP-GARCH models and find that, statistically speaking, these two

models are equivalent. That is, during no period were the variances from the DP-GARCH model (with confidence bands) *not* overlapping the variance of the GARCH model. We also compare the DP-GARCH model with OLS/SUR. The results of this comparison are summarized in Table 5. The upper panel of Table 5 illustrates that 37.1 per cent (33.7 per cent) of the time OLS/SUR model is statistically equivalent to DP-GARCH for the cocoa (sugar), suggesting that a trader, for this percentage of the time, would be indifferent between the approaches. The lower panels of Table 5 provide more insight into the merits of each approach. The middle panel of Table 5 illustrates the percentage of times OLS/SUR variance is significantly less than the DP-GARCH approach for cocoa and sugar portfolios (47.2 per cent and 48.9 per cent, respectively). Furthermore, the greatest improvement in portfolio variances of DP-GARCH vis-à-vis GARCH amounts to 78.18 per cent and 73.99 per cent for, respectively, cocoa and sugar.

These results seemingly contradict findings presented earlier whereby DP-GARCH outperforms OLS on average. The implication is there must be times when OLS/SUR approach is beaten dramatically by DP-GARCH. Indeed, a cursory glance at the lower panel of Table 5 suggests this very result. While DP-GARCH exhibits fewer times that it actually dominates OLS/SUR (15.7 per cent and 17.9 per cent in the cocoa and sugar models, respectively), there are times—as illustrated by the maximum improvement numbers—that the OLS/SUR model proves quite detrimental to the trader. Indeed, in the case of sugar OLS/SUR is outperformed on one occasion by 608.67 per cent in the sugar model and 212.96 per cent in the cocoa model.¹⁴

Figures 3 and 4 shed more light on this issue. Figure 3 illustrates over the 1992 – 1998 period the difference between sugar DP-GARCH variance (with confidence bands) and sugar OLS/SUR variance. Also depicted (in the insert) is a time-series plot of the cash and futures sugar prices for 1995, resulting cost variances, and bootstrapped confidence intervals. An

interesting pattern emerges in that when cash and sugar prices are moving consistently together, as they are at the beginning of 1995, OLS/SUR performs satisfactorily. As the co-movement relationship breaks down, however, and the basis becomes more volatile, DP-GARCH tends to outperform OLS/SUR. Specifically, OLS/SUR variance, complete with confidence bands, is greater than the variance (and bands) generated by DP-GARCH. Toward the end of the year as the basis normalizes OLS/SUR approach improves and eventually outperforms DP-GARCH.

As presented in Figure 4, a similar (but less vivid) scenario is illustrated for cocoa. Figure 4 depicts DP-GARCH variance and confidence bands and OLS portfolio variance over the 1990 – 1996 period, along with a plot of the basis (cash and futures prices). The cocoa market, unlike sugar, has a somewhat more predictable basis; due to mild basis volatility only in 1990 is OLS/SUR beaten by DP-GARCH. From 1991 – 1996 the basis was more predictable, hence the acceptable performance of OLS/SUR.

6. Conclusion

In recent years numerous empirical hedging studies—typically relying upon variants of ARCH/GARCH-type models—have adopted the term ‘dynamic’ to describe the portfolio updating that occurs when the assumption of a time-invariant cash-futures covariance matrix is relaxed. While these studies have established the potential benefits of relaxing this assumption, they have continued to rely on one-period-ahead variance and covariance forecasts. In this sense the extant ARCH/GARCH hedging literature has not truly utilized a ‘dynamic’ approach. Conversely, the DP hedging literature, while allowing for multi-period-ahead hedge ratio calculations, has not employed modern methods for estimating time-varying covariance matrices.

Our paper has attempted to bridge the gap between these two related yet distinct approaches. Importantly, we present a truly time-varying *and* dynamic hedging model by

combining DP recursion relations with GARCH time-series models. To illustrate the DP-GARCH framework empirical applications were reported for hypothetical cocoa and sugar merchandisers. Results indicate that for both commodities gains were achieved by using the DP-GARCH model relative to the standard GARCH approach, with greater rewards accruing for sugar. Importantly, while sample summary statistics suggest that modest improvements may be made by employing techniques that optimally update—like the DP-GARCH approach—we also isolate particular times that such a model outperforms simpler alternatives. Specifically, through a bootstrap experiment we isolate the number of times DP-GARCH approach dominates (in both a statistical and economic sense) OLS. We conclude that in periods of low volatility a trader may be just as well off following the simpler (and probably cheaper) OLS/SUR approach. The caveat is, of course, that the simpler approach may fail dramatically relative to more sophisticated techniques in periods of excessive volatility. Consequently, a DP-GARCH approach is likely, at times, to yield greater rewards in more volatile markets like sugar.

