Hedging with Two Futures Contracts: Simplicity Pays*

Van Thi Tuong Nguyen†
Katelijne A.E. Carbonez‡
Piet Sercu§

Monday 21st June, 2010

*We gratefully acknowledge financial support from FWO and ICM, and help with the data from the USDA. Comments from Van's committee (Michael Brennan, Geert Dhaene, and Gunther Wuyts) have substantially improved earlier drafts. We also received useful comments from Klyuev Vladimir, Eric de Bodt, Michel Levasseur and other participants at the All China Economics International Conference (Hongkong), workshop at K.U. Leuven and the joint financial economics seminar ESC/Universite de Lille. All remaining errors are the authors' responsability.

†FWO PhD Fellow at Leuven School of Business and Economics, K.U.Leuven, Naamsestraat 69, B-3000 Leuven; +32 16 32 64 62; thituongvan.nguyen@econ.kuleuven.be

‡ICM PhD Fellow at Leuven School of Business and Economics, K.U.Leuven, Naamsestraat 69, B-3000 Leuven; +32 16 32 66 59; Katelijne.Carbonez@econ.kuleuven.be

§Leuven School of Business and Economics, K.U.Leuven, Naamsestraat 69, B-3000 Leuven; +32 16 32 67 56; Piet.Sercu@econ.kuleuven.be

Abstract

We propose to use two futures contracts in hedging an agricultural commodity commitment to solve either the standard delta hedge or the roll-over issue. Most current literature on dual-hedge strategies is based on a structured model to reduce roll-over risk and is somehow difficult to apply for agricultural futures contracts. Instead, we propose to apply a regression based model and a naive rules of thumb for dual-hedges which are applicable for agricultural commodities.

The naive dual strategy stems from the fact that in a large sample of agricultural commodities, De Ville, Dhaene and Sercu (2008) find that GARCH-based hedges do not perform as well as OLS-based ones and that we can avoid estimation error with such a simple rule. Our semi naive hedge ratios are driven from two conditions: omitting exposure to spot price and minimizing the variance of the unexpected basis effects on the portfolio values. We find that, generally, (i) rebalancing helps; (ii) the two-contract hedging rules do better than the one-contract counterparts, even for standard delta hedges without rolling-over; (iii) simplicity pays: the naive rules are the best one–for corn and wheat within the two-contract group, the semi-naive rule systematically beats the others and GARCH performs worse than OLS for either one-contract or two-contract hedges and for soybeans the traditional naive rule perform nearly as good as OLS. These conclusions are based on the tests on unconditional variance (Diebold and Mariano (1995)) and those on conditional risk (Giacomini and White (2006)).

JEL classification: G11, Q11, Q14
Key words: hedging strategy, hedge ratio, convenience yield

Introduction

This paper considers the problem of hedging a commodity commitment using two futures contracts rather than one. We find that two-contract hedges are useful not just to manage roll-over risk, an area where they are well established, but also for standard delta hedges, especially when hedging is done dynamically. In addition, we find that simple regressions tend to outclass sophisticated ones, and that a two-contract variant of the naive one-to-one hedge ratio typically does even better than simple regressions. In that sense, simplicity seems to pay.

In general, hedgers use two futures contracts either because they face a combination of two risks for each of which a separate hedge is available, or because the hedge instrument is driven by two sources of risk which the hedger wants to recombine in different proportions. An example where the hedger faces a combination of two separately hedgeable risks would be the scenario where, say, a European buyer of vegetable oils hedges the risk of the Euro price of soy oil via a CBOT contract for the dollar price risk, and a dollar/euro forward or futures for the exchange risk. A familiar example of the second type, where the complication stems from the presence of two risks in the hedge instruments, is the roll-over hedge. This occurs when the horizon of the hedger's commitment is longer than the lives of the available contracts or perhaps of the sufficiently liquid ones. The hedger then needs to roll-over the hedge when the shorter futures contracts mature and later-dated futures contracts start trading. In the familiar analysis, the roll-over introduces basis risk: we face uncertainty about the size of the difference between the final futures price of the expiring contract (which is close to the then prevailing spot price) and the futures price of the new contract at the time of the roll-over. However, if the basis of the new contract is driven by substantially the same factors as the basis of currently traded contracts, and if the relationship is known, we can approximately replicate the new contract already before it is traded. If, for example, the hedger wants a twelve-month contract but he/she can work with six- and nine-month versions only, he/she tries to best replicate the twelve-month contract using the available material.

Standard delta hedges, as defined in most of the literature, differ from roll-over hedges in that the exposed position expires before rather than after the hedge instruments. The usual approach is to adopt a static hedge and adjust the hedge ratio, for instance on the basis of regression analysis. But also here a potential role is present for a two-contract strategy. If the ideal contract has a life of, for instance, two months while only three- and six-month instruments are traded, one can still try to optimally replicate the desired hedge using the available contracts, exactly as in the case of a roll-over problem. This paper deals with both these applications, i.e. roll-over and delta hedge problems.

Setting up a hedging policy involves more, though, than choosing between single- or dual-contract hedges. At least equally crucial is the rule for setting the hedge ratio(s) and the related issue of static versus dynamic hedging. Most applications go for either a structured model, a regression model or a naive rule of thumb. A structured model is usually set in continuous time and prescribes a detailed dynamic hedging policy. Naive rules avoid both analysis and statistics. Examples are hedging one ton spot by one ton forward. Like regression hedges, the naive rules are often static. We discuss them in more detail below.

In a structured model we specify the equations of motion, we then derive a corresponding hedge policy and apply it dynamically, like delta hedging of options. On paper, i.e. within the model's assumptions, the hedge works perfectly. With a regression model, on the other hand, we go for a static minimum-variance hedge ratio, commonly ascribed to Ederington (1979) but already known to Stein (1961) and Johnson (1960). In a one-contract hedge, such a hedge ratio is given by the conditional slope coefficient of a regression of ΔS (the future change in the spot price to which one is exposed) on Δf (the change in the futures price that is being used as the hedge). If the second joint moments of spot and futures price changes are constant, this hedge ratio is best estimated by OLS on past data. Many refinements to this strategy have been proposed, including dynamic hedging strategies (which exploit either the time-varying volatility of the spot and futures returns–GARCH; see Baillie and Myers, 1991), or error-correction in the spot-futures basis–ECM (Kroner and Sultan, 1993), or the random-coefficients regression model (Bera, Garcia and Roh, 1997). With a naive rule of thumb, lastly, we hedge one to one (bushel by bushel), or hedge one unit of currency i by h units of currency j and set h equal to the current cross exchange rate. Results from performance races are mixed. Some studies conclude that among the single-contract hedges the naive rule performs worse than regression, while in the latter class GARCH, possibly combined with an error-correction-model, seems to do best (see the above references). However, in a large sample of agricultural commodities, De Ville *et al.* (2008) find that GARCH-based hedges do not perform as well as OLS-based ones, while in Sercu and Wu (2000) the rules of thumb beat regression hedges for exchange risks, especially for high-R^2 currency pairs.

In contrast to the wide literature on one-contract hedges, prior research on two-contract strategy in commodity markets[1] mostly focus on structural models set up for the specific purpose of reducing the roll-over risk (Gibson and Schwartz, 1990; Brennan, 1991; Brennan and Crew, 1995; Schwartz, 1997; Neuberger, 1999; Veld-Merkoulova and De Roon, 2003; and Bülher, Korn and Schöbel, 2004). Even though these strategies often do quite well in oil or metals markets, they may be more difficult to apply in the case of agricultural commodities. For example, agricultural contracts have fixed delivery dates rather than the fixed lives we see in currency, oil or metal forward markets. Worse, there are strong jumps in the basis at the time of the harvest, so it is hard to see how the basis could be a simple process of the type assumed in the roll-over literature. In this paper we accordingly propose to use two-contract hedging strategies for agricultural commodities and focus on regression models and naive rules of thumb, which are easily applied to agricultural futures contracts. For two-regressor minimum-variance models, we use OLS and GARCH to related changes in spot price (ΔS) to changes in first-nearest futures price (Δf_1) and second-nearest futures price (Δf_2). Below, we refer to them as the OLS-2 and GARCH-2 policies. The naive rules we consider set the sum of the two hedge ratios quite mechanically, close to one-bushel-for-one-bushel in fact, and then decide on the mix of the two futures contracts, on the basis of minimizing the variance of the hedged portfolio relating to the unexpected change in the convenience yield. We refer to this hedge policy as semi-naive rule or as naive-2. Our other contending rules are single-regressor OLS and GARCH, and the one-for-one rule of thumb for a one-contract hedge.

On the whole, our empirical results show that all hedges work rather well, regardless of the number of hedges and the rule for fixing the ratio(s). Still, dual hedges do work better than single ones, and simplicity pays: among dual hedges, the semi-naive rule works best, before OLS and then GARCH for corn and wheat. Only for soybeans, GARCH-2 comes out on top. The naive rules of thumb (semi and traditional) do well even though they are almost surely biased, relative to a regression-based strategy.[2] But while this bias is avoided by regression, any statistics-based hedge ratio will always involve an estimation error. In addition, these particular regressions use bad data. There is bid-ask bounce in the futures data, for instance, and there are synchronization problems. In addition, for agricultural futures contracts, the time to maturity and basis in the historic data base change all the time, while one ideally

[1] For fixed-income markets, two-factor models are far more dominant.

[2] Sercu and Uppal (1995), for instance, show that the cross-rate rule in currency cross-hedges emerges if the regression coefficient between the two currencies is assumed to be unity, which is generally an overestimation (Sercu, 2009).

should use data with the time to maturity and initial basis of the hedge problem at hand. For instance, if one wants a two-month hedge using a futures with initial life of three months and if, initially, the basis equals 15 cents, then one should ideally use data on 60-day changes in the spot price and in a futures price with an initial basis of 15 cents and initial life of three months. The standard regressions ignore these subtleties, and treat data with all sorts of lives and basis levels as equally relevant. The resulting estimation errors in the hedge ratio obviously harm the performance of the hedge in a direct way. An indirect effect to be added to the picture is their impact on transaction costs. Thus, estimation errors are potentially serious handicaps that plague regression approaches relative to the naive rules of thumb. Also simplicity and low transaction costs may help explain why real-world traders often prefer naive rules over theoretically better grounded hedging policies. Among its strong sides we also note that the naive rule easily handles the jump in basis around harvest and the ever-changing lives of contracts, and avoids reliance on a particular time-series model.

To sum up, the paper has the following research questions: (1) does a dual-hedge work better than a single-hedge for agricultural commodities? (2) does rebalancing help? (3) does the simple naive rules outperform a more complicated model based on GARCH or OLS, and (4) do the dual rules still work better when transaction costs are considered?

