



Do Inventories Really Yield a Convenience? An Empirical Analysis of the Cost-adjusted Basis.

Katelijne A.E. Carbonez, Van T.T. Nguyen, Piet Sercu

DEPARTMENT OF ACCOUNTANCY, FINANCE AND INSURANCE (AFI)

Do Inventories Really Yield a Convenience? An Empirical Analysis of the Cost-adjusted Basis.*

Katelijne A.E. Carbonez[†]
Van T.T. Nguyen[‡]
Piet Sercu[§]

Tuesday 23rd February, 2010

Abstract

This paper considers four competing propositions to explain the convex relationship between inventories and the cost-adjusted basis called the ‘Working curve’: the convenience yield, the risk premium, data aggregation and the imbedded options value inherent to a futures contract. We use 70 years of weekly data from two different periods, 1885:01 to 1935:52 and 1989:01 to 2008:34, and we study contracts for storable agricultural commodities, namely wheat, corn and oats. The historical data are particularly interesting in this respect because they allow us to test the risk premium hypothesis, quality aggregation issues, and the potential impact of the imbedded options.

The convenience yield hypothesis emerges as having the best explanatory power for the ‘Working curve’. Although some of the other hypotheses predict a relationship between inventories and the cost-adjusted basis to some degree, the explanatory power is weak at best. For instance, the risk premium is often related to initial inventory levels but the result is not consistent across commodities and maturity months, the relationship is weak, and it does not subsume the strong negative-sloping/convex link with inventories that is present in the cost-adjusted basis. Similarly, we confirm empirically that data aggregation can indeed lead to a ‘Working curve’ even if the underlying disaggregated qualities show no such curve. Still, we also find a direct yield/inventory relation for the majority of disaggregated qualities, thus undermining the data aggregation hypothesis as the sole or major explanation. Finally, we show that also the imbedded option value is unlikely to explain the relationship between the cost-adjusted basis and inventories.

JEL classification: G12, N32, N52, Q11

Key words: backwardation, convenience yield, theory of storage, risk premium, data aggregation, imbedded options

*We gratefully acknowledge financial support from FWO and ICM, and help with the data from the USDA and dr. Scott Irwin. We also thank Thi Ngoc Tuan Bui and Fang Liu for useful suggestions. All remaining errors are the authors’.

[†]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

[‡]FWO PhD Fellow at Leuven School of Business and Economics, K.U.Leuven, Naamsestraat 69, B-3000 Leuven; +32 16 32 66 61; thituongvan.nguyen@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

Introduction

Backwardation, or ‘storage at a loss’, has been somewhat controversial among economists for many decades. Backwardation occurs when current spot prices adjusted for carrying charges such as storage costs and interest charges are trading at a positive spread over futures prices, i.e. when:

$$\psi_{t,T} := [S_t + \text{PV}(C, \mathbf{r}, t, T)](1 + r_{t,T}) - F_{t,T} > 0. \quad (1)$$

with $F_{t,T}$ the futures price at t for a contract maturing at T , S_t the current spot price, $S_t \times r_{t,T}$ the opportunity cost of investing in the spot market instead of the futures market and $\text{PV}(C, \mathbf{r}, t, T)$ the discounted costs of storage from t to T given interest rates \mathbf{r} . Familiarly, the futures price represents the risk-adjusted expected future spot price. In that light, a positive cost-adjusted basis ψ is somewhat puzzling because it seems to indicate a negative risk-adjusted return, i.e. inventory holders are willing to invest in inventory even though they expect to earn a negative risk-adjusted return on their investment.

Early researchers on backwardation (see work by Kaldor, 1939; Working, 1933, 1948 and 1949; Brennan 1958; and Telser 1958) have shown that the cost-adjusted basis is related to the level of inventories. When inventories are high, the cost-adjusted basis is zero—that is, there is no backwardation—but when inventories are low, the cost-adjusted basis is positive and backwardation occurs. This relationship between the cost-adjusted basis and inventory levels has been termed the ‘Working curve’ after Holbrook Working, who first demonstrated this empirical relationship.

A number of explanations have been derived since then to explain the shape of the Working curve: the convenience yield, the risk premium and mismeasurements due to data aggregation. The main explanation stems from the supply-of-storage theory first derived by Kaldor (1939) and Working (1948, 1949). The theory of storage builds upon the assumption of a convenience yield: inventory holders not only face storage costs, but also gain a convenience yield that results, among others, from avoiding search and replacement costs, from cost savings related to smoothing the production process or from a timing option. The convenience yield then is the missing link between the spot and futures prices and accounts for the shape of the Working curve: when inventories are low, the convenience yield is high and vice versa.

The risk premium theory, an alternative explanation for the shape of the Working curve, was developed by Keynes (1930) and Hicks (1939). They argue that hedgers, who are net short and have a strong desire to reduce price risks, have to pay a risk premium to the speculators,

who are net long and want to be compensated for the risk they undertake in providing a hedging instrument. In short, the holders of a long position expect the spot price at maturity to be above the futures price so that they can buy the commodity at the futures price and resell immediately at the spot price and make a profit equal to the ex post risk premium:

$$E_t(\tilde{\pi}_{t,T}) := E_t(\tilde{S}_T) - F_{t,T} > 0, \quad (2)$$

with $\tilde{\pi}_{t,T}$ denoting the ex post risk premium. The risk premium has been considered an alternative for the convenience yield in the past when researchers assumed that the best forecast of the future spot price is today's spot price, a relationship that is unlikely to hold. Nowadays, researchers on the risk premium theory rarely refer to the Working curve. However, researchers on the theory of storage typically discuss the risk premium as an alternative way to explain backwardation. This is why we also include the risk premium theory in this paper.

A third candidate explanation for the Working curve is data aggregation and transaction costs. Wright and Williams (1989) note that the observation that storage is related to backwardation might be an aggregation phenomenon. They show that averaging across different qualities can lead to a relationship like the Working curve even though there is no such relation if each quality is analyzed separately. Benirschka and Binkley (1995) and Brennan, Wright and Williams (1997) show that the same can hold true for location differentials. Chavas, Despins and Fortenberry (2000) similarly show that transaction costs can generate a Working-like curve.

We propose a fourth candidate explanation for the Working curve: the imbedded options value inherent in a futures contract. Consider an inventory holder that sells a futures contract. At maturity, he not only receives the futures price but also implicitly gets the option with a positive expected cashflow. If the imbedded options value is ignored, futures prices may appear to be in backwardation even though they are not.

Our empirical work supports the Working curve as a valid empirical relationship for storable commodity markets, more specifically for wheat, corn and oats in the Chicago area in both the periods 1885-1935 and 1989-2008. We also find that this Working curve is mostly driven by a convenience yield. We test the convenience yield hypothesis by exclusion: it is an unobservable variable and as such, it cannot be analyzed directly. To bypass that, we verify whether the risk premium, data aggregation and the imbedded option value can account for the observed strong relationship between the cost-adjusted basis and inventory levels, and we conclude that they fail to explain much of the cost-adjusted basis and certainly do not subsume the role of

inventories.

In the remainder of this introduction, we briefly summarize our contributions to the literature and our main results. We contribute to the literature firstly by setting up and analyzing an unusually extensive dataset, containing two periods, 1885-1935 and 1989-2008. Tests on the Working curve for the pre-World War II data are not new. The reason why we revisit the older sample is that it provides useful advantages relative to more recent data. First, there have been relevant changes in the characteristics of the data. The spot price is arguably of better quality in the older sample than in the modern sample. Notably, the daily settlement price was set in an organized (and liquid) exchange, the Chicago Board of Trade (CBOT). Nowadays, delivery in Chicago is rare, and spot prices differ substantially depending on where delivery takes place. Datastream even cautions the user about the CBOT spot prices, and provides, as a substitute, averages of the prices paid by the Central Illinois Country Elevators to the producers, recorded during the afternoons. Academics often use the nearest futures price as a stand-in for the spot price. Moreover, the interpretation of the inventory data has changed due to the waning role of Chicago as a warehousing center. In the past, most of the warehouses were in Chicago, and all deliveries were made in Chicago. Relatedly, inventory levels in Chicago accounted for a far higher proportion of nationwide totals than they do nowadays. This raises the question whether the Working curve was just a phenomenon of the past, a reflection of now-vanished circumstances, or whether it still exists today, and in what form. Second, in some respects the older sample contains more detailed information than is currently available. For example, the pre-World War II data contain spot prices for different qualities of each commodity, which allows us to directly test the empirical relevance of the data aggregation hypothesis and to better analyze the potential impact of the imbedded options value.