The framework developed here should not only be of interest as a contribution to the empirical hedging literature, but should also be of interest to risk managers. This said, more work is clearly needed. Specifically, the methodologies developed in this paper should be applied to other commodities (and assets) and for different hedging horizons. Also, transaction and liquidity costs should be included.¹⁵ Only then will it be possible to truly assess the value of the DP-GARCH approach.

Table 1.
Bivariate GARCH Models with Positive Definite Parameterization

$$\Delta \mathbf{p}_t = \mathbf{e}_t, \Delta \mathbf{p}_t = (c_t, f_t)^T$$

$$\mathbf{e}_t | \Omega_t \sim N(\mathbf{0}, H_t)$$

$$H_t = W^T W + A^T \mathbf{e}_{t-1} \mathbf{e}_{t-1}^T A + B^T H_{t-1} B$$

<i>Cocoa Portfolio</i>			<i>Sugar Portfolio</i>		
Parameter	Estimate	Std. Errors	Parameter	Estimate	Std. Errors
w_{11}	-18.381	2.647	w_{11}	-0.052	0.022
w_{12}	-10.710	2.466	w_{12}	0.090	0.035
w_{22}	-10.642	1.985	w_{22}	0.129	0.038
a_{11}	0.642	0.089	a_{11}	0.430	0.077
a_{12}	0.162	0.062	a_{12}	-0.467	0.090
a_{21}	-0.389	0.077	a_{21}	-0.151	0.067
a_{22}	0.095	0.058	a_{22}	0.771	0.080
b_{11}	0.679	0.069	b_{11}	0.762	0.057
b_{12}	-0.133	0.042	b_{12}	0.353	0.073
b_{21}	0.186	0.049	b_{21}	0.173	0.055
b_{22}	1.038	0.030	b_{22}	0.567	0.068
Log-likelihood -6481.71			Log-likelihood 1097.18		
Cash Equation			Cash Equation		
m_3	0.484		m_3	0.271	
m_4	1.760		m_4	2.213	
$Q(20)$	20.099 (0.452)		$Q(20)$	21.341 (0.377)	
$Q^2(20)$	26.985 (0.136)		$Q^2(20)$	24.252 (0.232)	
Futures Equation			Futures Equation		
m_3	0.403		m_3	0.011	
m_4	2.646		m_4	2.366	
$Q(20)$	13.249 (0.866)		$Q(20)$	28.385 (0.101)	
$Q^2(20)$	27.859 (0.113)		$Q^2(20)$	25.933 (0.168)	
Covariance Equation			Covariance Equation		
$Q^2(20)$	24.454 (0.223)		$Q^2(20)$	30.778 (0.060)	

Note: m_3 is sample skewness and m_4 is sample kurtosis. $Q(20)$ denotes the test statistic for twentieth-order serial correlation in the normalized residuals and $Q^2(20)$ is the test statistic for twentieth-order serial correlation in squared normalized residuals. Numbers in parentheses are asymptotic p -values.

Table 2.
Log-likelihood Estimates and Tests of Parameter Restrictions for the
Bivariate GARCH Commodity Models

$$\begin{aligned}\Delta \mathbf{p}_t &= \mathbf{m} + \mathbf{e}_t, \Delta \mathbf{p}_t = (c_t, f_t)^T \\ \mathbf{e}_t | \Omega_t &\sim N(\mathbf{0}, H_t) \\ H_t &= W^T W + A^T \mathbf{e}_{t-1} \mathbf{e}_{t-1}^T A + B^T H_{t-1} B\end{aligned}$$

Model/Test	Log-likelihood/LR Test	No. of Parameters/ Degrees of Freedom
<i>Cocoa Portfolio</i>		
No GARCH, $\mu = 0$	-7518.47	3
Diag A and B , $\mu = 0$	-6647.06	7
Full A and B , $\mu = 0$	-6481.71	11
Full A and B and $\mu \neq 0$	-6481.67	13
LR Test $A = B = 0$	2073.52 (0.000)	d.f. = 8
LR Test A and B Diagonal	330.79 (0.000)	d.f. = 4
LR Test $\mu = 0$	0.08 (0.961)	d.f. = 2
<i>Sugar Portfolio</i>		
No GARCH, $\mu = 0$	824.63	3
Diag A and B , $\mu = 0$	1042.45	7
Full A and B , $\mu = 0$	1097.18	11
Full A and B and $\mu \neq 0$	1096.87	13
LR Test $A = B = 0$	549.09 (0.000)	d.f. = 8
LR Test A and B Diagonal	109.47 (0.000)	d.f. = 4
LR Test $\mu = 0$	0.63 (0.730)	d.f. = 2

Note: Values in parentheses are asymptotic p -values. d.f. denotes degrees of freedom. All restrictions are imposed on parameters of the A and B matrices from the conditional covariance structure.