To answer these questions and to evaluate performance of competing strategies, we adopt the following test set-up: we hedge (i) either a long-term agricultural commodity commitment using short-maturity futures contracts (roll-over issue) or (ii) a short-term agricultural commodity commitment using long-maturity futures contracts (delta hedge issue). We consider 1, 2, 4, 6, 8, 10 and 12-week hedging, and our commodities are corn, soybeans and wheat. We apply both static hedges and dynamic versions with weekly rebalancing. Our testing period is from 1/1989 to 8/2007. The out-of-sample part runs from 13/12/2002 to 1/8/2007. We use daily futures price data for corn, wheat and soybeans for the March, May, July, September and December contracts (corn and wheat) and for the January, March, May, July, August, September and November contracts (soybeans). The futures price is the daily settlement price reported at The Chicago Board of Trade (CBOT). Data for inventory and storage cost are from CBOT (provided by the University of Illinois). We use prices of the nearest-to-maturity futures contract instead of cash prices (spot prices) in calculating convenience yield and in testing our models. The reason is that cash prices can differ dramatically between locations only 30 miles apart, while the nearest futures and the deferred futures are based on the same location and so avoid location issues. We do not include the maturity month of the contract into the analysis because the delivery date is not fixed precisely.[3] For the interest rate, we use the 3-month LIBOR rates. We use non-overlapping data to avoid bias in the hedging effectiveness. To compare the hedging performance of alternative strategies, we consider both the variance reduction for the portfolio and the utility gains after transaction cost. Comparisons of competing hedges are made unconditionally as well as conditionally, applying Diebold and Mariano (1995) and Giacomini and White (2006), respectively, to test the significance of the difference in the implied conditional variances and the expected utility gains of two strategies.

We conclude this introduction with more information on the results. First, in terms of variance reduction, the two-contract strategies tend to perform better than the one-contract counterparts even for delta hedges without rolling over: while GARCH-2 comes up abreast with GARCH — each wins in 9 cases out of 18 — OLS-2 beats OLS in 15 cases and the semi-naive rule outperforms the traditional naive version in 16 cases. Second, also for delta hedges rebalancing does help to reduce the variance. Third, we find that for corn and wheat the naive rules are in the lead: with rebalancing they end up first in 7 cases out of 12, and in 11 out of 18 cases without rebalancing, while OLS, the runner-up, wins just 4 times in each situation. The lone exception is soybeans with rebalancing, for which GARCH wins in all cases. Taking into account that the difference in variance reduction among strategies are not economically large, there certainly is no evidence here that hi-tech statistical techniques add value. Our simplicity-pays conclusion contradicts some earlier findings in the literature—which, it must be added, typically used much shorter time series. Finally, when transaction costs are considered, the above conclusions are reinforced: the naive rules result in the highest utility gains in all cases for corn and wheat. For corn and wheat, naive-2 leads 20 times out of 22, while the traditional naive rule comes up first in the remaining 2 cases, leaving no role for OLS- and GARCH-regression-based strategies. Again, the only exception is soybeans for which the one-contract hedges perform significantly better than the dual-hedges: OLS is the winner in all cases. However, in this case the traditional naive performs nearly as good as OLS strategy with very tiny difference in utility gains. Lastly, bringing up significance, it is worth noting that nearly all the statistically clear differences for either variance reduction or utility gains (if any) are in favor of the naive rules. There is only one case (one-week horizon, soybeans) of a

[3] As long as the model is linear in TTM, errors in $T-t$ equally affect all observations, whether TTM is large or small. But in the non-linear models that follow, errors in $T-t$ disproportionably affect observations near the end of the contract's life. For the sake of comparability, we omit the final month everywhere.

significant result against the naive rules.[4] Overall, all results confirm the better performance of the simplicity-pays rules (the naive rules) over the regression based strategies.

The remainder of this paper is structured as follows. In Section 1, we summarize the current literature on dual-hedge with roll-over hedging purpose. We describe the way to derive the semi-naive hedge ratios in Section 2. In Section 3, the evaluation method for the hedge performance is presented. Section 4 is about the empirical results. Section 5 concludes.

1 The literature on roll-over (dual) hedging

As mentioned in the introduction, most of the current literature on two-contract hedges focuses on reducing roll-over risk. Research on this topic, like Gibson and Schwartz (1990), Brennan (1991), Brennan and Crew (1995), Schwartz (1997), Neuberger (1999), Veld-Merkoulova and De Roon (2003) and Bülher, Korn and Schöbel (2004)—all focusing on crude oil, orange juice and lumber—demonstrates that the problem of hedging error, when rolling over the short-dated futures contracts, can be reduced by holding, at every moment, positions in two futures contracts.

To model the term structure of commodity prices, Gibson and Schwartz (1990) and Brennan (1991) assume two factors, the spot price and the convenience yield that follow a joint diffusion process. By imposing the usual no-arbitrage condition, they obtain a partial differential equation which all contingent claims must satisfy. Provided that there are two futures contracts at each moment, any contingent claim can be valued and hedged perfectly. This approach is powerful in terms of creating a consistent framework for pricing and hedging all contingent claims. However, it has certain drawbacks if the aim is to apply it in agricultural markets. One is the assumption that all futures are priced with regard to each other via just two factors. While a one-factor convenience-yield process is perhaps not a bad approximation for minerals and metals, this is less likely to apply for agricultural commodities where the futures price also depends on whether the contract matures before or after the harvest time and where, accordingly, two distinct sources of risk are active, like demand and the size of the next harvest.

Veld-Merkoulova and De Roon (2003) use a one-factor term-structure model of the conve-

[4] A closer look at the sample characteristics reveals that for soybeans the basis is most highly correlated with the price. Thus, basis risk is well picked up by a one-contract hedge and the remaining risk is so small that the benefit from further reducing is usually wiped out by transaction costs.

nience yields and futures contracts for two different maturities to construct a hedging strategy. They study the market for oil forwards, which, like currency forwards, have fixed lives (e.g. always 30 or 60 days, regardless of when the trade is initiated) rather than fixed delivery dates (e.g. the third Wednesday of March). Given that this approach also requires that a new contract with identical time-to-expiration is listed periodically, it cannot directly be applied to agricultural commodity markets.

Neuberger (1999) proposes to solve the rollover problem by assuming that the price of a newly listed long-dated futures contract is a linear combination of the prices of the contemporaneous shorter-dated contracts. This approach is close to our naive rule and yields good results in hedging long-term exposure in the crude oil market. However, again, this approach cannot directly be applied to agricultural commodities because it also requires fixed time-to-maturity regressors (i.e. each month the newly listed futures contracts must have the same standard time to expiration): agricultural futures markets offer fewer contracts and new contracts are not opened at equally spaced dates.

Therefore, even if all of these models work well, they are rather difficult to apply for agricultural commodity futures contracts. So, one of the motivations for our research is to find a dual-hedge rule that can be applied for agricultural commodities and to test whether two-contract hedges also perform better than the one-contract counterparts.

2 Specification of the dual hedges

As usual, we assume that at time t, a hedger has perfect knowledge about the size and the maturity either of his/her future inventory holdings or of his/her future commitment in a certain commodity. Our purpose is to find a good strategy to hedge against the possible changes in the market value of that commodity investment or commitment. More specifically, we want to determine the amount of each futures contracts per unit of the given spot position to minimize the uncertainty about the hedged flow.

Current time is denoted as t, and the futures contract expires at T; $C_{t,T}$ refers to the current cost of storage in dollar terms per period; $r_{t,T}$ is the one-period risk-free rate of return in the financial markets; S_t and $f_{t,T}$ denote spot and futures prices, respectively, and $Y_{t,T}$ denotes the convenience premium paid for holding inventory from t to T rather than buying forward for delivery at T. Then, futures prices are related to spot prices through the costs of storage, time value, and the dollar convenience yield:

$$f_{t,T} = S_t \cdot [1 + r_{t,T} \cdot (T - t)] + C_{t,T} \cdot (T - t) - Y_{t,T}. \tag{1}$$

2.1 Rule 1 for the semi-naive hedge: matching spot-price exposure

The hedger wants to use two futures contracts to hedge his/her anticipated future purchases. Denote $h_{1,t}$ and $h_{2,t}$ as the hedge ratios of the nearest-to-maturity and second-nearest-to-maturity futures contracts, maturing at time T_1 and T_2 ($T_2 > T_1$) respectively. Suppose the hedging horizon is m periods. The proposed hedge ratios depend on whether there is rebalancing or not.

If the hedge horizon is shorter than the maturities of two futures ($t + m < T_1 < T_2$), it is possible (but not necessarily optimal) to keep the hedged portfolio from time t until the end of horizon $t + m$ without rebalancing. The portfolio value at the end of horizon ($V_{p,t+m}$) is shown below, followed by the expression for its change.[5] We use Δx_{t+m} to denote $x_{t+m} - x_t$:

$$\begin{aligned} \Delta V_{p,t+m} &= (S_{t+m} - S_t) - \sum_{i=1}^{2} h_{i,t} \cdot (f_{t+m,T_i} - f_{t,T_i}), \\ &=: \Delta S_{t+m} - \sum_{i=1}^{2} h_{i,t} \cdot \Delta f_{t+m,T_i}. \end{aligned} \tag{2}$$

From Equation (1), assuming that the interest rate and the storage cost have not changed, the futures price change from t to $t + 1$ is given by the two prime sources of uncertainty, changes in spot prices and convenience yields:

$$\begin{aligned} \Delta f_{t+1,T} &= [S_{t+1} \cdot [1 + r_{t,T} \cdot (T - t - 1)] + C_{t,T} \cdot (T - t - 1) - Y_{t+1,T}] \\ &\quad - [S_t \cdot [1 + r_{t,T} \cdot (T - t)] + C_{t,T} \cdot (T - t) - Y_{t,T}] \\ &= \Delta S_{t+1} \cdot [1 + r_{t,T} \cdot (T - t - 1)] - S_t \cdot r_{t,T} - C_{t,T} - Y_{t+1,T} + Y_{t,T}. \end{aligned} \tag{3}$$

Below, we substitute this into (2), we denote the one-period risk-free rate of return by r_t

[5] In keeping with the literature, all cash flows from marking to market are treated as if they were concentrated at the end. This will not systematically affect the cashflow as long as the daily price changes are independent of the interest rates, which for commodities is not an unreasonable assumption.

instead of $r_{t,t+1}$, regroup, and finally compress the notation:

$$\begin{aligned}
\Delta V_{p,t+m} &= \Delta S_{t+m} - \sum_{i=1}^{2} h_{i,t} \cdot \{\Delta S_{t+m} \cdot [1 + r_t \cdot (T_i - t - m)] - S_t \cdot r_t \cdot m - C_{t,T_i} \cdot m - \Delta Y_{t+m,T}\}, \\
&= \Delta S_{t+m} \left[1 - \sum_{i=1}^{2} h_{i,t} \cdot [1 + r_t \cdot (T_i - t - m)]\right] \\
&\qquad + \sum_{i=1}^{2} h_{i,t} \cdot [\Delta Y_{t+m,T_i} + S_t \cdot r_t \cdot m + C_{t,T_i} \cdot m] \\
&=: \Delta S_{t+m} \cdot A_t + \sum_{i=1}^{2} h_{i,t} \cdot [\Delta Y_{t+m,T_i} - B_{t,T_i}] \qquad (4)
\end{aligned}$$

with $A_t := [1 - \sum_{i=1}^{2} h_{i,t} \cdot (1 + r_t \cdot (T_i - t - m))]$ and $B_{t,T_i} := -S_t \cdot r_t \cdot m - C_{t,T_i} \cdot m$, the total cost of carry over horizon $[t, T_i]$.