The second contribution comes from revisiting the Working curve. Our analysis of the modern sample tells us that a simple regression of the cost-adjusted basis on a cubic spline of inventory levels is too coarse a model and ignores more subtle features present in the Working relationship. We show that the cost-adjusted basis is better modeled using not just inventories but also time-to-harvest and time-to-maturity as discussed by Carbonez, Nguyen and Sercu (2008), both modeled as pure interaction terms. The Working curve still emerges as an empirical regularity in both periods, but it does change shape and curvature depending on time to maturity and time to harvest.

In the literature, the above hypotheses have typically been tested in isolation, and the samples have differed across studies. Our third contribution is that, to our knowledge, we are

the first to test all hypotheses simultaneously in the same (large) samples, and the first to ask the question which of them are subsumed by others. We are also the first, to our knowledge, to analyze whether the imbedded options inherent in a futures contract can account for the Working curve.

From our evidence, the convenience yield emerges as the dominant explanation of the Working curve, for the following reasons. (i) The results for the risk premium are mixed, and weak relative to the effects of inventory. It is true that the risk premium is (weakly) related to inventory levels and that the risk premium seems to explain part of the cost-adjusted basis, but this is not observed for all contracts and for all maturities. On the other hand, the risk premium does not subsume the (much stronger) link between the cost-adjusted basis and inventory levels when it is included as a variable in the regression. Thus, the risk premium explains, at best, only a tiny part of the Working effect. (ii) Our empirics confirm that when price and inventory data are averages or aggregates across qualities, a Working-like curve can effectively emerge even if it is absent, or weakly present only, in the underlying individual series. However, we also find Working curves for the majority of time series with disaggregated data, implying that aggregation cannot be the sole or major explanation. (iii) The delivery options imbedded in a futures contract are unlikely to explain the Working curve either. The imbedded options value depends on inventory levels at maturity and on variance and pairwise correlations, but none of these is well proxied for by inventories observed during the contract's life. Thus, the general conclusion is that the role of inventories seems to be largely direct rather than spurious—via risk premia, aggregation or option features.

This paper is structured as follows. Section 1 details the data used, revisits the Working curve and provides empirical evidence in favor of the relationship. Section 2 discusses the risk premium and shows empirically that the risk premium cannot account for the Working curve. Section 3 looks at the mismeasurement hypothesis and finds that quality and location do not change the shape of the Working curve. Section 4 analyzes the impact of the imbedded options (location, quality and timing) inherent in a futures contract and tests whether these imbedded options can account for the Working curve. Finally, Section 5 concludes.

1 Empirical Support for the Working Curve

1.1 Data

As stated before, we use an unusually comprehensive dataset to analyze the Working curve for storable agricultural commodities. The tests are performed using weekly CBOT spot, futures data and inventory data for corn, wheat and oats between 1885:01 and 1935:52,¹ and for corn and wheat between 1989:01 and 2008:34. The older sample is useful because there have been relevant changes in the characteristics of the data and because it contains more detailed information than is currently available. First, the spot price in the older sample is arguably of better quality than in the modern sample: the spot price used to be the daily settlement price from an organized (and very liquid) exchange, the Chicago Board of Trade (CBOT), while the published spot prices nowadays are the prices paid at the elevators (before 2000 for corn and 2008 for wheat) or barge loading locations along the Illinois river (after 2000 for corn and 2008 for wheat). As such, there is a location differential between the spot and futures prices. Moreover, spot prices these days are extremely illiquid, which may make prices more noisy. The second difference between the old and new data pertains to the role of Chicago, with implications for the interpretation of the inventory data. In the past, most of the warehouses were in Chicago, and all deliveries were centralized there. Relatedly, inventory levels in Chicago accounted for a far higher proportion of nationwide totals than they do nowadays. In 1885-1935 the inventory held in Chicago accounted, on average, to 33%, 15% and 24% of the total visible supply in the US for corn, wheat and oats respectively. In contrast, for the current period, Chicago's share in countrywide visible supplies has dwindled to 1.3% for wheat and 7.2% for corn. To get similar coverage as before, nowadays one would have to use Great Lakes Area inventory figures, but the commodity holdings reported in these data are, physically, scattered across the area and therefore subject to data aggregation issues compared to the ideal data, where spot and futures prices and inventory would all refer to the same location. Third, the old records contain data that are no longer published but are of great importance to study the impact of data aggregation on the 'Working curve'. More specifically, the old data contain detailed information on spot prices for various qualities of the same commodity quoted in parallel.

¹Please note that we exclude data from 1917 to 1920 because trading restrictions and price ceilings were in place in the aftermaths of the first World War.

More specifically, we use weekly data for wheat, corn and oats from 1885:01 to 1935:52 for the May, July, September and December contracts and weekly data for wheat and corn from 1989:01 to 2008:34 for the March, May, July, September and December contracts. Pre-world-war-II price data have been hand collected from the CBOT annual report microfilms. The spot prices and futures prices reported in the microfilm are the lowest and the highest price reached every day. The spot price is the quote for the standard quality deliverable on the futures contracts. We also have access to weekly low-high spot prices for various qualities traded on the CBOT. Due to the issues with the spot price nowadays, we follow the common practice in the literature and use the nearest futures as the spot price and the deferred futures as the futures price for the recent sample. The prices of the nearest futures and the deferred futures contracts are the daily settlement prices reported by the CBOT and were downloaded from Datastream. For the interest rate, we use the 3-Month libor rate for the 1989-2008 sample and the 3-Month interest rate provided by the NBER for the 1885-1935 sample. We deflate all prices by the US CPI so as to eliminate the near-unit-root problems present in the raw prices.² Weekly Chicago inventory data and storage costs have been collected from the microfilms for the pre-WWII sample. The recent Chicago inventory data pertain to inventory levels available in elevators cleared for delivery by the CBOT (the data have been kindly provided by the University of Illinois at Urbana-Champaign).

1.2 The base case

The theory of storage predicts a negative relationship between the inventory level and the cost-adjusted basis ψ . Without any economic analysis, the priors about the sign of ψ would be diffuse. Familiarly, in currency markets, arbitrage arguments can be invoked to show that only zero values are sustainable. In the case of commodities, however, arbitrage conditions and short selling restrictions in the spot market merely rule out negative ψ s. In fact, if going short would always be possible in both money and underlying cash markets, as is the case for currencies, arbitrage would ensure that ψ equals zero. Now borrowing money and buying spot should always be possible to a great many agents; so if the futures price would be at a premium relative to spot plus storage and interest cost, one would be able to borrow, buy spot, and sell forward at a profit until the negative ψ disappears. Eliminating a positive ψ , in contrast, could be problematic because going short in commodity markets is difficult or impossible whenever

²Upon request, more information on the data cleaning procedures is available in an extensive report.

inventories are tight. Thus,

$$\psi_{t,T} \geq 0 \text{ with } \begin{cases} \psi_{t,T} \approx 0 \text{ when inventories are abundant,} \\ \psi_{t,T} \geq 0 \text{ when inventories are tight.} \end{cases} \quad (3)$$

This has implications for the clientele attracted by spot v futures markets. An equally familiar result indeed is that, whenever ψ equals zero, both price-risk hedgers and speculators are indifferent between investing forward or buying and holding inventory financed by a loan. But if ψ is positive and shorting is impossible, then a speculator seeking a long position or a user who wants to reduce price risk will steer clear of the spot market and take position in the forward market instead, to avoid the assumedly positive premium ψ . The sole customers in the spot market would then be users that need to cover their consumption and safety stock for the upcoming period. They would be willing to pay a backwardation premium as long as the expected marginal benefit of holding inventory does not fall below the marginal expected costs of running out of stock. Obviously, the marginal benefit from holding inventory depends on the probability of an individual stock-out during the next period. Since the probability of a stock-out is a nonlinear function of the inventory, ψ is also a nonlinear function of the inventory. The marginal expected gain of not running out of stock during the next period is known at the beginning of the period, but obviously depends on circumstances. Because the probability of a stock-out decreases when the inventory increases, the more inventory there is in stock, the less positive the cost-adjusted basis is. This indicates a negative relationship between ψ and the inventory.