Table 3.
Descriptive Statistics for Hedge Ratios for Variance
Minimizing DP-GARCH, GARCH
and OLS/SUR Objectives, Four-Week Hedging Horizon

	Cocoa Portfolio				Sugar Portfolio			
In- Sample								
Hedge Model	b_{t-4}	b_{t-3}	b_{t-2}	b_{t-1}	b_{t-4}	b_{t-3}	b_{t-2}	b_{t-1}
DP-GARCH								
Avg.	-0.918	-0.920	-0.921	-0.923	-0.778	-0.780	-0.781	-0.783
SD	0.138	0.138	0.139	0.139	0.198	0.198	0.199	0.198
Min	-2.214	-2.128	-2.132	-2.136	-1.122	-1.125	-1.127	-1.129
Max	-0.064	-0.064	-0.065	-0.065	0.190	0.190	0.190	0.191
GARCH		b_{GARCH}				b_{GARCH}		
Avg.		-0.918				-0.778		
SD		0.005				0.198		
Min		-2.124				-1.122		
Max		-0.064				0.190		
OLS/SUR		$b_{\text{OLS/SUR}}$				$b_{\text{OLS/SUR}}$		
		-0.899				-0.723		
Out-of-Sample								
Hedge Model	b_{t-4}	b_{t-3}	b_{t-2}	b_{t-1}	b_{t-4}	b_{t-3}	b_{t-2}	b_{t-1}
DP-GARCH								
Avg.	-0.945	-0.947	-0.949	-0.950	-0.797	-0.801	-0.801	-0.801
SD	0.081	0.080	0.080	0.081	0.180	0.179	0.184	0.184
Min	-1.208	-1.210	-1.213	-1.219	-0.968	-0.974	-0.976	-0.978
Max	-0.792	-0.794	-0.796	-0.797	-0.027	-0.042	-0.039	-0.049
GARCH		b_{GARCH}				b_{GARCH}		
Avg.		-0.945				-0.797		
SD		0.081				0.180		
Min		-1.208				-0.968		
Max		0.792				-0.030		
OLS/SUR		$b_{\text{OLS/SUR}}$				$b_{\text{OLS/SUR}}$		
Avg.		-0.900				-0.727		
SD		0.002				0.002		
Min		-0.897				-0.723		
Max		-0.903				-0.730		

Note: The annualized discount rate, r , is 0.10. Avg. denotes sample average, SD is the corresponding standard deviation of the average of the hedge ratios, Min is the sample minimum and Max is the sample maximum. The GARCH hedge ratio b_{GARCH} represents the average hedge ratio that would be used by the trader over the four-week period, and is equal to the hedge ratio used at $t - 4$ by the DP-GARCH. The OLS/SUR hedge ratio, $b_{\text{OLS/SUR}}$, used each week is not, like the GARCH counterpart updated each week.

Table 4.
Summary Statistics for DP-GARCH, GARCH, and OLS/SUR
Portfolios, Four-Week Hedging Horizon.

	<i>Cocoa Portfolio</i>	<i>Sugar Portfolio</i>
	<i>In Sample</i>	
DP-GARCH		
Avg.	30228.77	2.248
SD	18736.00	0.051
Min	15688.70	0.833
Max	231871.40	13.406
GARCH		
Avg.	30461.37	2.501
SD	20359.31	0.060
Min	17147.21	1.002
Max	274274.50	12.938
OLS/SUR		
Avg.	30714.32	2.448
SD	20001.74	0.070
Min	16526.30	1.017
Max	224806.90	21.257
% Reduction in Variance of DP-GARCH Relative to:		
GARCH	0.769%	11.254%
OLS/SUR	1.606%	8.897%
	<i>Out-of-Sample</i>	
DP-GARCH		
Avg.	24211.25	2.055
SD	5049.20	0.742
Min	17327.74	1.120
Max	40431.00	4.628
GARCH		
Avg.	24211.13	2.227
SD	5196.19	0.765
Min	17148.44	1.355
Max	36393.43	4.432
OLS/SUR		
Avg.	23649.11	2.092
SD	4582.61	0.719
Min	16898.62	1.288
Max	37134.67	4.543
% Reduction in Variance of DP-GARCH Relative to:		
GARCH	-0.001%	8.370%
OLS/SUR	-2.322%	1.800%

Note: The annualized discount rate, r , is 0.10. (r was also adjusted to 0.01 and 0.05 and results were qualitatively unchanged). Avg. denotes sample average variance, SD denotes the corresponding standard deviation of the average, Min is the sample minimum and Max is the sample maximum. There are a total of 810 weekly hedging periods in sample and 57 out-of-sample.