Our first heuristic solution for a semi-naive hedge is to first and foremost eliminate all exposure to the prime source of variability, the spot price. Thus, the hedge ratios must satisfy $A_t = 0$ or

$$\sum_{i=1}^{2} h_{i,t} \cdot (1 + r_t \cdot (T_i - t - m)) = 1. \qquad (5)$$

This differs somewhat from a bushel-for-bushel naive hedge, which would have said that one bushel should be hedged by, in total, one bushel sold in the futures markets. The correction here takes into account time value for the remaining life of each of the futures contracts at the time the hedges are liquidated. Numerically, however, this comes close to a one-for-one rule for the sum of the two hedge ratios. The above relation is also the secret weapon of the rule: it gives us a hard equation for the sum of the hedge ratios. While the constraint is likely to be suboptimal (because, unlike a regression-based strategy, it ignores covariances with the convenience yield), it requires no estimation and is, therefore, free of estimation error.

2.2 Rule 2 for semi-naive policy: minimizing the residual variance

To get a unique solution for the two hs separately, we need a second condition. In keeping with the spirit of regression hedges we minimize the residual variance, but we try to improve on standard regression by working with a conditional mean for the convenience yield rather than the unconditional distribution that would have been used in a standard static regression hedge. At t, indeed, we do observe an initial Y, and we adopt a simple regression forecast for the expected change in the yield. At each t, past data provide us estimates for the conditional or residual (co)variances of the two yields. Thus, our second condition is to minimize the conditional variance, at t, for horizon m,

$$h_{1,t}^2 \cdot \mathrm{var}_{t,t+m}(e_{T_1}) + h_{2,t}^2 \cdot \mathrm{var}_{t,t+m}(e_{T_2}) + 2h_{1,t}h_{2,t} \cdot \mathrm{cov}_{t,t+m}(e_{T_1}, e_{T_2}), \tag{6}$$

subject to the condition in Equation (5). Simplifying the subscripts for the moments from $\mathrm{var}_{t,t+m}(e_{T_i})$ to $\mathrm{var}_t(e_{T_i})$, we get (see Appendix I for a proof):

$$\begin{cases} h_{1,t} = \frac{I_t[J_t \cdot \mathrm{var}_t(e_{T_2}) - \mathrm{cov}_t(e_{T_1}, e_{T_2})]}{\mathrm{var}_t(e_{T_1}) - 2J_t \cdot \mathrm{cov}_t(e_{T_1}, e_{T_2}) + J_t^2 \cdot \mathrm{var}_t(e_{T_2})}, \\ \\ h_{2,t} = I_t - J_t \cdot h_{1,t}. \end{cases}$$

with $I_t = \frac{1}{1+r_t\cdot(T_2-t-m)}$ and $J_t = \frac{1+r_t\cdot(T_1-t-m)}{1+r_t\cdot(T_2-t-m)}$.

The above considers a static delta hedge, never rebalanced. Now, suppose we want to rebalance the portfolio every period during the hedging horizon m, each period lasting k year where $k \ll 1$ and $t + k < T_1$, T_2.[6] In this case, the total portfolio value t to $t + m$ is equal to the cumulate change in the portfolio value realized in every subperiod.[7]

$$\Delta V_{p,t+m} = \sum_{j=1}^{m/k} (S_{t+kj} - S_{t+k(j-1)}) - \sum_{j=1}^{m/k} \sum_{i=1}^{2} h_{i,t+k(j-1)} \cdot (f_{t+kj,T_i} - f_{t+k(j-1),T_i}) \tag{7}$$

To reduce the unexpected variability, for every period we act in the same way as for the static hedge, except that the local horizon now is k rather than m. Such myopic period-by-period hedging is inspired by the near-absence of autocorrelation in the spot prices over short horizons and the relatively low variance of the changes in the convenience yields. Thus, for $j = 1$ to m/k,

$$\begin{cases} h_{1,t+k(j-1)} = \frac{I_{t+k(j-1)}[J_{t+k(j-1)} \cdot \mathrm{var}_{t+k(j-1)}(e_{T_2}) - \mathrm{cov}_{t+k(j-1)}(e_{T_1}, e_{T_2})]}{\mathrm{var}_{t+k(j-1)}(e_{T_1}) - 2J_{t+k(j-1)} \cdot \mathrm{cov}_{t+k(j-1)}(e_{T_1}, e_{T_2}) + J_{t+k(j-1)}^2 \cdot \mathrm{var}_{t+k(j-1)}(e_{T_2})}, \\ \\ h_{2,t+k(j-1)} = I_{t+k(j-1)} - J_{t+k(j-1)} \cdot h_{1,t+k(j-1)}. \end{cases}$$

We refer to this rule as semi-naive rule as it is half-way between a regression hedge and the standard rule of thumb. Like regression it does minimize variance; it even employs a regression model for the expected convenience yields at time $t+m$. However, it has the potential advantage

[6] One potential advantage of rebalancing is that we can apply new hedge ratios when new information becomes available. Relatedly, in the semi-naive hedge a series of one-week conditional forecasts is more precise than one unconditional ten-week prediction. Moreover, some rebalancing is actually inevitable in the case of a long-term commitment that needs rolling-over.

[7] This again ignores time value of the marking-to-market cash flows, as standard in this literature.

that it confines statistics to a well-specified area, *viz.* managing convience-yield risk without allowing any spill-overs to the spot-risk part of the hedge problem.

As $Y_{t,T_2} > Y_{t,T_1}$ (unless possibly if there is no harvest during the life of two contracts), this rule tends to produce values for h_1 above unity, and values for h_2 below zero, which is what one would expect. Given that, by assumption, the horizon m is shorter than T_1 and T_2,[8] the hedge-ratio pattern plausibly means that the combination of the two contracts seeks to back-extrapolate the contracts' properties towards a horizon shorter than T_1, like $t+m$. The soundness of this intuition is confirmed by average OLS- and GARCH-based hedge ratios, which do follow that pattern.

2.3 Forecasting the convenience yield

For the conditional variance/covariance of e_{T_i} we need an expected value for Y_{T_i}. In order to forecast the convenience yield, we follow Carbonez, Nguyen and Sercu (2008) (CNS). They show, analytically as well as empirically, that the convenience yield of agricultural commodities like corn, soybeans and wheat is well approximated by the product of time to maturity and a function of current scarcity, $\phi(x, S, ...)$, involving e.g. inventories x and/or the spot price S. In addition, if there is a harvest during the contract's life, a new term is introduced that involves a product of a similar function of scarcity, $\psi(x, S, ...)$ and the timespan from harvest to delivery. Finally, it turns out that published inventory data are lagging behind reality and add very little to the statistical fit, once the (CPI-deflated) price is in the model. Thus, $\phi()$ and ψ are close to linear in S, and we get a very tractable equation,

$$Y(S, t, T, T_h) = \alpha_t + [\beta \cdot (T-t) + \zeta \cdot \max(T - T_h, 0)]S_{t-1} + e_t, \tag{8}$$

with an intercept that seasonally depends on the delivery month and on the ongoing month interacted with time to maturity and time beyond harvest, if any. Table 6 summarizes the estimation results for the CNS price model for corn, soybeans and wheat. This estimation is with weekly data and for the initial-estimation period, 1989-2002.

[8]This is *a fortiori* the case with rebalancing, when the interim horizon shrinks to one week.

2.4 Regression-based Minimum-variance hedges

As benchmarks for the naive hedge we apply standard regression hedges: OLS and OLS-2, and GARCH and GARCH-2. For the first pair we regress the following equation with OLS:

$$P_{S,t} = \alpha - \beta P_{f_{T_1},t} - \gamma P_{f_{T_2},t} + \epsilon$$

with hedge ratios defined as: $h_1 = \beta$ and $h_2 = \gamma$; $P_{S,t} = \frac{\Delta S_t}{S_{t-1}}$; $P_{f_{T_i},t} = \frac{\Delta f_{t,T_i}}{f_{t-1,T_i}}$; and $V_{p,t+1} = \Delta S_{t+1} - \sum_1^2 h_i \Delta f_{t+1,T_i}$. For CCC-GARCH-2, we estimate CCC-GARCH(1,1) for the following system:

$$\begin{cases} P_{S,t} &= \alpha_1 + \beta_1 P_{S,t-1} + \gamma_1 P_{f_{T_1},t-1} + \epsilon_1 \\ P_{f_{T_1},t} &= \alpha_2 + \beta_2 P_{S,t-1} + \gamma_2 P_{f_{T_1},t-1} + \epsilon_2 \\ P_{f_{T_2},t} &= \alpha_3 + \beta_2 P_{S,t-1} + \gamma_3 P_{f_{T_2},t-1} + \epsilon_3 \end{cases} \tag{9}$$

with hedge ratios defined as the two-regressor slope coefficients:

$$h_{i,t} = \frac{\text{var}_t(f_{t+1,T_j})\text{cov}_t(P_{S,t+1}, P_{f_{T_i},t+1}) - \text{cov}_t(P_{f_{T_i},t+1}, P_{f_{T_j},t+1})\text{cov}_t(P_{S,t+1}, P_{f_{T_j},t+1})|I_t}{\text{var}_t(P_{f_{T_i},t+1})\text{var}_t(P_{f_{T_j},t+1}) - \text{cov}_t(P_{f_{T_i},t+1}, P_{f_{T_j},t+1})^2|I_t} \tag{10}$$

with $i = 1, j = 2$ and $i = 2, j = 1$. The GARCH regressions are of the constant conditional correlation (CCC) type. Too often, the more general BEKK version fails to converge.

We should note that also in the statistics-based hedges the role of the second contract in the two-contract rule seems to partially neutralize the convenience factor in the nearest hedge: we see a weight exceeding unity for the nearest-dated contract, and a negative weight for he further-dated one, the one with the higher convenience yield. One consequence is that the two-contract rule will be less useful, either when the convenience yields of two futures contracts are highly correlated and the difference between them is just a tiny or when they have high correlations with the spot price. In such cases, either the convenience yield risk is very small or most of risks are well picked up by a simple one-regressor hedgee already, and thus the role of the second contract or the role of eliminating basis risk is less useful. This may be even more so when the benefit from reducing basis risk must be traded off against the high transaction cost for two-contract hedges.

3 Evaluating the conditional hedging effectiveness

For each hedging strategy and hedge period $(t, t + m)$ we can compute a realized cash flow. We then compute squared cash flows. There is, of course, substantial variation over time in the spot and the closely related futures price, inducing a good dose of heteroscedasticity into the time series of hedged cash flows. Accordingly, for the purpose of economic and statistical evaluation we rescale each realized cash flow by dividing it by the spot price at t.[9] Therefore, the variance of the scaled cash flow is:

$$VSC := \frac{\sum_{t=t_1}^{T_N-m}\left(\frac{\Delta V_{p,t+m}}{S_t} - \overline{\frac{\Delta V_{p,.+m}}{S_t}}\right)^2}{T_{out}}, \tag{11}$$

with T_{out} the number of observations in the out-of-sample testing period for the particular hedging horizon.

To judge the relative performance of two competing hedging strategies that have a risk-minimization objective, normally one compares the percentage reduction between the variances of the hedged and the unhedged portfolio returns. The strategy that provides the highest relative variance reduction is then deemed to be the best one. However, this method suffers from two major problems. First, being an unconditional measure, it is not adequate when the objective is to evaluate a dynamic strategy resulting from the minimization of the conditional portfolio return variance, as done, for instance, under GARCH. If variances change, a statistical test that assumed constant variance is hard to justify. A second reason why it is not sufficient to simply check the unconditional performance is that a rule that does well on average may still have a poor conditional relative hedging performance at particular moments. An unconditional criterion would make sense only if, once the hedger selects a particular method, like OLS, he/she is committed to using it for ever. Now, OLS may do well on average, but there may still be days when, conditionally, GARCH is expected to do better because the conditional variance of its hedge happens to be lower on that particular day. Accordingly, a conditional test has, as its null, that at any moment the conditional variances of two contending strategies are equal. If this null is not rejected, there is no statistically convincing ground that switching pays off.