In the base-case regressions we first follow Gorton, Hayashi and Rouwenhorst (2007) to model the nonlinear cost-adjusted basis/inventory relationship. That is, $\psi_{t,t+1}$ is regressed on a cubic spline function of inventory. Unlike Gorton *et al.*, we take the inverse of inventory instead of inventory itself: in the proposed specification, a polynomial specification of $\psi_{t,t+1}$ automatically goes to zero when inventory levels go to infinity, and to infinity when inventory approaches a stock-out. Not surprisingly, also in terms of ex post goodness of fit this specification does better than a spline in inventory itself. In our test, the cubic spline function is:

$$\mathbf{f}(x, \mathbf{a}) = a_0 + a_1x^{-1} + a_2x^{-2} + a_3x^{-3} + a_4[\text{Max}(x^{-1} - k_1, 0)]^3, \quad (4)$$

where the knot point, k_1 , is a constant. We define k_1 later in the empirical section.

The inventory data and the prices used in the regressions need some pre-processing: quan-

Table 1: Adjusted R^2 of the regressions based on Equation 5 for the Wheat, Corn and Oats contracts in 1885-1935 and 1985-2005

| aR^2 | Sample 1: 1885-1935 | | | | Sample 2: 1985-2005 | | | | |
|--------|---------------------|------|------|------|---------------------|------|------|-----|-----|
| | May | Jul | Sep | Dec | Mar | May | Jul | Sep | Dec |
| Wheat | 19.1 | 14.0 | 9.4 | 10.6 | 20.4 | 14.0 | 0.0 | 2.3 | 0.7 |
| Corn | 26.5 | 30.5 | 13.4 | 8.2 | 1.1 | 5.0 | 17.3 | 7.1 | 0.4 |
| Oats | 8.5 | 1.2 | 3.2 | 21.3 | na | na | na | na | na |

Note: The table gives the adjusted R^2 of the following regression:

$$\mathbf{f}(x, \mathbf{a}) = a_0 + a_1 x^{-1} + a_2 x^{-2} + a_3 x^{-3} + a_4 [\max(x^{-1} - k_1, 0)]^3, \quad (5)$$

for the contracts of wheat, corn and oats. Only the significant terms of the regressions have been retained. k_1 is equal to 10 for all the contracts of the 1885-1935 period, and to 10 and 60 for the corn and wheat contracts of the 1989-2008 period respectively.

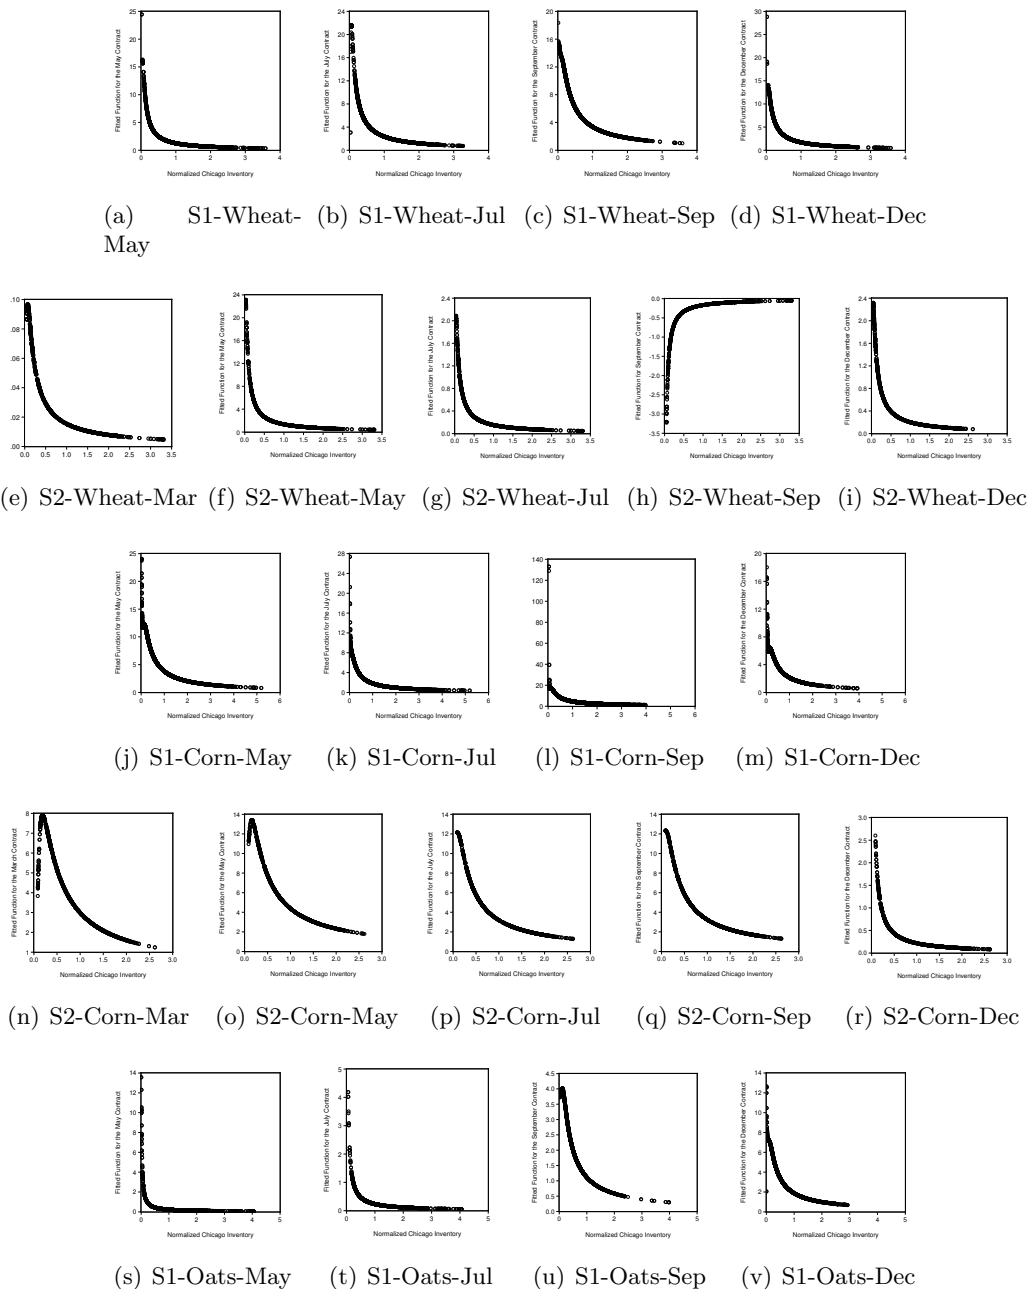
ties are normalized,³ and prices deflated. Normalized inventory equals inventory divided by its Hodrick-Prescott trend. On the basis of eyeball tests we settled for a smoothness parameter of $8E+8$. To interpret this, one can refer to the standard values for the smoothness parameter λ_q in quarterly series: 1,600 for quarterly series with peak-to-peak cycles of short duration (roughly 10 year), and 160,000 for cycles of about 30 years or longer. Gorton *et al.* (2007), for instance, use 160,000. Correcting for the effect of frequency (we have weekly instead of quarterly data), our λ is in-between these standards.⁴

The results for each commodity, each period and each contract are given in Figure 1. Overall, it looks like the Working curve is an empirical fact both in the past as in the present for the wheat, corn and oats markets in the Chicago area. A closer look at Figure 1 also shows anomalous results, though: one relation between the cost-adjusted basis and inventory relationships (wheat, september contract) is inverted relative to what (3) predicts and some working curves decline again when inventory achieves a critical value. Table 1 also shows that the inventory has a much lower explanatory power for inventories in the 1989-2008 sample than in the 1885-1935 sample. We thought of two possible quick-fix explanations for the fact that the Working curve seems to break down for some contracts in the recent data. First, the 1989-2008 sample may be contaminated by outlier-plagued cost-adjusted bases and thus to erratic Working-curve estimates. Most notably, in the 1989-2008 sample the cost-adjusted bases have a strange

³As a robustness check, we also perform the tests using raw inventories. The results are consistent for raw inventories and are available from the authors upon request.