Table 5.
Evaluating the Performance of the DP-GARCH and OLS/SUR
Portfolios: Results from the Bootstrap Experiment

	<i>Cocoa Portfolio</i>	<i>Sugar Portfolio</i>
Scenario	<i>OLS variance = DP-GARCH variance</i>	
Percentage of times	37.10%	33.70%
Avg. improvement	-	-
SD	-	-
Min	-	-
Max	-	-
Scenario:	<i>OLS variance < DP-GARCH variance</i>	
Percentage of times	47.20%	48.90%
Avg. improvement	32.28%	37.71%
SD	14.51	12.30
Min	10.64%	11.32%
Max	78.18%	73.99%
Scenario:	<i>OLS variance > DP-GARCH variance</i>	
Percentage of times	15.70%	17.90%
Avg. improvement	54.95%	95.92%
SD	46.45	98.68
Min	6.27%	18.54%
Max	212.96%	608.67%

Note: Percent of times represents: 1) Percentage of times that the DP-GARCH and OLS approaches are statistically equivalent (upper panel). 2) Percentage of times the OLS approach statistically outperforms the DP-GARCH approach (middle panel) and 3) Percentage of times the DP-GARCH approach statistically outperforms the OLS approach (lower panel). Avg. improvement denotes percentage improvement in variance reduction against the competing model (either DP-GARCH or OLS), SD denotes the corresponding standard deviation of that average improvement, and Min is the sample minimum and Max is the sample maximum improvement.

Figure 1. Weekly Sugar DP-GARCH (b_{t-4} , b_{t-3} , b_{t-2} , b_{t-1}) GARCH (b_{GARCH}) and OLS/SUR ($b_{OLS/SUR}$) Optimal Hedging Ratios

Weekly optimal hedging for the DP-GARCH, GARCH and OLS/SUR hedging models: January 1985 – August 2000. The GARCH hedge ratio b_{GARCH} represents the average hedge ratio that would be used by the trader over the four-week period. It is equal to the hedge ratio used at $t - 4$ by the DP-GARCH user as it is assumed that the GARCH user uses weekly data to form the hedge ratio to be applied at $t - 4$ and left in place until the commodity is purchased at the end of the horizon. The OLS/SUR hedge ratio, $b_{OLS/SUR}$, used each week is not, like the GARCH counterpart, updated each week and is constant over the entire time -frame.

Figure 2. Weekly Cocoa DP-GARCH ($b_{t-4}, b_{t-3}, b_{t-2}, b_{t-1}$) GARCH (b_{GARCH}) and OLS/SUR ($b_{OLS/SUR}$) Optimal Hedging Ratios

Weekly optimal hedging for the DP-GARCH, GARCH and OLS/SUR hedging models: January 1985 – August 2000. The GARCH hedger ratio b_{GARCH} represents the hedge average hedge ratio that would be used by the trader over the four-week period. It is equal to the hedge ratio used at $t - 4$ by the DP-GARCH user as it is assumed that the GARCH user uses weekly data to form the hedge ratio to be applied at $t - 4$ and left in place until the commodity is purchased at the end of the horizon. The OLS/SUR hedge ratio, $b_{OLS/SUR}$, used each week is not, like the GARCH counterpart, updated each week and is constant over the entire time -frame.

Figure 3. Comparing the Variance of the DP-GARCH and OLS Sugar Model

Large graph: Weekly optimal hedging variances for the DP - GARCH (with confidence bands) and OLS models over the period 1992 – 1998. Inserted graph displays the data for year 1995, and corresponds to the period of time that the variances are framed.

Figure 4. Comparing the Variance of the DP-GARCH and OLS Cocoa Model

Large graph: Weekly optimal hedging variances for the DP-GARCH (with confidence bands) and OLS models over the period 1990 – 1996. Inserted graph displays the data for year 1995, and corresponds to the period of time that the variances are framed.

Endnotes

1.This branch of the hedging literature also uses the term dynamic because the assumption of constant conditional variances and covariances is relaxed via ARMA-like specifications for conditional second moments. But in reality, these hedge ratios are typically *not* determined from a dynamic optimization set-up as applied to a multi-period hedging problem. In this sense GARCH-derived OHRs are not dynamic, but rather may be thought of as time varying.