In this paper, besides comparing the unconditional variances of portfolio returns from alternative hedging strategies, we apply the Diebold and Mariano (1995) (DM) and Giacomini and White (2006) (GW) statistical test to unconditionally and conditionally compare the conditional portfolio return variances implied by competing hedging strategies. The DM test is still an unconditional test, but it does take into account changing variances. In that sense, the test is complementary to the standard one and more in line with the idea of changing uncertainty. In addition, the test goes beyond a simple comparison of variances in that it allows inference

[9] Within a given hedge period with rebalancing, the cashflows are of course accumulated without rescaling: dollar price changes are additive, returns are not.

too. To measure conditional hedging effectiveness, the conditional portfolio return variance obtained from each strategy is compared. Denote by $cv_t(t^1)$ and $cv_t(t^2)$ the squared demeaned out-of-sample portfolio return obtained by hedging strategy t^1 and t^2 respectively. Andersen and Bollerslev (1998) and Diebold and Lopez (1996), among others, note that these squared demeaned portfolio returns are unbiased estimates of the true conditional variance. Let the difference in the squared demeaned returns be $\mathrm{d}v_t \equiv cv_t(t^1) - cv_t(t^2)$. The DM method is an unconditional test of the null hypothesis of equal conditional hedging effectiveness. Giacomini and White (2006) propose a conditional test. This GW test is large-sample, so we do not apply the GW test for horizons beyond one week because of insufficient observations. Both tests are described in Appendix II.

4 Data and results

4.1 Test procedure and descriptive statistics

4.1.1 Test set-up

As inventory data used to be released once a week only, the price data are sampled for the same days as those for which the inventory is known. Our total weekly sample, from 2/1/1989 to 1/8/2007, contains 969 observations for each commodity. Simple returns are always calculated from the same contract. This means that, if at time $t+1$ the spliced-together series is jumping from a shorter contract to a longer contract, the futures return at this time is computed from two prices for the longer contract, while the one at time t is computed from the last two prices for the shorter contract.

As mentioned and motivated in the introduction, we use the nearest-to-maturity futures prices to stand in for the spot prices, and thus the second-nearest-to-maturity and the third-nearest-to-maturity series are used to test the strategies. We test the rule for hedge horizons of one, four, six, eight, ten, and twelve weeks, with either weekly rebalancing or no rebalancing. To avoid statistical issues in assessing the confidence intervals for hedging effectiveness and so on, we use non-overlapping data. We divide the sample into two subperiods: the initial estimation period and the (out-of-sample) testing period. The initial estimation period (T_{in}) is from 2/1/1989 to 7/12/2002–about three fourths of the sample, a common procedure. The out-of-sample (T_{out}) testing period is from 13/12/2002 to 1/8/2007. With the non-overlapping method, this results in 727, 175, 113, 83, 65, and 52 observations in the initial estimation for the one-, four-, six-, eight-, ten-, and twelve-week hedging horizons respectively. The corresponding numbers are 242, 67, 48, 39, 32, and 28 for the out-of-sample testing. As the name suggests, the initial-estimation observations are used to estimate or calculate the hedge ratio for the first hedge experiment in the test period. Then, at time $t = T_{in} + 1$, we add the new observation realized at that time to the sample. We delete the first observation from the sample and we re-estimate the model to obtain the next hedge ratio, to be used at time $T_{in} + 1$. We continue repeating this process until we reach the end of the test data. With this moving window technique, the sample size for the estimation remains constant.

Relating to the GARCH based regression model, we choose the CCC-GARCH model as our representative for the GARCH class of models. CCC offers a simple way to avoid a possible non-positive semi-definite conditional variance-covariance matrix resulting from VECH-GARCH. We did experimented with the theoretically superior BEKK-GARCH variant but we sometimes ran into a singular-matrix problem. The moving window technique is also adopted for OLS and GARCH.

The strategies tested are the single-hedge strategies—*viz* the traditional naive one-to-one hedge, simple regression (OLS), and a CCC-GARCH—and the contending dual-hedge variants: the semi-naive rule, OLS-2 and CCC-GARCH-2.

4.1.2 Descriptive statistics of spot prices and convenience yields

The data sources were described in the introduction. Table 1 summarizes unconditional variances and correlations of the scaled weekly first differences of spot prices and the two convenience yields. The last column in the table compares the means of the scaled first difference of the two convenience yields. It is clear from the table that, while the spot-price variance for soybeans is somewhat below that for corn and wheat, the variances of its convenience yields are about twice those for corn and wheat. For soybeans, we also observe the highest correlation between spot price and convenience yields, that is, the (high) risk of the convenience yields can be picked up relatively well by a single hedge that also covers the spot price. Furthermore, for soybeans we also see the highest correlation between the two yields (0.934) compared with corn and wheat, and the mean difference is also smallest (only 5.7e-4 %). This means that the two hedge contracts offered essentially the same cash flows. As a result, hedge ratios that try to exploit the small differences might pick up mostly noise. All this can harm the effectiveness of the second contract in hedging for soybeans. For wheat, the correlation between spot price and the second-contract convenience yield is not much lower than for soybeans but, at least, the difference between two yields is half as large again as that of soybeans. In short, we expect

Table 1: Descriptive statistics of spot prices and convenience yields, out-of-sample

Commodity	$\sigma^2_{\frac{\Delta S_t}{S_t}}$	$\sigma^2_{\frac{\Delta Y_{t,T_1}}{S_t}}$	$\sigma^2_{\frac{\Delta Y_{t,T_2}}{S_t}}$	$\rho_{\frac{\Delta S_t}{S_t}\frac{\Delta Y_{t,T_1}}{S_t}}$	$\rho_{\frac{\Delta S_t}{S_t}\frac{\Delta Y_{t,T_2}}{S_t}}$	$\rho_{\frac{\Delta Y_{t,T_1}}{S_t}\frac{\Delta Y_{t,T_2}}{S_t}}$	$(\frac{\Delta Y_{t,T_1}-\Delta Y_{t,T_2}}{S_t})$
Corn	0.00160	0.00006	0.00011	0.115	0.292	0.875	0.00076
Soybeans	0.00131	0.00013	0.00021	0.387	0.459	0.934	0.00057
Wheat	0.00167	0.00006	0.00014	0.283	0.458	0.791	0.00108

Note: In this table, σ^2_x is weekly variance of x, and ρ_{xy} is correlation between x and y.

that the role for two-contract hedges is less promising for soybeans comparative to corn and wheat, especially when the transaction cost is considered.

4.2 Out-of-sample results for unconditional performance

4.2.1 Out-of-sample hedging performance with rebalancing: naive-2 pays

Table 2 reports the unconditional measures of relative out-of-sample hedging performance, as summarized by the percentage variance reduction of alternative hedging strategies over the no hedging one. Results with weekly rebalancing are in the upper panel while results without rebalancing are in the lower panel. The traditional naive and semi-naive strategies are reported in the first and second rows of each panel, respectively. To indicate the relative performance, we use boldface to indicate the best (highest) number.

In this subsection we discuss the results with rebalancing. At this first stage, we let the users rebalance every week. Later we have them weigh the cost of trading against the degree of sub-optimality of the outstanding hedge position, so that they can waive their option to rebalance at least some of the time.

First, it is clear that all hedging strategies considered here provide substantial (and, in fact, quite similar) variance reduction over the no hedging one, especially for long term hedging horizons. Second, and more central to the paper, the dual-hedge policies do systematically better than single-contract solutions, even for delta hedges when the hedging horizons (one or four weeks) are shorter than the maturities of hedged futures contracts.[10] For clean comparison, we contrast the single- and dual-hedge results per strategy. Our semi-naive rule beats the traditional naive rule in 16 cases out of 18 (i.e. for three commodities with 6 hedging horizons

[10] Recall we do not include the maturity month of futures contracts in our analysis, so maturities on the hedge side are at least 4 weeks.

each), while for the simple regressions OLS-2 performs better than OLS 15 times. Only for the more sophisticated estimates there is no clear gain: CCC-GARCH-2 offers higher variance reduction than CCC-GARCH in 9/18 cases (of which 5 are for soybeans). The third and, to practitioners, most interesting result is that among dual hedging rules, the semi-naive rule systematically beats the others, coming up ahead of OLS-2 10 out of 18 times and ahead of GARCH-2 11 times. When comparing the scores with all strategies considered here, whether single- or dual-hedge, for corn and wheat the semi-naive rule wins 5 out of 12 cases; OLS-2 wins 3 cases, traditional naive wins 2 cases and GARCH-2 and GARCH each wins 1 case. Only for soybeans, then, GARCH-2 does a good job.

In short, we find that regression does not really help for corn and wheat, whether single- or dual. Moreover, if one nevertheless goes for regression, the simpler variant seem to do best again. In fact, our fifth finding is that, for corn and wheat, GARCH wins against OLS in only 3 out of 12 cases for the single hedge and in 1 case for two-contract version. In the same simplicity-pays vein, we also find that a single-contract OLS-based hedge is beaten by the one-for-one naive rule in 10 cases. All this does not chime well with the positive conclusions about GARCH in some earlier studies (e.g. Ballie and Myers, 1991). Still, it is to be noted that they use a different data set, with much shorter series and sometimes a different GARCH specification.

Next, we address the issue whether rebalancing, which was applied in all the tests discussed thus far, really helps.

4.2.2 Does rebalancing help?

Now we consider the results without rebalancing in the lower panel of Table 2. As one would expect, for most horizons and hedging strategies, a static hedge achieves a lower variance reduction than a dynamic hedge, especially for long hedging horizons. Specifically, out of the 90 cells (commodity $\times$ strategy $\times$ 5 horizons) that we consider, the static hedge comes out first 29 times.

Note also that, even without rebalancing, the naive rules still systematically beat the regression-based strategies: the semi-naive wins in 9 out of 18 times, OLS-2 wins 4 times, GARCH-2 wins 3 times and the traditional naive wins 2 times. This again confirms the simplicity-pays rule even with without rebalancing.

4.2.3 Interim conclusions

To sum up, comparing the variance reduction among strategies, we find that: (i) rebalancing is better than a static hedge; (ii) dual hedges perform better than single ones, even for 'delta' problems without rolling over; (iii) the semi-naive rule systematically beats the other dual-hedge rules for corn and wheat in most cases; and (iv) simplicity pays. The exception is soybeans for which GARCH-2 is the best model. However, as mentioned earlier the differences in variance reduction among dual strategies are not overwhelming.

In the next steps we want to add significance statements for unconditional and conditional differences in variance reduction (subsection 4.3) and transaction-cost adjusted variance (subsection 4.4). As we shall see, up to one exception, the significant differences are all in support of the semi-naive rule rather than the other way around, and the clearest answers are obtained when the criterion is expected utility, i.e. variance adjusted for transaction costs.