⁴To translate λ_q into an equivalent for weekly data, Ravn and Uhlig (2002) recommend multiplication by the fourth power of the relative frequency. Taking 4 weeks in a month, we'd get $\lambda_w = \lambda_q \times 12^4 = 50,000 \times 20,736$, implying an equivalent quarterly smoothness of $\lambda_q = 38,500$, nicely in between the 1600 and 160,000 standards for 10- and 30-year cycles, respectively.

Figure 1: The average Working curve for the Wheat and Corn contracts



Note: The graphs show the fitted value of the regression of the cost-adjusted basis on the inverse of the normalized inventory $f(x, \mathbf{a})$ from Equation 5. The graphs are split by commodity (wheat, corn or oats), by contract (March, May, July, September or December) and by sample (S1: 1885-1935, S2: 1989-2008). We deleted datapoints with a normalized inventory lower than 0.025 for the contracts in S2. Though the results are basically the same, the graphs were difficult to read due to outliers arising from stock-out effects. k_1 is equal to 10 for all contracts of the pre-WWII sample and equal to 10 and 60 for the corn and wheat contracts of the recent sample.

distribution compared to the earlier sample: on the one hand, the cost-adjusted bases are very right-skewed; on the other hand, they are often negative which seems to indicate a breakdown in arbitrage. The peaking cost-adjusted bases are concentrated in the period 9/1995 to 10/1996 and correspond to peaking spot prices while futures prices still predicted a return to normality. A breakdown in arbitrage in a market that has become increasingly professional seems very unlikely. However, the period 2000-2008 has also been plagued by convergence problems which may have affected results. We test the above conjectures by deleting the period with peaking spot prices from the sample and re-estimating the curve on the one hand and by deleting the periods where arbitrage seems to breakdown on the other hand. Both actions lead to a better fit of the cubic spline to the cost-adjusted basis but do not solve the inverted Working curve problem.⁵

Our second conjectured explanation is that the inventory data may have been a noisy proxy of a true variable. In fact, Chicago's low share of regional or national inventory levels may mean that it is now less representative of what really is going on. For instance, if Chicago holdings are temporarily low but the traders know that there is plenty available in nearby silos, spot prices will not increase sharply compared to the futures price and thus the cost-adjusted basis will not surge, or not surge as much as the ignorant econometrician expects. We check this conjecture by adopting the inventory level in the Great Lakes area in our estimations. Again, this does not improve the graphs for the September contract of wheat in the 1989-2008 sample.⁶

In short, neither of the above two factors can explain the atypical shape of the Working curve for the recent sample. Our third conjectured explanation starts from the fact that the above Working curve models are rather rough. The simple spline model ignores time to maturity, an obvious weakness. The simple solution of adding this variable as an independent 'main effect', in the parlance of variance analysis, is not necessarily adequate, though; for instance, many papers model cost-adjusted bases as a per annum percentage, implying that the log of the spread is modeled as the product of a spline in x and time to maturity. A second objection to an additive model for time to maturity is that n million bushels m months prior to expiration may mean something very different depending on whether the next harvest comes during, or right after, or far after the contract's life. A less *ad hoc* model of inventory and time

⁵Results are available from the authors upon request

⁶Results are available from the authors upon request

to maturity may be better than the simple traditional models at explaining the cost-adjusted bases.

Carbonez, Nguyen and Sercu (2008) propose a more careful model. They show that the convenience yield over the contract's life $T - t$ is the result of capitalizing all expected period-by-period convenience yields towards T . They next assume that, at any moment s , the term structure of expected one-period forward convenience yields, $E_t(\tilde{\psi}_{s,s+1})$, is rising or falling more or less smoothly, except if there is a harvest during the contract's life. If there is such an intervening harvest, at time T_h , the level of $E_t(\psi_{s,s+1})$ is allowed to change discretely at $s = T_h$, and a new path of expected one-period yields starts. Lastly, the initial one-period convenience yield, $\psi_{t,t+1}$, is assumed to be a function of the inventories prevailing at t , and so is the expected one-period convenience yield right after any intervening harvest time. The result, after a partial linearization, is a model featuring two functions of x , one being weighted by the contract's total remaining life $T - t$, and the second capturing the discontinuity at the intervening harvest time, if any, weighted by the contract's remaining post-harvest life:

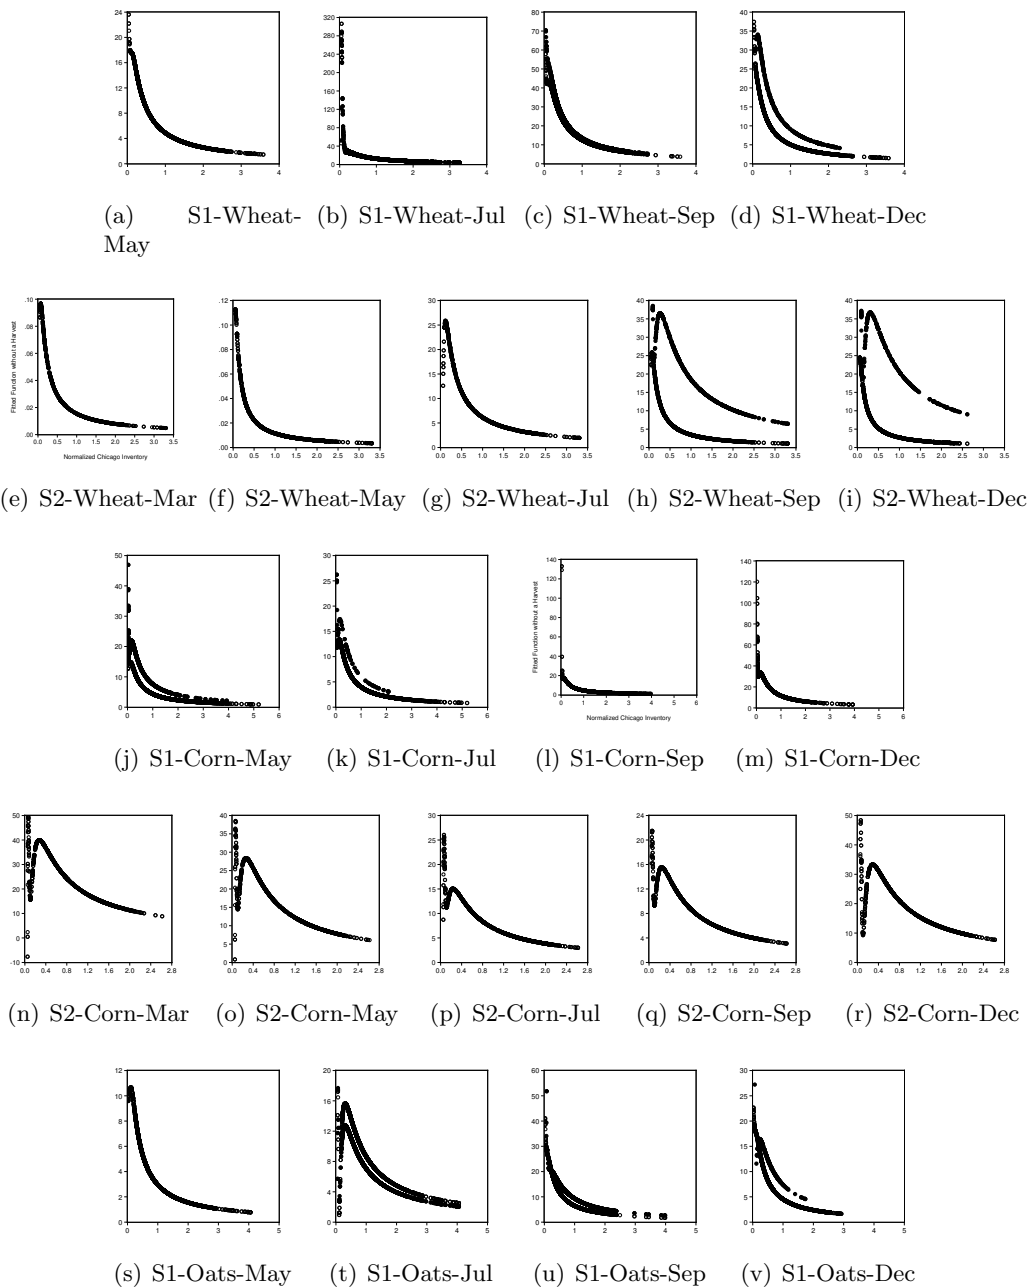
$$\psi_{t,T} = f(x_t)(T - t) + g(x_t) \max(T - T_h, 0). \quad (6)$$