2.ARCH is short for Autoregressive Conditional Heteroscedasticity while GARCH denotes Generalized ARCH.

3.There are a variety of methods for determining OHRs. Collins (1997) questions the validity of many of these approaches including the MV framework; he also reviews the empirical evidence about hedging behavior when evaluating competing models. In an attempt to explain ‘real life’ hedging behavior he concludes that the MV model is unlikely to fully capture a *short hedgers* true objective. Indeed, the vast majority of competing models evaluated did not pass such a test. The purpose of this study is to combine techniques (DP and GARCH) used in previous research (from a *long hedgers* perspective) that have employed the MV methodology. We therefore utilize the MV framework here, but we acknowledge that such an objective may not be truly optimal for all hedgers. Also, while it is plausible that commodity markets are unbiased, there is some empirical evidence suggesting that, at least for some markets, futures prices do exhibit some persistence and/or biasedness. See, e.g., Tong (1996); Rausser and Carter (1982); and Raynaud and Tessier (1984).

4.Myers and Thompson (1989) suggest that lags should be included in the regression procedure in order to include conditioning information. Pennings and Leuthold (2001) estimated hedging models with and without lags by using a procedure proposed by Britten-Jones (1999). Because

(1) Pennings and Leuthold (2001) found that including lags made no difference in their model; (2) because we wish to maintain the assumption of weak form efficiency; and (3) because we do not find any evidence of residual autocorrelation (see econometric results) we do not include lags in the mean regression equations. We therefore maintain the MV set up.

5. Most high-frequency asset price data are modeled in first difference form without any autocorrelation structure (the approach utilized here). As noted by Fama (1965) this martingale behavior is often interpreted as being consistent with weak form efficiency. This approach says nothing, however, about higher moments of asset returns, which are typically found to exhibit leptokurtosis. Consequently, some studies employ a distribution such as Student's t to account for the excess kurtosis (see Baillie and Myers, 1991). In this study we follow Haigh and Holt (2000) and maintain the normality assumption, accepting its somewhat stringent assumptions.

6. We focus on a weekly hedge for several reasons. First, our original motivation and hence model was based on conversations with a large food merchandiser who suggested that a weekly horizon was reasonable given their large amount of transactions (see Footnote 13). Second, as shown by Castelino (1992), Geppert (1995) and Pennings and Meulenberg (1997), the hedging effectiveness of this type of model is likely to increase with time, and so it is expected (although not explored here) that the model is likely to improve over a longer time horizon. We present the more conservative results based on this shorter (but more realistic) time-period. Extending the time frame is left for future research.

7. Indeed both the OLS and SUR models are nested within the GARCH framework (see econometric estimation results for full details).

8.To simplify the model, we follow Mathews and Holthausen (1991) in assuming that the hedger knows b_{t-1} at the initial trade date. If we did not make this assumption, b_{t-1} would be stochastic and additional variance and covariance terms would be involved. However, this assumption is not restrictive, as variance-covariance estimates based on historical relationships (like the GARCH framework) are relatively easy to forecast. Moreover, to operationally use a hedge ratio that is a function of future hedge ratios the merchandiser must forecast future ratios and thus consider them to be non-stochastic.

9.As pointed out by an anonymous reviewer complicated models that allow for more updating will incur higher transaction costs. As Haigh and Holt (2000) show, incorporating transaction costs reduces the appeal of more complicated approaches but not by enough to reduce the incentive to utilize these techniques, particularly when the hedger is confronted with large cash transactions. In the current analysis marginal transaction costs are unlikely to deter the representative hedger from employing a more sophisticated strategy. Such may not be the case, however, for a farmer or a smaller corporate entity. From an operational standpoint, incorporating transaction costs into a DP framework would require a numerical solution to the resulting recursion relations. Therefore, incorporating transaction costs is left for future research.

10.This time-line represents an in-sample comparison. In the out-of-sample analysis all models experience some variability. Therefore, unless the basis remains constant, the OLS hedge ratios will also change as new observations are added. However, the concepts underlying model comparisons remain the same. Results for OLS and basic GARCH (DP-GARCH) models are likely always going to differ, even if only slightly, due to the nature of estimation.

11.Baillie and Myers (1991) also undertook their analysis using a similar framework, thus employing a relatively parsimonious model. We base our structure for the time series generating process on residual diagnostic results. Because we find no evidence of residual autocorrelation in the markets studied here the MV framework may be appropriate.