4.3 Conditional and unconditional performance: significance tests

As mentioned in the methodology section, due to an insufficient number of observations, we can undertake the GW test for the one-week horizon only. Still, it is reassuring to note that, at this one-week horizon, the GW and DM test results are in full agreement. Tables 4 provides summaries of the results from the DM and GW tests for the out-of-sample hedging performance of alternative strategies with rebalancing for corn, soybeans and wheat. Both tests compare the variances reduction pairwise. In the table, a cell shows the result for the comparison of the model mentioned in the column header and the row header, respectively. The entry in each cell is the name of the strategy that delivers the greater reduction in the conditional variance, as gauged by the DM test. Using the p-values of the GW test for the one-week horizon and DM test for other horizons, we also report via the familiar asterisks whether a strategy does significantly better than the other. Finally, competing strategies are ranked according to their pairwise comparison (6 strategies in total). The best strategy (ranked 1) is the strategy that does better than all the other strategies most often. Next, the second best strategy (ranked 2) is the one that most often does better than all other remaining strategies (after taking out the best-ranking strategy). All strategies are ranked like that. The last and worst strategy is ranked 6.

First, and surely least unexpectedly, all hedging strategies do significantly better than the no-hedging one for all horizons and all commodities.[11] Second, the results confirm the conclusion that we obtain from comparing the variance reduction of alternative strategies. For most of the horizons, the top spots are filled by the dual hedges. Among those, the naive hedges hold the top positions in most cases except for soybeans. In case any competing strategy outperforms the semi-naive rule, it does not significantly better at any horizon, with the single exception of soybeans at the one-week horizon. That is, in all cases but one, where the hedge performance differs significantly from that of the semi-naive rule, the difference is in its favor. Although CCC-GARCH-2 is the best model for soybeans, it does not significantly beat other strategies except at 1-week horizon. In contrast, for corn and wheat GARCH-2 is significantly outperformed by other strategies in 3 cases (at 4 and 6-week horizon for corn and at 10-week horizon for wheat).

At this stage, we still conclude that dual-hedges seem to do better than single-ones, and that, given the choice for either one or two contracts in the hedge, the semi-naive method seems to do better than complicated ones. We now turn to economic relevance: is any utility gained after transaction costs?

4.4 What's left after transaction cost?

Because the dual-hedge rule involves positions in two futures contracts, the transaction cost is likely to be higher than under a one-contract strategy, potentially even wiping out the benefit from the reduced variance. We test for this by evaluating the utility value obtained by the hedger from the dual-hedges, letting him/her decide when to revise, weighing the execution cost against the likely gain in terms of risk. To that end, we suppose that the hedger has a mean-variance utility function. That is, the hedger's criterion after taking into account the transaction cost is:

$$U_{t-1} = [\mathrm{E}_{t-1}(R_{p,t}) - TC_{t-1}] - \frac{\lambda}{2} \cdot \mathrm{var}_{t-1}(R_{p,t}). \tag{12}$$

Following Kroner and Sultan (1993), we suppose that, at time $t-1$, the hedger will only rebalance his/her portfolio when the expected utility gain at time t from rebalancing at $t-1$ is higher than the loss from the transaction cost for doing it at $t-1$. In handling the expected return, we likewise follow Kroner and Sultan (1993) and take the expected portfolio return to be equal to zero. So, the hedger will not rebalance his/her portfolio at time $t-1$ (and thus

[11] We can provide these results upon request.

maintain the same hedge ratios inherited from $t-2$) when the rebalancing does not improve the variance enough to justify the trading cost. Thus, relative risk aversion (λ) allows us to translate risk changes into equivalent terms of expected return after percentage cost and vice versa.

Stated positively, the hedger will rebalance when:

$$-TC_{t-1} - \frac{\lambda}{2} \cdot \mathrm{var}_{t-1}(R_{p,t}|h_{t-1}) \quad > \quad -\frac{\lambda}{2} \cdot \mathrm{var}_{t-1}(R_{p,t}|h_{t-2}) \tag{13}$$

In Appendix III we describe the details on how the rule is applied in some special situations, namely the start of the test and the dates when a contract expires.

The forward-looking conditional variance at time t is calculated from the CCC-GARCH model as estimated at $t-1$. Following Lien and Yang (2007), we set the transaction cost equal to 75 USD per contract[12] per round trip for 5000 bushels. This is converted into a percentage cost per bushel using the day's spot price. We test with four values for relative risk aversion ($\lambda = 2$, 3, and 4). We do not consider for λs of 6 and higher because when $\lambda = 6$ the general stock-market risk premium would implausibly exceed 13.5%.[13] To evaluate the hedging performance of models, we compare the utility, and the best strategy is the one that generates the highest utility. We also perform the GW and DM tests to check whether the differences of the utility gains between alternative strategies are significant. In this case, $\mathrm{d}u_t$, defined as the utility gap $U_t(t^1) - U_t(t^2)$, replaces the variance gap $\mathrm{d}v$ in the DM test. The utility at $t+m$ with rebalancing every k period equals the sum of the utilities at every subperiod:[14]

$$U_{t+m} = \sum_{j=1}^{m/k} U_{t+kj}. \tag{14}$$

Table 3 summarizes the results on cost-adjusted variances for various hedging horizons with weekly rebalancing. The numbers shown under each hedging-horizon header are the percentage increases in the utility which the corresponding strategies provide over the no-hedging alternatives. The results show that the costs do not eclipse the gains from reduction,

[12]This cost includes the brokerage commission fees and bid-ask spreads.

[13]Setting the market volatility as low as 15%, the variance would still be 0.0225 and the risk premium $0.0225 \times 6 = .135$.

[14]One could argue that (12) is not a regular Von Neumann-Morgenstern utility function like, say, expected utility under negative exponential preferences and normality. So, here, time-additive utility for (12) does not have the usual meaning. An alternative interpretation is that (12), with zero expectations, a variance is adjusted for transaction costs in a theoretically justified way, with relative risk aversion as the appropriate weight. Then, the sum is basically the sum of the variances adjusted for trading costs, which is close to the variance of the total return adjusted for costs.

except for the combination of corn or wheat with the lowest risk aversion ($\lambda = 2$). At higher risk aversion levels, the semi-naive rule results in the highest utility increase, except for wheat at the one and 4-week horizons when the traditional naive rule is the winner. In addition, the GW test and DM test indicate that all significant differences are again in favor of the semi-naive rule. The only exception again is soybeans for which the two-contract strategies are significantly beaten by their one-contract counterparts for all values of λ. The results are consistent with our earlier diagnosis that, for soybeans, the convenience yield risk is quite small and closely associated with price risk, such that the benefit from variance reduction via a second contract is wiped out by the transaction cost. Therefore, for soybeans the two-contract hedges are not useful. For soybeans, OLS is the best strategy which provides highest utility gain. However, we find that the traditional naive is the second one and that utility gains from this strategy are nearly as good as that from OLS. This is confirmed by the DM test: only at the 1-week horizon utility gain from OLS strategy is significantly higher than that from the traditional naive rule. In Tables 5, we report the results from GW test and DM test for $\lambda = 4$.[15]

4.5 Further Diagnostics

From the above, our mixture of *a priori* restrictions for the prime source of risk and variance-minimization for the convenience-yield part seems to work. We did some experimenting as to which of these ingredients is most crucial. First we ran OLS regressions with constraint (5) imposed but leaving the regression otherwise unchanged. The result was that the hedge consisted almost exclusively of the nearest contract, and its performance was very comparable to the standard one-for-one hedge. We then tried two rules that, like ours, lead to above-unity positions in the nearest contract and a short position in the longer-dated one. One was, hyper-naively, just based on time to maturity: we created a position with a weighted life equal to the horizon m. This would have worked if convenience were a constant and exact linear function of time to maturity, but this is clearly not appropriate: this hedge has a middling performance only. The last variant we tried aimed at elimination of the expected change of the convenience yield, the complement to the residual-risk minimization we are applying now.[16] This worked almost as well as the current rule, but the hedge ratios were extremely noisy (including 3300,

[15] We can provide the results for other values of λ upon request.

[16] The total uncertainty about the yield-related part of the hedge consists of the squared expected yield change and the residual risk. The semi-naive rule in this paper considers the latter component, the alternative focused on the former.

in an extreme case) and required two rounds of smoothing before they became acceptable. Even then the performance remained marginally below the semi-naive one considered here.

The conclusion, then, is that a substantial part of the improvement stems from the better modeling of the yield, using the seasonals, market tightness, time to maturity, and time beyond the harvest. Regular regressions regard any deviation from the grand mean as unpredictable, which seems to be a non-trivial shortcoming. This, lastly, raises the question whether the two-step approach cannot be merged into one. Currently we are first modeling yields, and computing moments for unexpected yield changes which then serve as inputs into the yield-risk minimization rule. The alternative, left for future research, would be to run a constrained Ederington regression which also contains, as a 'control', the first-difference version of the yield equation.

5 Conclusion

This paper studies the use of two futures contracts in hedging an agricultural-commodity commitment or inventory position. We propose to use dual hedges not just for roll-over problems but also for delta hedges. Next to regression-based hedges using two contracts, we also propose a semi-naive rule, half-way between a regression hedge and the standard rule of thumb. Like regression it does minimize variance; it even employs a regression model for the expected convenience yields at time $t+m$. However, it has the potential advantage that it confines statistics to a well-specified area, *viz.* managing convience-yield risk without allowing any spill-overs to the spot-risk part of the hedge problem. The spot exposure is set a priori at a level (close to) unity, like in a naive hedge. This feature has the advantage that it confines estimation errors to the secondary source of risk, the convenience yield. The hedge also uses conditioning information about that second source of risk, again unlike the standard regression hedge. Given its generality, this semi-naive strategy is easily applied to agricultural commodities for which some recently proposed strategies are not immediately suitable. To forecast convenience yields, we adopt the simplest version of the Carbonez, Nguyen and Sercu (2008) model, with price as the sole measure for current scarcity. In order to compare the semi-naive rule with its competitors, we evaluate both the variance reduction of the portfolio value and the utility gain after considering execution cost (i.e. the transaction-cost adjusted variance). We also add significance test for unconditional and conditional equality of the remaining risks (Diebold and Mariano, 2002; Giacomini and White, 2006). Lastly, we implement the hedging experiments both in a static and a dynamic fashion, i.e. without and with interim rebalancing.

Our results confirm that for hedging horizons exceeding one period (one week, in this paper), rebalancing does help to reduce the portfolio variance. Therefore, in our analysis we focus on this version. We find that two-contract hedges do better than single-contract ones, even in a pure delta-hedge situation where the horizon is shorter than the lives of the two futures used as hedge instruments. More interestingly, the results indicate that simplicity pays: in terms of variance reduction, for corn and wheat, we find that the naive rules (traditional and semi) do better than the regression based hedges, OLS and CCC-GARCH (in that order). The naive rules are, in fact, the winners 7 times out of 12, and in the remaining 5 cases they are not systematically and significantly beaten by one and the same contender. The only exception is for soybeans for which CCC-GARCH-2 does best—but even there it significantly ahead of naive rules only at the one-week horizon. In contrast, for corn and wheat GARCH-2 is significantly outperformed by other rules in some cases while the semi-naive rule never is.

In terms of utility gains after transaction costs, lastly, for corn and wheat the naive rules win too. However, for soybeans, the utility gains from the one-contract strategies are significantly higher than from the two-contract strategies, a feature that we can trace to the rather special covariance matrix. Even for soybeans utility gains are highest from OLS strategy, but the differences between OLS and traditional naive rule are not significant except for the one-week horizon.

Overall, then, for agricultural commodities, two-contract hedges do better than one-contract counterpart and the simple rules are the best one.