The usual Working-curve pattern (positive-valued, negative-sloped and convex) should apply to the entire expression and to $f(x)$; for $g(x)$, which is actually a difference between two Working curves, we have no priors.⁷

Figure 2 displays the revised 'Working curves' for each commodity, each contract and each sample period based on the above specification. Each subfigure shows the graphs for two different times to maturity: one before the harvest and one after the harvest. The following conclusions stand out. First, though the recent cost-adjusted bases are still occasionally negative, taking into account time to maturity and harvest timing seems to eliminate the inverted curves for the recent wheat data. Second, the 'Working curve' seems to change shape and curvature during the contract life. Unfortunately, we still find working curves that decline when inventories are very low in the recent sample.⁸ Table 2 shows how the adjusted R^2 improves dramatically compared to Table 1 when the interaction term with time to maturity and harvest

⁷To complete the model, we also add seasonal dummies to the regressions in a next step. The reason is that reported Chicago inventories turn out to be peaking not right after harvest but some months later (see Figure 4 in Section 4). Given that there is some time lag between harvest and actual inventories in Chicago, scarcity cannot be fully described by the reported inventory levels. Just after harvest, inventories will be very low while agents already know that scarcity has been resolved, which in its turn will affect prices and convenience yields. Thus, we predict negative dummies for the first months after the harvest, correcting for the effect of underreported availabilities. Working curves still arise when seasonals are included. (Results are available from

Figure 2: The revised Working curve for the Wheat, Corn and Oats contracts



Note: The graphs show the fitted value of the cost-adjusted bases based on Equation (6). The graphs are split by commodity (wheat, corn or oats), by contract (March, May, July, September or December) and by sample (S1: 1885-1935, S2: 1985-2008).

Table 2: Adjusted R^2 of the regressions based on Equation 6 for the Wheat, Corn and Oats contracts in 1885-1935 and 1989-2008

| aR^2 | Sample 1: 1885-1935 | | | | Sample 2: 1989-2007 | | | | |
|--------|---------------------|------|------|------|---------------------|------|------|------|------|
| | May | Jul | Sep | Dec | Mar | May | Jul | Sep | Dec |
| Wheat | 34.3 | 30.4 | 35.9 | 33.9 | 44.0 | 22.2 | 8.0 | 18.2 | 34.5 |
| Corn | 60.2 | 52.5 | 26.0 | 39.3 | 31.6 | 33.5 | 39.4 | 11.4 | 15.6 |
| Oats | 33.8 | 37.9 | 43.0 | 54.7 | na | na | na | na | na |

Note: The table gives the adjusted R^2 of the following regression:

$$\psi_{t,T} = f(x_t)(T-t) + g(x_t) \max(T - T_h, 0), \quad (7)$$

for the contracts of wheat, corn and oats. Only the significant terms of the regressions have been retained. k_1 is equal to 10 for all the contracts of the 1885-1935 period, and to 10 and 60 for the corn and wheat contracts of the 1989-2008 period respectively.

timing is included in the model.

2 The Risk Premium is Not Related to Inventories

In the preceding section, we have shown that the Working curve is still an empirical fact in storable agricultural commodity markets, both in pre-world-war-II markets and in more recent markets. The question is whether this empirical relationship is really due to a convenience yield, rather than being the reflection of a risk premium, data aggregation and/or imbedded option values.

In this section, specifically, we analyze whether the Working curve may be just the reflection of a risk premium. As discussed above, the risk premium hypothesis assumes that hedgers, who are net short and have a strong desire to reduce price risks, have to pay a risk premium to the speculators, who are net long and want to be compensated for the risk they undertake in providing a hedging instrument. In short, the holders of a long position expect the spot price at maturity to be above the futures price so that they can buy the commodity at the futures price and resell immediately at the spot price and make a profit equal to the risk premium.

From the above, it must be obvious that the risk premium per se cannot explain the Working curve. Indeed, the risk premium relates the futures price to the expected future spot price, not the current one. The risk premium literature has been linked to the Working curve because Keynes (1930) stated that ‘if supply and demand are balanced, the spot price must exceed the forward price by the amount which the producer is ready to sacrifice in order to

the authors upon request.

⁸When we pool the regressions of all maturities, we find a perfect working curve for wheat, corn and oats.

Table 3: Estimates of Variance Ratios $VR(q)$ and Variance-ratio Test Statistics $Z(q)$ and $Z^*(q)$

| | | Sample: 02/01/1885 to 10/11/1928 | | | | | | |
|-------|-------|--|--------|--------|--------|--------|--------|--------|
| | | Number q of base observations forming variance ratio | | | | | | |
| | | 30 | 60 | 90 | 120 | 150 | 180 | 180 |
| Wheat | VR | 0.60 | 0.50 | 0.45 | 0.41 | 0.37 | 0.34 | 0.32 |
| | Z | -7.35* | -6.47* | -5.77* | -5.37* | -5.11* | -4.89* | -4.70* |
| | Z^* | -2.60* | -2.69* | -2.67* | -2.68* | -2.70* | -2.69* | -2.69* |
| Corn | VR | 1.00 | 0.93 | 0.92 | 0.92 | 0.89 | 0.84 | 0.78 |
| | Z | -0.03 | -0.90 | -0.83 | -0.70 | -0.93 | -1.15 | -1.52 |
| | Z^* | -0.01 | -0.42 | -0.43 | -0.39 | -0.55 | -0.71 | -0.95 |
| Oats | VR | 0.90 | 0.86 | 0.85 | 0.80 | 0.76 | 0.73 | 0.70 |
| | Z | -1.82 | -1.88 | -1.57 | -1.80 | -1.97* | -2.03* | -2.05* |
| | Z^* | -0.65 | -0.84 | -0.78 | -0.97 | -1.13 | -1.22 | -1.29 |

Notes: Variance-ratio test of the random walk hypothesis for daily real commodity price series. The variance ratio $VR(q)$ is defined as $\frac{\sigma_c^2(q)}{\sigma_a^2(q)}$; where $\sigma_c^2(q)$ is an unbiased estimator of $\frac{1}{q}$ of the variance of the q -th difference and $\sigma_a^2(q)$ is an unbiased estimator of the variance of the first difference of the real commodity price. The variance ratios $VR(q)$ are reported as well as the test statistic $Z(q)$ and the heteroscedasticity-robust test statistic $Z^*(q)$, which test the null hypothesis that $VR(q)$ equals one. * denotes rejection of the random walk at the 5% level.

“hedge” himself, i.e. to avoid the risk of price fluctuations during his production period. Thus in normal conditions the spot price exceeds the forward price, i.e. there is backwardation.’ Note that Keynes implicitly assumes that the best predictor of the future spot price is the current spot price, which then links the future risk premium to the current price. Formally, if \tilde{S}_T denotes the spot price at maturity T , $F_{t,T}$ the current t price of a futures contract maturing at T and $\tilde{\pi}_{t,T}$ the ex post risk premium, we have


$$0 < E_t(\tilde{\pi}_{t,T}) := E_t(\tilde{S}_T) - F_{t,T} > 0, \quad (8)$$

$$= S_t - F_{t,T} \text{ provided that } E_t(\tilde{S}_T) = S_t. \quad (9)$$

The martingale assumption is a priori unconvincing since we know that spot prices are seasonal, mean reverting, and related to storage costs and interest charges; however, deviations from martingale behavior may still not be so strong as to totally eclipse the hypothesized risk-premium effect. Unfortunately, tests on the risk premium can only be performed on the pre-WWII sample as recorded spot data for the recent sample are flawed.