12.As suggested by an anonymous reviewer, simulations were conducted by using three different levels of interest rates, 1 per cent, 5 per cent and 10 per cent. Varying the rate had no qualitative effect on model orderings; however, higher the interest rates eroded the performance of the DP-GARCH and GARCH approaches relative to OLS. Therefore, to provide conservative estimates of the relative performance of these models we employ the higher interest rate.

13.Therefore, in this application, and over this hedging horizon, substantial gains, in terms of risk reduction appear to be available by adopting the DP-GARCH relative to alternatives. Indeed according to a large U.S. based food manufacturer approximately \$75 million dollars and \$350 million dollars are spent annually on sugar and cocoa purchases, respectively, in their manufacturing processes (and the food manufacturer regards itself as a large corporate commodity hedger). According to this manufacturer, purchases and hedges are undertaken frequently. Given the large amount of purchases it is unlikely that extra transaction costs associated with more complicated methods would deter a large hedger. Transaction costs could, however, alter the risk management strategy of a smaller hedger.

14.Collins (2000) undertakes out-of-sample analysis on a variety of competing models. Similar to Collins(2000), we isolate the proportion of times our model outperforms potentially inferior models. In our analysis, however, we investigate the *conditions* that arise when a model better/worse than a simpler approach. While Collins (2000) suggests the naïve (i.e., one-to-one) hedge outperforms the risk-minimizing hedge most of the time, he did not investigate what

happens when the market exhibits short-run unpredictable volatility (i.e., erratic basis behavior) as we do here.

15.As shown by Pennings and Leuthold (2001), liquidity might also have an important bearing on hedging effectiveness, particularly in thinly traded markets. Researchers might therefore take this into account when developing optimal hedging strategies. For the commodities analyzed here trading volume between July 2000 and June 2001 was, for example, 1,408,945 cocoa contracts and 759,828 sugar contracts respectively. While these markets are not traded as heavily as larger financial contracts like the FTSE (9,033,641 contracts traded), they experience heavier trading volume than other commodities like wheat, which had a trading volume of 97,705 contracts. We therefore conclude that liquidity may not be a serious issue in our application. As pointed out by an anonymous reviewer, however, liquidity costs may be updated in manner similar to that for updating variances/covariances. Such an updating scheme and would likely be more important in thinly traded markets.

References

- Anderson, R.W. and Danthine, J.P. (1983). The time pattern of hedging and the volatility of futures prices. *Review of Economic Studies* 50: 249-266.
- Baba, Y., Engle, R.F., Kraft D.F. and Kroner, K.F. (1989). Multivariate simultaneous generalized ARCH. mimeograph, University of California, San Diego, Dept. of Economics.
- Baillie, R.T. and Myers, R.J. (1991). Bivariate GARCH estimation of the optimal commodity futures hedge. *Journal of Applied Econometrics* 6: 109-124.
- Berck, P. (1981). Portfolio theory and the demand for futures: the case of California cotton. *American Journal of Agricultural Economics* 63: 466-474.
- Binswanger, H.P. (1982). Empirical estimation and the use of risk preferences: discussion. *American Journal of Agricultural Economics* 64: 391-393.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedacity. *Journal of Econometrics* 31: 301-327.
- Britten-Jones, M. (1999). The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance* 54: 655-671.
- Castelino, M.G. (1992). Hedge effectiveness: basis risk and minimum-variance hedging. *The Journal of Futures Markets* 12: 187-201.
- Cecchetti, S.G., Cumby, R.E. and Figlewski, S. (1988). Estimation of the optimal futures hedge. *Review of Economics and Statistics* 70: 623-630.
- Collins, R.A. (1997). Toward a positive economic theory of hedging. *American Journal of Agricultural Economics* 79: 488-499. (May 1997).
- Collins, R.A. (2000). The risk management effectiveness of multivariate hedging models in the U.S. soy complex. *The Journal of Futures Markets* 20: 189-204.

- Dickey, D.A. and Fuller, W. (1981). Likelihood ratio tests for autoregressive time series with a unit root. *Econometrica* 49: 1057-1072.
- Ederington, L.H. (1979). The hedging performance of the new futures markets. *Journal of Finance* 34: 157-170.
- Engle, R.F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica* 50: 987-1008.
- Engle, R.F. and Granger, C.W.J. (1987) Co-integration and error correction: representation, estimation, and testing. *Econometrica* 55: 251-276.
- Engle, R.F. Kroner, K.F. (1995). Multivariate simultaneous generalized ARCH. *Econometric Theory* 11: 122-150.
- Fama, E.F. (1965). The behavior of stock market prices. *Journal of Business* 38: 34-105.
- Gagnon, L., Lypny, G.J. and McCurdy, T.H. (1998). Hedging foreign currency portfolios. *Journal of Empirical Finance* 5: 197-220.
- Garcia, P., Roh, J.S. and Leuthold, R.M. (1995) Simultaneously determined, time-varying hedge ratios in the soybean complex. *Applied Economics* 27: 1127-1134.
- Geppert, J.M. (1995). A statistical model for the relationship between futures contract hedging effectiveness and investment horizon length. *The Journal of Futures Markets* 15:507-536.
- Haigh, M.S. and Holt, M.T. (2000). Hedging multiple price uncertainty in international grain trade. *American Journal of Agricultural Economics* 82: 881-896.
- Kahl, K.H. (1983). Determination of the recommended hedging ratio. *American Journal of Agricultural Economics* 65: 603-605.