References

Andersen, T.G. and T. Bollerslev (1998). "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economics Review.* 39: 885-905.

Baillie, R., R. Mayers (1991). "Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge." *Journal of Applied Economics.* 6: 109-124. Bera, A.K, Garcia, P., Roh, J.S. (1997)."Estimation of Time-Varying Hedge Ratios for Corn and Soybeans: BGARCH and Random Coefficient Approaches". OFOR paper Number 97-06.

Bollerslev, T. (1990). "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model." *Review of Economics and Statistics.* 72: 498-505.

Brennan, M. (1991). "The Price of Convenience and The Valuation of Commodity Contingent Claims." *D.Lund and B. Oksendal (Editors), Stochastic models and option values.* 33-71. Amsterdam: North-Holland.

Brennan, M., and N. Crew (1997). "Hedging Long Maturity Commodity Commitments with Short-Dated Futures Contracts." *M. Dempster and S. Pliska (Editors), Mathematics of derivative securities.* 165-190. Cambridge: Cambridge University Press.

Carbonez, K.A.E., T.T.V. Nguyen and P. Sercu (2008). "Remodeling the Working Curve: The Roles of Scarcity, Time to Maturity and Time to Harvest". Working paper.

Lien, D. and Li Yang (2008). "Asymmetric Effect of Basis on Dynamic Futures Hedging: Empirical Evidence from Commodity Markets." *The Journal of Banking and Finance.* 32:187-198.

Diebold, F. X. and R.S. Mariano (1995). "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics.* 13:253-263.

Diebold, F. X. and J.A. Lopez (1996). "Forecast Evaluation and Combination." in *Handbook of Statistics.* 14: Statistical Methods in Finance, ed. by G.S. Maddala and C.R. Rao. Amsterdam: North-Holland. 241-268.

Engle, R.F. (2002). "Dynamic Conditional Correlation: A New Simple Class of Multivariate GARCH Models." *The Journal of Business Economic Statistics.* 20:339-350.

Ederington L. H. (1979). "The Hedging Performance of the New Futures Markets." *The Journal of Finance.* 20(1):157-170.

Fama, E.F., and K.R. French (1987). "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage." *The Journal of Business.* 60(1):55-73.

Gelper, S., R. Fried and C. Croux (2007). "Robust Forecasting with Exponential and Holt-Winters Smoothing."

Giacomini, R., and H. White (2006). "Test of Conditional Predictive Ability." *Econometrica.* 74(6):1545-1578.

Gibson, R., and E. Schwartz (1990). "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims." *The Journal of Finance.* 45(3):959-976.

Kroner, K.F., J. Sultan (1993). "Time Varying Distribution and Dynamic Hedging with Foreign Currency Futures." *The Journal of Financial Quantitative Analysis.* 28: 535-551.

Schwartz, E.S. (1997). "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging." *The Journal of Finance.* 52:923-973.

Stein Jerome L. (1961). "The Simultaneous Determination of Spot and Futures Prices." *The American Economic Review.* 51(5):1012-1025.

Johnson Leland L. (1960). "The Theory of Hedging and Speculation in Commodity Futures." *The Review of Economic Studies.* 27(3):139-151.

Neuberger, A. (1999). "Hedging Long-Term Exposures with Multiple Short-Term Futures Contracts." *The Review of Financial Studies.* 12(3):429-459.

Ng, V.K. and S.C. Pirrong (1994). "Fundamental and Volatility: Storage, Spreads and the Dynamics of Metal Prices." *Journal of Business.* 67(2):203-230.

Pesaran, M.H., and Smith, R. (1995). "Estimating Long-run Relationship from Dynamic Heterogeneous Panels." *Journal of Econometrics.* 68, 79-113.

Sercu, P and R. Uppal (1995). *International Financial Markets and The Firm.* Southwestern Publishing/Thomson International, Cincinnatti, Ohio.

Sercu, P. and Wu, X. (2000). Cross- and delta-hedges: Regression- versus price-based hedge ratios. *Journal of Banking and Finance*, 24(5): 735 - 757.

Sercu, P. (2009). *International Finance: Theory into Practice.* Princeton University Press, Princeton, New Jersey.

Working, H. (1949). "The Theory of the Price of Storage." *American Economic Review.* 39:1254-1262.

Veld-Merkoulova, Y.V. and Frans A. De Roon (2003). "Hedging Long-term Commodity Risk." *The Journal of Futures Markets.* 23(2):109-133.

Zellner, A. (1969)."On the Aggregation Problem, a New Approach to a Troublesome Problem, Estimation and Risk Programming: Essays in honor of Gerhard Titner, Berlin, Springer."

Appendix I: Derivation of the semi-naive hedge ratios:

Starting from Equation (5), we can write:

$$\begin{aligned} h_{2,t} &= \frac{1 - h_{1,t} \cdot (1 + r_t \cdot (T_1 - t - m))}{(1 + r_t \cdot (T_2 - t - m))} \\ &= I_t - J_t \cdot h_{1,t}. \end{aligned} \quad (15)$$

with $I_t = \frac{1}{1+r_t \cdot (T_2 - t - m)}$ and $J_t = \frac{1+r_t \cdot (T_1 - t - m)}{1+r_t \cdot (T_2 - t - m)}$.

Substitute Equation (15) to Equation (6), it becomes:

$$h_{1,t}^2 \cdot \text{var}_{t,t+m}(e_{T_1}) + (I_t - J_t \cdot h_{1,t})^2 \cdot \text{var}_{t,t+m}(e_{T_2}) + 2h_{1,t} \cdot (I_t - J_t \cdot h_{1,t}) \cdot \text{cov}_{t,t+m}(e_{T_1}, e_{T_2}). \quad (16)$$

Set the first derivative of Equation (16) to zero to solve the minimization problem, we get:

$$\begin{aligned} 0 &= \frac{\partial}{\partial h_1} \\ &= 2h_{1,t} \cdot \text{var}_{t,t+m}(e_{T_1}) - 2I_t J_t \cdot \text{var}_{t,t+m}(e_{T_2}) + 2h_{1,t} J_t^2 \cdot \text{var}_{t,t+m}(e_{T_2}) \\ &\quad +2I_t \cdot \text{cov}_{t,t+m}(e_{T_1}, e_{T_2}) - 4h_{1,t} J_t \cdot \text{cov}_{t,t+m}(e_{T_1}, e_{T_2}) \\ &= h_{1,t}[\text{var}_{t,t+m}(e_{T_1}) + J_t^2 \cdot \text{var}_{t,t+m}(e_{T_2}) - 2J_t \cdot \text{cov}_{t,t+m}(e_{T_1}, e_{T_2})] \\ &\quad -I_t[J_t \cdot \text{var}_{t,t+m}(e_{T_2}) - \text{cov}_{t,t+m}(e_{T_1}, e_{T_2})]. \end{aligned} \quad (17)$$

Set $\text{var}_{t,t+m}(e_{T_i}) = \text{var}_t(e_{T_i})$ and $\text{cov}_{t,t+m}(e_{T_1}, e_{T_2}) = \text{cov}_t(e_{T_1}, e_{T_2})$, we get:

$$\begin{cases} h_{1,t} = \frac{I_t[J_t \cdot \text{var}_t(e_{T_2}) - \text{cov}_t(e_{T_1}, e_{T_2})]}{\text{var}_t(e_{T_1}) - 2J_t \cdot \text{cov}_t(e_{T_1}, e_{T_2}) + J_t^2 \cdot \text{var}_t(e_{T_2})}, \\ \\ h_{2,t} = I_t - J_t \cdot h_{1,t}. \end{cases}$$

Appendix II: The DM and GW tests

The DM method is an unconditional test of the null hypothesis of equal conditional hedging effectiveness.

$$H_{0,dm} : E_t[\text{d}v_t] = 0. \quad (18)$$

We define T_{out} := the number of observations used out-of-sample; $\overline{\text{d}v} := T_{out}^{-1} \sum_{t=T_{in}+1}^{T} \text{d}v_t$; $\widehat{LRV}(\text{d}v_t)$:= the estimate of long run variance of $\text{d}v_t$; and:

$$dm := \frac{\sqrt{T_{out}} \cdot \overline{\text{d}v}}{\sqrt{\widehat{LRV}(\text{d}v_t)}}, \quad (19)$$

which is a standard t-test on a mean (of dv, here). Diebold and Mariano (1995) show that $dm \stackrel{\mathrm{d}v}{\to} N(0,1)$. In an application, the DM statistic can be computed as the t-statistic in a regression of dv_t on a constant with the Newey-West standard error. Note that the sign of $\overline{\mathrm{d}v}$ indicates the direction of rejection of the two-sided DM test. If in our test we observe, for instance, $\overline{\mathrm{d}v} < 0$, this implies that $cv(t^1) < cv(t^2)$, i.e. that strategy t^1 does better than t^2 and vice versa.

Giacomini and White (2006) propose a conditional test. They construct a test of $H_{0,z}$ against the two-sided alternative:

$$H_{1,z} : E[Z_t]^{'} E[Z_t] > 0. \tag{20}$$

The GW test statistic can be written as:

$$gw_z = T_{out} R^2, \tag{21}$$

where R^2 is the uncentered square multiple correlation coefficient for the artificial regression $\iota = Z\theta + \epsilon$. Define $Z \equiv (\mathrm{d}v_{T_{in}+1} z^{'}_{T_{in}} ... \mathrm{d}v_T z^{'}_{T-1})^{'}$, a $T_{out} \times q$ matrix and ι, a $T_{out} \times 1$ vector of ones. Under $H_{0,z}$, $gw_z \stackrel{\mathrm{d}v}{\to} \chi^2_{(q)}$. We choose the test function z_{t-1} as follows:

$$z_{t-1} \equiv (1, S_{t-1}, Y_{t-1,T_1}, \mathrm{d}v_{t-1}, f_{t-1,T_1}, S_{t-1}\, f_{t-1,T_1})^{'} \tag{22}$$

Appendix III: Details of tests with endogenous hedging

There are some special situations that deserve a detailed discussion. First, at the beginning of a new hedging period, we do not rebalance the portfolio since the hedge is just starting. However, we still apply the condition in Equation (13), with inherited hedge ratios from time $t-2$ equal to zero, to decide whether the hedger should hold any hedged position.[17] If the above condition is not met, then for the one-contract strategies the agent does not hedge, while for the two-contract strategies, the hedger switches from a two- to a one-contract strategy. For OLS-2 and CCC-GARCH-2, the hedger obviously uses the hedge ratio from OLS and CCC-GARCH strategies when switching. For our the semi-naive rule, the one-contract hedge ratio is calculated from Equation (5): $1 - \sum_{i=1}^{2} h_{i,t} \cdot [1 + r_t \cdot (T_i - t - m)] = 0$ with $h_{2,t} = 0$. In case the new calculated one-contract hedge ratio still does not help enough, in terms of expected risk

[17] Thus, for the one-week horizon, each week the hedger decides whether to hold a hedged position or not, not whether to rebalance the position.

reduction, in order to justify the expense, the hedger holds the spot position, unhedged. This way of switching from two- to one-contract strategies is also applied during the hedging period when no position is held in the long futures contract (i.e when $h_{2,t-1} = 0$). More specifically, such situation occurs: (i) when the short-dated contract has expired and the transaction cost condition prevents the hedger from holding any position with the new long contract or (ii) when the hedger decides not to rebalance the hedged portfolio and the previous hedge ratio for the long contract is zero.