To test the martingale assumption, we perform a series of Variance Ratio tests and test whether commodity spot prices are random walks. Following Liu and He (1991), we study variances of sums over various time windows. We allow the window size q to range from 30 days to 210 days because it allows for the kind of observation interval we are interested in. Thus, by using one day as our base observation interval, Z - and Z^* -statistics are calculated for each q by comparing the variance of the base interval with that of 30-day, 60-day, 90-day, 120-day, 150-day, 180-day and 210-day observation intervals. Interestingly, the random walk cannot be rejected in the majority of the cases, certainly when heteroscedasticity is accounted

Figure 3: Seasonality patterns of inventory levels for wheat, corn and oats



Note: Given that inventory data are volatile and have a trend, it is not possible to measure seasonality on the raw inventory data. Instead, we compute the deviation from the long-term trend $((\text{raw inventory} - \text{trend})/\text{trend})$ and fit monthly dummies to this deviation to find the seasonality. S1: 1885-1935.

for. This makes an analysis of the risk premium potentially relevant in the context of this paper.

An additional reason to test the risk premium hypothesis is the empirical and theoretical finding of Gorton, Hayashi and Rouwenhorst (2007) that risk premia of commodity futures vary with the state of inventories. Their predictions, though derived from the theory of storage, seem to provide evidence in favor of the risk premium theory in explaining the Working curve. However, a closer look at their model shows that the issue is far from settled. What the authors show is that the state of inventory at the end of the month is a key predictor of the excess return from the end of the month to the next and that the mean excess return and inventory are inversely related. Their argument is that when inventories are low, the ability of inventories to absorb demand and supply shocks is reduced, thus raising the conditional volatility of future spot prices. If the risk premium is indeed compensation paid by the hedgers to the speculators for providing insurance against price risk (and if this risk is not fully diversifiable to the speculator), the mean excess return from commodity futures should increase when future spot price risk increases. As we have seen, inventory levels are seasonal and relatively sticky. Thus, when looking at risk premia only one month ahead, as Gorton *et al* do, the level of inventories does not change dramatically (except of course around harvest time). As a result, current inventories may actually be proxying for inventories at maturity, one month later. This is a testable conjecture. Specifically, when we look at the relationship for longer times to maturity, the state of inventories should typically change more substantially; for example, the future spot-price risk at T is unlikely to increase even when inventories are low at t if there is harvest prior to T . We test our conjecture by performing similar tests to Gorton *et al* (2007) for various holding-period lengths. The first test is a regression test: the n -day ex post risk premium is

regressed on normalized inventory levels up to six months before the maturity of the contract, and the results are compared to similar regressions for the convenience yield:

$$\pi_{T_i-n, T_i} := S_{T_i} - F_{T_i-n, T_i} = \alpha_n + \beta_n x_{T_i-n} + e_i, \quad (10)$$

$$\psi_{T_i-n, T_i} := [(S_{T_i-n} - \text{PV}(C))(1 + r_{T_i-n, T_i}) - F_{T_i-n, T_i}] = \alpha'_n + \beta'_n x_{T_i-n} + e'_i, \quad (11)$$

where x is the inverse of normalized inventory. Note that every calendar year i provides one observation, so that there is no issue of overlapping holding periods.

These regressions can be run per contract or after pooling across all contracts, with the aim of obtaining an estimate and a confidence interval for the average regression coefficient.⁹ Since in these tests each time series has at most 51 observations, the power of the contract-by-contract regressions is low; we accordingly discuss the pooled regressions, and relegate the results per contract to Appendix B. The results are summarized in Table 4. From the pre-war data a very clear picture emerges: for all commodities, yields are strongly related to inventories, but risk premia are not (wheat and oats) or only weakly so (corn). This implies that in our earlier regressions, inventories are unlikely to have proxied for the risk premium.

The second test, also modified from Gorton *et al*, is a conditional means test which can also be interpreted as a trading-rule test. We sort the observations on the basis of normalized inventory level, and check whether the corresponding risk premia are significantly different, on average, depending on whether inventory is more than 50 percent above or 50 percent below the normalized level. As before, we test this strategy for holding periods up to six months before the maturity of the contract. The results, in Table 5, do not provide any strong evidence of profits from taking a long (short) futures position when inventories are low (high). Although the mean is always higher for the low inventory case, the standard deviation of the mean is higher as well, and the resulting t-ratios on the difference of the two means are insignificant. This observation holds for all maturities considered and for both samples.

Finally, to strengthen our results so far, we compare the regression results of fitting a Working curve to the risk premium and the cost-adjusted basis based on Equation (6). The results are in Tables 6. We also fit a regression of Equation (6) with the ex post risk premium

⁹According to Pesaran and Smith (1995), there are four procedures that can be used to estimate this average effect: the mean group estimator (estimating separate regressions for each group and averaging the coefficients over groups), pooled regression, aggregate time-series regressions, and cross-section regressions on group means. In the static case, where the regressors are strictly exogenous and the coefficients differ randomly and are distributed independently of the regressors across groups, all four procedures provide a consistent and unbiased estimate of the coefficient means (Zellner, 1969).

Table 4: Does the risk premium depend on inventory levels as the cost-adjusted basis? Evidence for wheat, corn and oats

| $T - t$ | Sample: 1885-1935 | | | | | | | | | |
|---------|-------------------|-------|---------|---------|-------|---------|-------|-------|---------|----|
| | Wheat | | | Corn | | | Oats | | | |
| | RP | R^2 | ψ | RP | R^2 | ψ | RP | R^2 | ψ | |
| 0w | 0.44*** | 19 | 0.47*** | 0.06 | 08 | 0.05 | -0.00 | 05 | -0.00 | 06 |
| 2w | 0.22 | 02 | 0.74*** | 0.05 | 05 | 0.18*** | 0.02* | 06 | -0.00 | 06 |
| 4w | -0.02 | 02 | 0.28*** | 0.23*** | 05 | 0.52*** | 0.02 | 04 | 0.17*** | 22 |
| 6w | 0.03 | 02 | 0.42*** | 0.14 | 02 | 0.43*** | -0.00 | 03 | 0.07*** | 48 |
| 8w | 0.29* | 02 | 1.02*** | 0.31* | 03 | 0.61*** | -0.02 | 04 | 0.49*** | 34 |
| 10w | 0.49* | 03 | 1.33*** | 0.06 | 01 | 0.27*** | -0.08 | 03 | 0.23*** | 35 |
| 12w | 0.63** | 04 | 1.32*** | 0.06 | 02 | 0.31*** | -0.05 | 04 | 0.00 | 34 |
| 14w | 0.22 | 02 | 1.65*** | 0.40** | 06 | 0.60*** | -0.00 | 04 | 0.00 | 41 |
| 16w | 0.73 | 02 | 2.52*** | 0.08 | 05 | 0.32*** | -0.01 | 02 | 0.08*** | 41 |
| 18w | 0.03 | 01 | 0.75*** | 0.66*** | 10 | 0.70*** | 0.01 | 01 | 0.03*** | 29 |
| 20w | -0.02 | 01 | 3.45*** | 0.35** | 06 | 0.22** | 0.01 | 03 | 0.51*** | 31 |
| 22w | 0.31 | 01 | 3.13*** | 0.20 | 03 | 0.14* | -0.02 | 03 | 0.38 | 33 |
| 24w | 0.18 | 00 | 3.49*** | 0.41 | 03 | 0.43*** | 0.28 | 02 | 0.91*** | 39 |
| 26w | 0.56 | 01 | 5.44*** | 0.58** | 06 | 0.80*** | 0.45 | 10 | 0.40 | 42 |
| 28w | 0.29 | 01 | 4.08*** | 0.23 | 02 | 0.32*** | 0.07 | 06 | 0.17 | 42 |
| 30w | 0.07 | 00 | 2.05*** | 0.24 | 04 | 0.10 | 0.18 | 08 | 0.09 | 49 |
| 32w | 0.12 | 01 | 1.23*** | 0.21 | 10 | 0.18* | 0.12 | 05 | 0.01 | 49 |
| 34w | 0.15 | 02 | 0.74*** | 0.34* | 14 | 0.30* | 0.00 | 13 | 0.00 | 59 |