- Karp, L.S. (1988). Dynamic hedging with uncertain production. *International Economic Review* 29: 621-637.
- Kroner, K.F. and Sultan, J. (1993). Time varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis* 28: 535-551.
- Lence S., Kimle, K. and Hayenga, M. (1993). A dynamic minimum variance hedge. *American Journal of Agricultural Economics* 76: 1063-1071.
- Lin, J.W., Najand, M. and Yung, K. (1994). Hedging with currency futures: OLS vs GARCH. *Journal of Multinational Financial Management* 4: 45-67.
- Martin, L. and Garcia, P. (1981). The price-forecasting performance of futures markets for live cattle and hogs: a disaggregated analysis. *American Journal of Agricultural Economics* 63: 209-215.
- Martinez, S. and Zering, K.D. (1992). Optimal dynamic hedging decisions for grain producers. *American Journal of Agricultural Economics* 74: 879-888.
- Mathews, K.H. and Holthausen, Jr., D.M. (1991). A simple multiperiod minimum risk hedge model. *American Journal of Agricultural Economics* 73: 1020-1026.
- Myers, R.J. (1991). Estimating time-varying optimal hedge ratios on futures markets. *Journal of Futures Markets* 11: 39-53.
- Myers, R.J. and Thompson, S.R. (1989). Generalized optimal hedge ratio estimation. *American Journal of Agricultural Economics* 71: 858-861.
- Nance, D.R., Smith, C.W. and Smithson, C.W. (1993). On the determinants of corporate hedging. *The Journal of Finance* 48: 267-284.
- Park, T.H. and Switzer, L.N. (1995). Bivariate GARCH estimation of the optimal hedge ratios for stock index futures: a note. *Journal of Futures Markets* 15: 61-67.

- Pennings, J.M.E. and Meulenberg, M.T.G. (1997). Hedging efficiency: a futures exchange management approach. *The Journal of Futures Markets* 17: 599-615.
- Pennings, J.M.E. and Smidt, A. (2000). Assessing the construct validity of risk attitude. *Management Science* 46: 1337-1348.
- Pennings, J.M.E. and Leuthold, R.M. (2001). Introducing new futures contracts: reinforcement versus cannibalism. *Journal of International Money and Finance* 20: 659-675.
- Rausser, G.C. and Carter, C. (1983) Futures market efficiency in the soybean complex. *Review of Economics and Statistics* 65: 469-478.
- Raynaud, J. and Tessier, J. (1984). Risk premiums in futures markets: an empirical investigation. *The Journal of Futures Markets* 4: 188-209.
- Sephton, P.S. (1993). Optimal hedge ratios at the Winnipeg Commodity Exchange. *Canadian Journal of Economics* 26: 175-193.
- Shapiro, B.I. and Brorsen, B.W. (1988). Factors influencing farmer's decisions of whether or not to hedge. *North Central Journal of Agricultural Economics* 10: 145-153.
- Tong, W.H.S. (1996). An examination of dynamic hedging. *Journal of International Money and Finance* 15: 19-35.
- Vukina, T. and Anderson, J.L. (1993). A state-space forecasting approach to optimal intertemporal cross-hedging. *American Journal of Agricultural Economics* 75: 416-424.
- Yeh, S.C. and Gannon, G.L. (2000). Comparing trading performance of the constant and dynamic hedge models: a note. *Review of Quantitative Finance and Accounting* 14: 155-160.

Appendix

Derivation of the Cost and Variance Expressions at Time Period $t - 1$ and $t - 2$ and the Resulting Optimal Hedging Ratios.