Second, consider the situation where the hedger does not rebalance the portfolio but the previous hedge ratio for the long contract is different from zero. In this case, for our semi-naive rule, the hedger re-calculates the hedge ratio for a short futures contract at $t - 1$ with $h_{2,t-1} = h_{2,t-2}$ based on Equation (5).

Table 2: Out-of-sample percentage of annualized variance reduction

Hedging Strategy	Panel A: Corn						Panel B: Soybeans						Panel C: Wheat					
	Hedging Horizon (in weeks)						*Hedging Horizon (in weeks)*						*Hedging Horizon (in weeks)*					
	1	4	6	8	10	12	1	4	6	8	10	12	1	4	6	8	10	12
	Weekly rebalance																	
Naive hedge	**67.53**	69.44	**73.64**	77.97	81.13	80.02	68.78	68.28	64.59	65.29	68.60	66.38	71.31	79.23	78.10	80.13	77.62	82.07
Semi-naive	67.46	**69.82**	73.60	78.54	**81.63**	80.39	70.57	71.08	66.57	67.34	70.43	67.54	**71.52**	**79.44**	**79.10**	81.97	79.19	83.89
OLS	67.43	69.19	73.09	78.83	80.82	80.19	68.47	67.62	64.12	64.85	68.30	66.13	70.91	79.41	77.85	81.61	77.92	82.55
OLS-2	67.53	69.75	73.47	78.76	81.57	**80.49**	70.57	71.14	66.45	67.37	70.49	67.49	71.04	79.08	78.64	**83.32**	**80.20**	84.63
GARCH-CCC	66.67	65.65	70.91	**79.43**	79.06	80.19	68.09	70.33	70.23	69.02	70.97	70.62	70.09	78.07	77.92	81.48	76.90	84.22
GARCH-CCC-2	65.66	65.23	69.31	76.35	76.01	76.54	**70.88**	**74.95**	**73.43**	**73.09**	**74.35**	**73.14**	70.67	77.92	77.49	81.58	75.51	**84.76**
	No rebalance																	
Naive hedge	**67.53**	69.64	73.31	75.01	**79.78**	78.09	68.78	69.15	64.17	71.58	72.75	69.53	71.31	77.45	79.41	74.74	79.46	67.77
Semi-naive	67.46	**70.34**	**74.55**	75.43	79.25	**78.74**	70.57	71.02	**64.31**	**71.82**	**72.84**	70.08	**71.52**	**78.62**	**80.58**	76.91	81.28	70.17
OLS	67.43	69.06	73.31	73.72	77.16	75.67	68.47	68.50	63.94	70.88	72.19	68.62	70.91	77.99	79.57	73.82	79.24	66.86
OLS-2	67.53	69.90	73.38	74.42	76.87	75.39	70.57	**71.73**	63.77	70.95	71.83	69.02	71.04	78.53	79.82	**78.09**	**82.67**	**74.50**
GARCH-CCC	66.67	65.54	46.24	41.20	52.48	43.92	68.09	71.63	51.59		66.70		70.09	57.60	66.46	70.88	65.09	
GARCH-CCC-2	65.66	69.28	52.90	**76.95**			**70.88**	64.63	59.37	66.27	67.12	**73.24**	70.67	64.77	58.50	67.56	70.61	

Notes:
1. The table reports the annualized percentage variance reduction of alternative hedging strategies over the no hedging one. The upper panel is for weekly rebalance while the lower panel is for no rebalance method. To annualize, we divide the standard deviation by the square root of the length of the hedging period (in years). We use 52 weeks per year to annualize portfolio return variance.
2. In each panel (upper and lower), a number in bold in a column is the highest variance reduction within this column.
3. For corn, soybeans and wheat, in some horizons, there is a problem for CCC-GARCH estimation (negative of estimated variance) for without rebalancing. So in the report, these numbers are missing.

Table 3: Percentage increase of utility over the no-hedging strategy with the effect of conditional transaction cost

Hedging Strategy	Panel A: Corn						Panel B: Soybeans						Panel C: Wheat					
	Hedging Horizon (in weeks)						*Hedging Horizon (in weeks)*						*Hedging Horizon (in weeks)*					
	1	4	6	8	10	12	1	4	6	8	10	12	1	4	6	8	10	12
	$\lambda = 2$																	
Naive hedge	0.00	0.00	0.00	0.00	0.00	0.00	**9.68**	**22.87**	**27.61**	**29.16**	**30.47**	**31.11**	0.00	0.00	0.00	0.00	0.00	0.00
Semi-naive	0.00	0.00	0.00	0.00	0.00	0.00	6.07	23.19	25.99	28.67	29.69	29.48	0.00	0.00	0.00	0.00	0.00	0.00
OLS	0.00	0.00	0.00	0.00	0.00	0.00	**10.01**	**24.19**	**27.66**	**29.18**	**30.47**	**31.11**	0.00	0.00	0.00	0.00	0.00	0.00
OLS-2	0.00	0.00	0.00	0.00	0.00	0.00	7.62	23.34	26.32	28.33	29.13	29.77	0.00	0.00	0.00	0.00	0.00	0.00
GARCH-CCC	0.00	0.00	0.00	0.00	0.00	0.00	8.66	21.72	24.25	24.87	26.28	26.99	0.00	0.00	0.00	0.00	0.00	0.00
GARCH-CCC-2	0.00	0.00	0.00	0.00	0.00	0.00	8.66	21.72	24.25	24.87	26.28	26.99	0.00	0.00	0.00	0.00	0.00	0.00
	$\lambda = 3$																	
Naive hedge	0.00	4.34	4.34	5.72	5.72	4.34	**19.66**	**38.18**	**40.51**	**41.61**	**43.55**	**42.87**	0.00	0.19	0.67	0.67	0.67	0.67
Semi-naive	0.00	**4.39**	**4.39**	**5.77**	**5.77**	**4.39**	13.09	34.78	38.10	38.27	42.15	41.02	0.00	**1.33**	**0.67**	**2.48**	**2.48**	**0.67**
OLS	0.00	0.00	0.00	0.00	0.00	0.00	**20.14**	**38.31**	**40.59**	**41.67**	**43.59**	**42.90**	0.00	0.16	0.65	0.65	0.65	0.65
OLS-2	0.00	0.00	0.00	0.00	0.00	0.00	14.78	36.52	38.22	41.15	42.12	40.59	0.00	0.16	0.65	0.65	0.65	0.65
GARCH-CCC	0.00	1.73	1.73	1.73	1.73	1.73	19.11	36.38	40.04	39.99	41.99	42.45	0.00	0.00	0.00	0.00	0.00	0.00
GARCH-CCC-2	0.00	1.73	1.73	1.73	1.73	1.73	12.17	33.52	36.13	38.45	39.21	38.85	0.00	0.00	0.00	0.00	0.00	0.00
	$\lambda = 4$																	
Naive hedge	1.53	6.18	7.24	6.85	7.92	10.08	**27.32**	**46.37**	**47.65**	**51.79**	**53.25**	**52.17**	**0.48**	**1.68**	2.10	2.77	2.77	2.10
Semi-naive	**1.66**	**6.24**	**7.30**	**6.91**	**7.96**	**10.14**	17.14	41.63	44.55	48.68	49.64	49.31	0.16	1.46	**2.49**	**3.17**	**3.17**	**2.49**
OLS	1.16	5.38	5.38	6.08	6.08	5.38	**27.80**	**46.93**	**48.18**	**52.90**	**53.73**	**52.65**	0.41	1.61	2.06	2.74	2.74	2.06
OLS-2	1.16	5.38	5.38	6.08	6.08	5.38	19.07	42.62	45.13	48.70	51.96	50.25	0.11	1.30	2.28	2.95	2.95	2.28
GARCH-CCC	1.46	5.61	5.61	6.32	6.32	5.61	27.11	46.13	47.48	51.90	52.82	52.00	0.16	1.25	1.67	2.26	2.26	1.67
GARCH-CCC-2	1.46	5.61	5.61	6.32	6.32	5.61	16.17	40.40	43.23	47.08	49.82	47.64	0.16	1.25	1.67	2.26	2.26	1.67

Notes:
1. The table reports the annualized percentage utility increase over the no-hedging strategy for alternative hedging strategies for different hedging horizons with weekly rebalancing. To annualize the utility, we divide the utility by the length of the hedging period (in years). We use 52 weeks per year to annualize.

Table 4: Pairwise comparisons of variance reductions from GW and DM test - Weekly rebalancing

	Panel A: Corn						Panel A: Soybeans						Panel A: Wheat					
	semi	OLS	OLS-2	GARCH	GARCH-2	Rank	semi	OLS	OLS-2	GARCH	GARCH-2	Rank	semi	OLS	OLS-2	GARCH	GARCH-2	Rank
1 week hedging horizon																		
Naive hedge	naive	naive	naive	naive**	naive	1	naive2	naive	ols2	naive**	ccc2	4	naive2	naive	naive	naive	naive	2
Semi-naive		naive2	ols2	naive2	naive2	3		naive2	naive2	naive2*	ccc2*	2		naive2	naive2	naive2*	naive2	1
OLS			ols2	ols*	ols	4			ols2	ols**	ccc2	5			ols2	ols	ols	4
OLS-2				ols2*	ols2	2				ols2	ccc2*	3				ols2	ols2	3
GARCH-CCC					ccc	5					ccc2	6					ccc2	6
GARCH-CCC-2						6						1						5
4 week hedging horizon																		
Naive hedge	naive2	naive	ols2	naive**	naive**	3	naive2	naive	ols2*	ccc	ccc2	5	naive2	ols	naive	naive	naive	3
Semi-naive		naive2	naive2	naive2**	naive2**	1		naive2	ols2	naive2	ccc2	3		naive2	naive2	naive2	naive2	1
OLS			ols2	ols**	ols**	4			ols2	ccc	ccc2	6			ols	ols	ols	2
OLS-2				ols2**	ols2**	2				ols2	ccc2	2				ols2	ols2	4
GARCH-CCC					ccc	5					ccc2**	4					ccc	5
GARCH-CCC-2						6						1						6
6 week hedging horizon																		
Naive hedge	naive	naive	naive	naive*	naive*	1	naive2	naive	ols2	ccc	ccc2	5	naive2*	naive	ols2	naive	naive	3
Semi-naive		naive2	naive2	naive2	naive2*	2		naive2	naive2	ccc	ccc2	3		naive2	naive2	naive2	naive2	1
OLS			ols2	ols	ols	4			ols2	ccc	ccc2	6			ols2	ccc	ols	5
OLS-2				ols2	ols2*	3				ccc	ccc2	4				ols2	ols2	2
GARCH-CCC					ccc	5					ccc2	2					ccc	4
GARCH-CCC-2						6						1						6
8 week hedging horizon																		
Naive hedge	naive2	ols	ols2	ccc	naive	5	naive2	naive	ols2	ccc	ccc2	5	naive2	ols	ols2	ccc	ccc2	6
Semi-naive		ols	ols2	ccc	naive2	4		naive2	ols2	ccc	ccc2	4		naive2	ols2	naive2	naive2	2
OLS			ols	ccc	ols	2			ols2	ccc	ccc2	6			ols2*	ols	ols	3
OLS-2				ccc	ols2	3				ccc	ccc2	3				ols2*	ols2	1
GARCH-CCC					ccc	1					ccc2	2					ccc2	5
GARCH-CCC-2						6						1						4
10 week hedging horizon																		
Naive hedge	naive2	naive	ols2	naive	naive	3	naive2	naive	ols2	ccc	ccc2	5	naive2	ols	ols2	naive	naive	4
Semi-naive		naive2	naive2	naive2	naive2	1		naive2	ols2	ccc	ccc2	4		naive2	ols2	naive2	naive2	2
OLS			ols2	ols	ols	4			ols2	ccc	ccc2	6			ols2	ols	ols	3
OLS-2				ols2*	ols2	2				ccc	ccc2	3				ols2**	ols2*	1
GARCH-CCC					ccc	5					ccc2	2					ccc	5
GARCH-CCC-2						6						1						6
12 week hedging horizon																		
Naive hedge	naive2	ols	ols2	ccc	naive	5	naive2	naive	ols2	ccc	ccc2	5	naive2	ols	ols2	ccc	ccc2	6
Semi-naive		naive2	ols2	naive2	naive2	2		naive2	naive2	ccc	ccc2	3		naive2	ols2	ccc	ccc2	3
OLS			ols2	ccc	ols	4			ols2	ccc	ccc2	6			ols2	ccc	ccc2	5
OLS-2				ols2	ols2	1				ccc	ccc2	4				ols2	ccc2	2
GARCH-CCC					ccc	3					ccc2	2					ccc2	4
GARCH-CCC-2						6						1						1