Notes: To find the above results, we regress the discounted risk premium and the cost-adjusted basis on the inverse of normalized inventories for up to 34 weeks before the maturity of the contract. We perform a different regression for each time to maturity. More specifically, we regress:

$$\delta_{t,T} = \alpha + \beta \cdot \frac{1}{norminv_t} + \varepsilon_t, \quad (12)$$

where $\delta_{t,T}$ is either the risk premium $\pi_{t,T}$ or the cost-adjusted basis $\psi_{t,T}$. The regressions are cross-sectional and contemporaneous, thus there are no overlapping observations in this analysis. We also pool the results for the May, July, September and December contract for ease of exposition. *, ** and *** denotes significance at the 10%, 5% and 1% level respectively.

as an additional regressor. The results clearly show that the risk premium cannot explain the Working curve. Although both the risk premium and the cost-adjusted basis are related to inventories, the former relationship is not consistent across contracts and commodities. In fact, often the risk-premium/inventory relation does not even have the expected sign. Equally interesting is the fact that the relation between the risk premium and inventories is weak, as evidenced by the adjusted R^2 , and that it does not subsume the relationship between the cost-adjusted basis and inventories when both are included into one regression. Thus, we conclude that the risk premium is not a major factor behind the Working curve.

3 Disaggregated Data Still Exhibit a Working Curve

In this section, we analyze whether the Working curve may be an artefact of data aggregation. We discuss both location aggregation and quality aggregation.

In neither of our samples, location aggregation can be an issue. Indeed, we estimate the Working curve from prices quoted on the Chicago Board of Trade and for inventory levels available in the city of Chicago, not in a wider area. What has changed over time is the representativeness of the Chicago storage data. Before the second World War, Chicago was one of the main grain warehouses of the United States. In the nineteenth century, grain was shipped by the farmer to Chicago to be stored in one of the city's public or private elevators. The farmer then received an elevator receipt to claim his grain when he needed it.¹⁰ This system became less important in the course of the twentieth century with the farmer storing grain on the farm as well, but Chicago kept attracting a lot of grain from the countryside. For the recent data, the picture is a little different in the sense that Chicago is no longer a large warehousing center. Thus, our recent inventory data must be far less representative for true availabilities than before. But they remain free of aggregation bias. Actually, if despite the resulting errors-in-variables bias we still find a Working curve for prices quoted on the CBOT and inventories stored in Chicago, our conclusion that yield is related to scarcity becomes quite conservative.

Thanks to the wealth of our database in the period 1885 to 1935, we can test whether quality aggregation is an issue. First we review how quality aggregation can induce a spurious

¹⁰Originally the grain was stored by the owner. Later on, to rationalize things, grain was inspected and stored by quality type.

Table 5: Can we make a profit by taking a long futures position when inventories are low as compared to when inventories are high?

Sample: 1885-1935

| $T - t$ | Wheat | | | | | Corn | | | | | Oats | | | | | |
|---------|---------|------------|---------|------------|----|---------|------------|---------|------------|------|---------|------------|---------|------------|-------|------|
| | Mean | | St. dev | | t | Mean | | St. dev | | t | Mean | | St. dev | | t | |
| | < 0.5 | ≥ 1.5 | < 0.5 | ≥ 1.5 | | < 0.5 | ≥ 1.5 | < 0.5 | ≥ 1.5 | | < 0.5 | ≥ 1.5 | < 0.5 | ≥ 1.5 | | |
| 0w | 3.71 | 1.50 | 4.94 | 2.59 | 67 | 42 | 0.40 | 0.24 | 2.06 | 0.87 | 1.32 | 3.98 | 1.32 | 72 | 37 | 0.28 |
| 2w | 2.67 | 0.47 | 7.83 | 5.17 | 64 | 46 | 0.40 | 0.24 | 1.95 | 0.74 | 4.99 | 2.58 | 66 | 41 | 0.22 | |
| 4w | 1.25 | 0.07 | 9.21 | 6.09 | 59 | 36 | 0.11 | 0.36 | 3.20 | 0.64 | 6.31 | 3.41 | 56 | 38 | 0.36 | |
| 6w | 1.09 | 0.40 | 10.72 | 5.56 | 70 | 36 | 0.06 | 0.27 | 3.54 | 1.17 | 6.85 | 5.77 | 53 | 44 | 0.27 | |
| 8w | 0.19 | 0.18 | 10.78 | 6.09 | 57 | 38 | 0.00 | 0.31 | 2.98 | 0.43 | 6.07 | 5.65 | 33 | 46 | 0.31 | |
| 10w | 2.56 | -1.14 | 11.28 | 7.65 | 66 | 43 | 0.27 | 0.16 | 2.25 | 0.77 | 7.09 | 6.07 | 42 | 54 | 0.16 | |
| 12w | 3.40 | -1.08 | 10.50 | 9.21 | 59 | 42 | 0.32 | 0.01 | 1.00 | 0.86 | 7.61 | 6.37 | 42 | 41 | 0.01 | |
| 14w | 2.20 | 0.14 | 10.67 | 8.46 | 50 | 40 | 0.15 | 0.13 | 3.21 | 1.64 | 9.56 | 7.32 | 47 | 39 | 0.13 | |
| 16w | 2.22 | -1.06 | 10.99 | 8.68 | 52 | 37 | 0.24 | 0.09 | 1.49 | 0.48 | 8.23 | 7.04 | 38 | 41 | 0.09 | |
| 18w | 0.89 | 0.48 | 10.87 | 9.24 | 54 | 29 | 0.03 | 0.05 | 2.19 | 1.56 | 8.49 | 8.29 | 36 | 42 | 0.05 | |
| 20w | 1.86 | -0.93 | 11.44 | 9.67 | 61 | 38 | 0.19 | 0.20 | 3.12 | 0.67 | 8.83 | 8.70 | 44 | 41 | 0.20 | |
| 22w | 3.47 | 0.90 | 11.06 | 9.43 | 51 | 34 | 0.18 | 0.16 | 3.75 | 1.77 | 9.02 | 8.35 | 45 | 34 | 0.16 | |
| 24w | 4.35 | 0.50 | 12.80 | 9.15 | 44 | 35 | 0.25 | 0.10 | 3.43 | 2.27 | 8.65 | 7.94 | 45 | 33 | 0.10 | |
| 26w | 1.53 | -0.14 | 13.37 | 10.18 | 30 | 30 | 0.10 | 0.24 | 3.99 | 0.84 | 9.76 | 9.13 | 37 | 24 | 0.24 | |
| 28w | 2.91 | 0.61 | 15.29 | 9.28 | 30 | 30 | 0.13 | 0.19 | 5.64 | 2.86 | 11.61 | 8.81 | 43 | 21 | 0.19 | |
| 30w | 5.12 | 2.05 | 13.05 | 11.57 | 26 | 28 | 0.18 | 0.27 | 5.70 | 1.69 | 11.49 | 10.22 | 34 | 15 | 0.27 | |
| 32w | 3.86 | 0.69 | 13.23 | 11.84 | 22 | 28 | 0.18 | -0.06 | 4.84 | 5.67 | 10.94 | 9.39 | 38 | 11 | -0.06 | |
| 34w | 2.17 | -1.19 | 12.97 | 12.52 | 12 | 22 | 0.19 | 0.22 | 4.56 | 2.17 | 9.51 | 5.79 | 34 | 5 | 0.22 | |

Notes: In this table we compare the discounted risk premium that can be gained by taking a long position when normalized inventories are lower than 0.5 and when they are larger than or equal to 1.5. We give mean and standard deviation of the risk premium for up to 34 weeks until maturity to analyze whether the timing of the position can have an effect. We then compare the means and compute the t-statistic for the difference of the means. We use a simple t-statistic for the difference of two means with varying standard deviation. Since the data are non-overlapping and we compare risk premia in different years or different months, we do not need to correct for autocorrelation as in Gorton, Hayashi and Rouwenhorst.