The cost function facing the merchandiser at $t - 1$ is

$$Cost_t = -c_t + b_{t-1}(f_{t-1} - f_t), \quad (A1)$$

where c_t is the cash price in period t , f_t is the corresponding futures price in period t , and b_{t-1} is the proportion of the cash purchase hedge in period $t-1$. Therefore, the variance of cost, Var_{t-1} , may be written as

$$Var_{t-1}(Cost_t) = Var(c_t) + b_{t-1}^2 Var(f_t) + 2b_{t-1} Cov(c_t, f_t).$$

The first order condition for an extremum associated with the above is, after simplifying,

$$\frac{\partial Var_{t-1}(Cost_t)}{\partial b_{t-1}} = b_{t-1} Var(f_t) + Cov(c_t, f_t) = 0. \quad (A2)$$

Solving (A2) for the OHR, b_{t-1} , yields

$$b_{t-1} = -\frac{Cov(c_t, f_t)}{Var(f_t)}, \quad (A3)$$

the optimal hedging ratio to be used at $t-1$. At time $t-2$, the variance of cost is:

$$\begin{aligned} Var_{t-2}(Cost_t) &= Var_{t-2}(-c_t + b_{t-1}(f_{t-1} - f_t) + rb_{t-2}(f_{t-2} - f_{t-1})) \\ &= Var_{t-2}(c_t) + b_{t-1}^2 Var_{t-2}(f_{t-1}) + r^2 b_{t-2}^2 Var_{t-2}(f_{t-1}) \\ &\quad + b_{t-1}^2 Var_{t-2}(f_t) + 2rb_{t-2} Cov_{t-2}(f_{t-1}, c_t) \\ &\quad - 2rb_{t-2} b_{t-1} Var_{t-2}(f_{t-1}) + 2rb_{t-2} b_{t-1} Cov_{t-2}(f_{t-1}, f_t) \\ &\quad - 2b_{t-1} Cov_{t-2}(c_t, f_{t-1}) + 2b_{t-1} Cov_{t-2}(c_t, f_t) \\ &\quad - 2b_{t-1}^2 Cov_{t-2}(f_{t-1}, f_t). \end{aligned} \quad (A4)$$

where r is a discount factor. The first order conditions corresponding to the minimization of

(A4) with respect to b_{t-2} and b_{t-1} are:

$$\begin{aligned}\frac{\partial Var_{t-2}(Cost_t)}{\partial b_{t-2}} &= 2r^2 b_{t-2} Var_{t-2}(f_{t-1}) + 2r Cov_{t-2}(f_{t-1}, c_t) \\ &\quad - 2r b_{t-1} Var_{t-2}(f_{t-1}) + 2r b_{t-1} Cov_{t-2}(f_{t-1}, f_t) = 0.\end{aligned}\tag{A5}$$

$$\begin{aligned}\frac{\partial Var_{t-2}(Cost_t)}{\partial b_{t-1}} &= 2b_{t-1} Var_{t-2}(f_{t-1}) + 2b_{t-1} Var_{t-2}(f_t) \\ &\quad - 2r b_{t-2} Cov_{t-2}(f_{t-1}, f_t) - 2Cov_{t-2}(f_{t-1}, c_t) \\ &\quad + 2Cov_{t-2}(c_t, f) - 4b_{t-1} Cov_{t-2}(f_{t-1}, f_t) = 0.\end{aligned}\tag{A6}$$

Expressions (A5) and (A6) represents a system of two equations in the two unknowns b_{t-1} and b_{t-2} . Conveniently, given unbiased commodity markets the OHR at $t - 2$ is reduced to

$$b_{t-2} = -\frac{Cov_{t-2}(c_t, f_{t-1})}{r Var_{t-2}(f_{t-1})}.$$

This result holds because with unbiased markets—each series is represented by a martingale—so the futures price at $t-1$ may be represented as

$$f_{t-1} = f_{t-2} + \xi_{t-1}.$$

In this case the variance of f_{t-1} , taken in $t - 2$, is simply

$$Var_{t-2}(f_{t-1}) = E_{t-2}(\xi_{t-1})^2.$$

The futures price at time t can be written as $f_t = f_{t-1} + \xi_t$, and using the fact

$$f_t = f_{t-2} + \xi_{t-1} + \xi_t,$$

we have:

$$Cov_{t-2}(f_{t-1}, f_t) = E_{t-2}(\xi_{t-1}, \xi_{t-1} + \xi_t) = E_{t-2}(\xi_{t-1})^2$$

Therefore, from (A6) $2rb_{t-1}Var_{t-2}(f_{t-1}) = 2rb_{t-1}Cov_{t-2}(f_{t-1}, f_t)$, which obtains the hedge ratio represented by

$$b_{t-2} = -\frac{Cov_{t-2}(c_t, f_{t-1})}{rVar_{t-2}(f_{t-1})}$$

Using similar arguments it can also be shown that the hedging ratios at $t-3$ and $t-4$ are as presented in (5).