Notes:
1. The table provides summaries of the results from the DM test to compare the pairwise variance reductions with rebalancing. In the table, a cell shows the result for the comparison of the model mentioned in the column header and the row header, respectively. The entry in each cell is the name of the strategy that delivers the greater reduction in the conditional variance, as gauged by the DM test. This result is based on the sign of the coefficient for constant term on the regression of the difference of the two conditional variance dv_t on constant with OLS (Newey-West) estimation. *naive*, *naive2*, *ols*, *ols2*, *ccc* and *ccc2* means traditional naive hedge, semi-naive, one-contract OLS, two-contract OLS, one-contract CCC- GARCH and two-contract CCC- GARCH is the better one respectively.
2. Using the p-values of the GW test for one-week horizon and DM test for other horizons, we also report via the familiar asterisks whether a strategy does significantly better than the other. ***, ** and *indicate the significant of the better model compared to the other at 1%. 5% and 10% consecutively.
3. Competing strategies are ranked according to their pairwise comparison (6 strategies in total). The best strategy (ranked 1) is the strategy that does better than all other strategies. Next, the second best strategy (ranked 2) is the one that does better than all other remaining strategies (after taking out the ranked 1 strategy). All strategies are ranked like that. The last and worst strategy is ranked 6.

Table 5: Pairwise comparisons of utility gains from GW and DM test - Weekly rebalancing - $\lambda = 4$

	Panel A: Corn						Panel A: Soybeans						Panel A: Wheat					
	semi	OLS	OLS-2	GARCH	GARCH-2	Rank	semi	OLS	OLS-2	GARCH	GARCH-2	Rank	semi	OLS	OLS-2	GARCH	GARCH-2	Rank
1 week hedging horizon																		
Naive hedge	naive2*	naive*	naive*	naive	naive	2	naive***	ols***	naive**	naive	naive***	2	naive	naive	naive	naive	naive	1
Semi-naive		naive2*	naive2*	naive2	naive2	1		ols***	ols2	ccc***	naive2***	5		ols	naive2	ccc	ccc2	5
OLS			ols2*	ccc	ccc2	6			ols**	ols	ols***	1			ols***	ols	ols	2
OLS-2				ccc	ccc2	5				ccc**	ols2***	4				ccc	ccc2	6
GARCH-CCC					ccc2	4					ccc**	3					ccc2	4
GARCH-CCC-2						3						6						3
4 week hedging horizon																		
Naive hedge	naive2	naive	naive	naive	naive	2	naive***	ols	naive*	naive	naive***	2	naive	naive	naive	naive	naive	1
Semi-naive		naive2	naive2	naive2	naive2	1		ols***	ols2	ccc**	naive2	5		ols	naive2	naive2	naive2	3
OLS			ols2	ccc	ccc2	6			ols**	ols	ols**	1			ols***	ols	ols	2
OLS-2				ccc	ccc2	5				ccc	ols2	4				ols2	ols2	4
GARCH-CCC					ccc2	4					ccc**	3					ccc2	6
GARCH-CCC-2						3						6						5
6 week hedging horizon																		
Naive hedge	naive2	naive	naive	naive	naive	2	naive**	ols	naive	naive	naive*	2	naive2	naive	ols2	naive	naive	3
Semi-naive		naive2	naive2	naive2	naive2	1		ols**	ols2	ccc	naive2	5		naive2	naive2	naive2	naive2	1
OLS			ols2	ccc	ccc2	6			ols*	ols	ols**	1			ols2***	ols	ols	4
OLS-2				ccc	ccc2	5				ccc	ols2	4				ols2	ols2	2
GARCH-CCC					ccc2	4					ccc	3					ccc2	6
GARCH-CCC-2						3						6						5
8 week hedging horizon																		
Naive hedge	naive2	naive	naive	naive	naive	2	naive**	ols	naive	ccc	naive	3	naive2	naive	ols2	naive	naive	3
Semi-naive		naive2	naive2	naive2	naive2	1		ols**	ols2	ccc	naive2	5		naive2	naive2	naive2	naive2	1
OLS			ols2	ccc	ccc2	6			ols*	ols	ols**	1			ols2***	ols	ols	4
OLS-2				ccc	ccc2	5				ccc	ols2	4				ols2	ols2	2
GARCH-CCC					ccc2	4					ccc*	2					ccc2	6
GARCH-CCC-2						3						6						5
10 week hedging horizon																		
Naive hedge	naive2	naive	naive	naive	naive	2	naive**	ols	naive	naive	naive	2	naive2	naive	ols2	naive	naive	3
Semi-naive		naive2	naive2	naive2	naive2	1		ols**	ols2	ccc	ccc2	6		naive2	naive2	naive2	naive2	1
OLS			ols2	ccc	ccc2	6			ols	ols	ols*	1			ols2***	ols	ols	4
OLS-2				ccc	ccc2	5				ccc	ols2	4				ols2	ols2	2
GARCH-CCC					ccc2	4					ccc	3					ccc2	6
GARCH-CCC-2						3						5						5
12 week hedging horizon																		
Naive hedge	naive2	naive	naive	naive	naive	2	naive*	ols	naive	naive	naive*	2	naive2	naive	ols2	naive	naive	3
Semi-naive		naive2	naive2	naive2	naive2	1		ols**	ols2	ccc	naive2	5		naive2	naive2	naive2	naive2	1
OLS			ols2	ccc	ccc2	6			ols	ols	ols**	1			ols2***	ols	ols	4
OLS-2				ccc	ccc2	5				ccc	ols2	4				ols2	ols2	2
GARCH-CCC					ccc2	4					ccc	3					ccc2	6
GARCH-CCC-2						3						6						5

Notes:
1. The table provides summaries of the results from the DM test to compare the utility pairwise after taking into account the conditional transaction cost with weekly rebalancing. In the table, a cell shows the result for the comparison of the model mentioned in the column header and the row header, respectively. The entry in each cell is the name of the strategy that delivers the greater increase in the utility, as gauged by the DM test. This result is based on the sign of the coefficient for constant term on the regression of the difference of the two utility du_t on constant OLS (Newey-West) estimation. *naive*, *naive2*, *ols*, *ols2*, *ccc* and *ccc2* means traditional naive hedge, semi-naive, one-contract OLS, two-contract OLS, one-contract CCC- GARCH, and two-contract CCC- GARCH strategy is the better one respectively.
2. Using the p-values of the GW test for one-week horizon and DM test for other horizons, we also report via the familiar asterisks whether a strategy does significantly better than the other. ***, ** and *indicate the significant of the better model compared to the other at 1%. 5% and 10% consecutively.
3. Competing strategies are ranked according to their pairwise comparison (6 strategies in total). The best strategy (ranked 1) is the strategy that does better than all other strategies. Next, the second best strategy (ranked 2) is the one that does better than all other remaining strategies (after taking out the ranked 1 strategy). All strategies are ranked like that. The last and worst strategy is ranked 6.

Table 6: Initial weekly estimation result for convenience yields

Coefficient	Corn	Soybeans	Wheat
$Alpha$	2.583	−0.676	−4.376
$Alpha - jan$		2.511	
$Alpha - mar$	−2.050	−3.050	−8.265
$Alpha - may$	−3.012	−5.331	−1.436
$Alpha - jul$	−1.240	−6.528	10.593
$Alpha - aug$		−3.631	
$Alpha - sep$	2.860	5.675	4.006
$Alpha - nov/dec$	3.417	10.170	−6.233
$(T-t) * spot$	0.002	$7.0e-4$	0.001
$(T-T_h) * spot$	0.001	$8.4e-4$	0.001
$(T-t) * Jan$	−0.385	−0.228	−0.011
$(T-t) * Feb$	−0.391	−0.238	−0.107
$(T-t) * Mar$	−0.394	−0.214	−0.233
$(T-t) * Apr$	−0.383	−0.202	−0.234
$(T-t) * May$	−0.317	0.131	−0.272
$(T-t) * Jun$	−0.320	0.215	−0.264
$(T-t) * Jul$	−0.345	0.107	−0.242
$(T-t) * Aug$	−0.345	−0.220	−0.256
$(T-t) * Sep$	−0.366	−0.259	−0.194
$(T-t) * Oct$	−0.362	−0.261	−0.188
$(T-t) * Nov$	−0.358	−0.221	−0.174
$(T-t) * Dec$	−0.369	−0.234	0.030
$(T-T_h) * Jan$	−0.173	−0.487	−0.714
$(T-T_h) * Feb$	−0.173	−0.490	−0.623
$(T-T_h) * Mar$	−0.180	−0.509	
$(T-T_h) * Apr$	−0.187	−0.516	
$(T-T_h) * May$		−0.858	
$(T-T_h) * Jun$		−0.933	
$(T-T_h) * Jul$		−0.782	−0.470
$(T-T_h) * Aug$			−0.441
$(T-T_h) * Sep$	−0.145	−0.196	−0.495
$(T-T_h) * Oct$	−0.163	−0.343	−0.520
$(T-T_h) * Nov$	−0.177	−0.466	−0.519
$(T-T_h) * Dec$	−0.169	−0.485	−0.725

Notes: 1. The table summarizes results of the initial estimation for convenience yields with weekly frequency from price model proposed by Carbonez, Nguyen and Sercu (2008): with scarcity measured by the price ('Price' model).

$$(\text{'Price'}) \quad : \quad Y_{t,T} = \alpha + (\sum_{t=1}^{12} \delta_m \mathbf{1}_{M(t)=m} + \beta S_{t-1})(T-t) + (\sum_{t=1}^{12} \rho_m \mathbf{1}_{M(t)=m} + \zeta S_{t-1}) max(T-T_h, 0), \tag{23}$$

where $\mathbf{1}_{x>k} = 1$ if 1/norminv $> k$, otherwise $\mathbf{1}_{x>k} = 0$.
2. For soybeans, $T - T_h = 150$ for January; $T - T_h = 210$ for March; 270 for May; 330 for July; 360 for August; 30 for September and 90 for November contract. For wheat, $T - T_h = 315$ for March; 375 for May; 75 for July; 135 for September and 225 for December contract.