Table 6: Does the risk premium explain the behavior of the cost-adjusted basis? evidence for wheat, corn and oats between 1885 and 1935

| Panel A: Wheat | | | | | | | | | | | | |
|-----------------|------------|----------|------------|-----------|------------|----------|------------|------------|------------|----------|-----------|--|
| | May | | Jul | | Sep | | Dec | | | | | |
| | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M3 | |
| $(T-t) * x_1$ | 0.007*** | | 0.001*** | | 0.019*** | | 0.022*** | | 0.008*** | | 0.008*** | |
| $(T-t) * x_2$ | -3.2e-4*** | -1.5e-5* | 0.001*** | 4.8e-4*** | -3.6e-4*** | -0.003** | -0.003** | -3.6e-4*** | -1.2e-4*** | -7.3e-5* | 0.008*** | |
| $(T-t) * x_3$ | 3.8e-6*** | | -7.8e-5*** | | 0.001*** | | 9.9e-5*** | | | | 1.2e-4*** | |
| $(T-t) * x_4$ | | | 0.056*** | | -7.8e-5*** | | -9.9e-5*** | | | | | |
| $(T_h-t) * x_1$ | | | -0.003*** | | 2.7e-6*** | | -9.9e-5*** | | 0.005* | | 0.005* | |
| $(T_h-t) * x_2$ | | | | | 0.056*** | | -0.007** | | | | | |
| $(T_h-t) * x_3$ | | | | | -0.003*** | | 5.9e-4*** | | | | | |
| $(T_h-t) * x_4$ | | | | | -0.002*** | | -0.002** | | | | | |
| $Mthdum$ | ✓ | ✓ | ns | ✓ | ns | ✓ | ns | ✓ | ✓ | ✓ | ns | |
| αR^2 | 36.66 | -0.29 | 33.90 | 2.86 | 33.90 | 43.09 | 2.85 | 43.09 | 41.33 | 0.82 | 41.33 | |

| Panel B: Corn | | | | | | | | | | | | |
|-----------------|------------|----------|------------|----------|------------|----------|-----------|----------|------------|--------|------------|--|
| | May | | Jul | | Sep | | Dec | | | | | |
| | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M3 | |
| $(T-t) * x_1$ | 0.011*** | | 0.007*** | | 0.013*** | | 0.011*** | | 0.006*** | | 0.021*** | |
| $(T-t) * x_2$ | -0.001*** | 0.001*** | -0.001*** | 0.002*** | -0.001*** | 0.010*** | -0.001*** | 0.010*** | -0.002*** | 0.003* | -0.002*** | |
| $(T-t) * x_3$ | 4.7e-5*** | | 2.5e-5*** | | 4.1e-5*** | | -3.0e-5* | | 8.7e-5*** | | 8.7e-5*** | |
| $(T-t) * x_4$ | -4.8e-5*** | | -2.5e-5*** | | -4.4e-5*** | | 8.4e-5** | | -2.2e-4*** | | -9.1e-5*** | |
| $(T_h-t) * x_1$ | 0.005 | | -0.003*** | | -0.002*** | | - | | 0.031*** | | 0.031*** | |
| $(T_h-t) * x_2$ | -0.001*** | | 2.4e-4*** | | -0.002*** | | 0.051*** | | -0.005*** | | -0.005*** | |
| $(T_h-t) * x_3$ | 4.6e-5*** | | 3.4e-4*** | | 2.3e-4*** | | ✓ | | 2.0e-4*** | | 2.0e-4*** | |
| $(T_h-t) * x_4$ | -6.2e-5*** | | - | | 0.065*** | | - | | -2.2e-4*** | | -2.2e-4*** | |
| $Mthdum$ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | |
| αR^2 | 66.51 | 2.73 | 55.94 | 6.91 | 54.75 | 28.85 | 7.46 | 30.44 | 43.46 | 4.19 | 50.32 | |

| Panel C: Oats | | | | | | | | | | | | |
|-----------------|------------|-----------|------------|----------|------------|-----------|------------|-----------|------------|-----------|------------|--|
| | May | | Jul | | Sep | | Dec | | | | | |
| | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 | M3 | |
| $(T-t) * x_1$ | 0.003*** | | 0.005*** | | 0.018*** | | 0.018*** | | 0.008*** | | 0.008*** | |
| $(T-t) * x_2$ | -2.5e-4*** | -0.006*** | -7.2e-4*** | -1.1e-4* | -0.002*** | -0.002*** | -7.0e-4*** | -0.002*** | -7.0e-4*** | -0.002*** | -7.0e-4*** | |
| $(T-t) * x_3$ | 8.5e-6*** | | 2.1e-5*** | | 5.6e-5*** | | 2.3e-5*** | | 5.6e-5*** | | 2.3e-5*** | |
| $(T-t) * x_4$ | -8.5e-6*** | | -2.1e-5*** | | -5.6e-5*** | | -2.3e-5*** | | -2.3e-5*** | | -2.3e-5*** | |
| $(T_h-t) * x_1$ | | | -6.2e-4*** | | 0.062*** | | -6.2e-4*** | | 0.031*** | | 0.031*** | |
| $(T_h-t) * x_2$ | | | -0.016*** | | -0.016*** | | -0.003*** | | -0.003*** | | -0.003*** | |
| $(T_h-t) * x_3$ | | | 0.001*** | | 0.001*** | | 3.4e-4*** | | 3.4e-4*** | | 3.4e-4*** | |
| $(T_h-t) * x_4$ | | | -0.008*** | | -0.008*** | | -0.003*** | | -0.003*** | | -0.003*** | |
| $Mthdum$ | ✓ | ✓ | ns | ✓ | ns | ✓ | ns | ✓ | ✓ | ✓ | 0.100* | |
| αR^2 | 55.76 | 4.98 | 43.35 | 1.66 | 43.35 | 48.88 | 0.64 | 48.88 | 58.75 | 3.36 | 59.31 | |

Notes:
 1. In this table, we regress the cost-adjusted basis and the discounted risk premium on a non-linear function of inventories as discussed earlier. More specifically, we regress the following:

$$M1 : \psi_t, T = f(x_t)(T-t) + g(x_t) \max(T-T_h, 0).$$

$$M2 : RP_t, T = f(x_t)(T-t) + g(x_t) \max(T-T_h, 0).$$

$$M3 : \psi_t, T = RP_t, T + f(x_t)(T-t) + g(x_t) \max(T-T_h, 0).$$
 (13)
 (14)
 (15)

*, ** and *** denotes significance at the 10%, 5% and 1% level respectively; x_1, x_2, x_3, x_4 stand for $\frac{1}{norm.inv.}$, $\frac{1}{(norm.inv.)^2}$, $\frac{1}{(norm.inv.)^3}$ and $\max(\frac{1}{norm.inv.} - k_1)$. Only the significant terms of the regressions have been retained. k_1 is equal to 10 for all the contracts. For wheat, $T - T_h$ is equal to 30, 30, 90 and 180 for the May, July, September and December contract respectively; for corn, $T - T_h$ is equal to 210, 270, 330 and 60 for the May, July, September and December contract respectively; and for oats, $T - T_h$ is equal to 330, 30, 90 and 180 for the May, July, September and December contract respectively.
 2. The number of observations is not always the same for the regression of the cost-adjusted basis and the regression of the risk premium. This is due to missing observations for the risk premium after December 1935 (due to our sample) for contracts maturing in 1936 while the cost-adjusted basis can still be computed.
 3. We use OLS corrected for Newey-West standard errors for the regression to account for the autocorrelation induced by overlapping data.

Working curve (Wright and Williams, 1989). Consider a commodity with two qualities A and B that are close substitutes to each other. Consider also a two-period model with following prices: quality A and B trade at price p_t^A , p_t^B , p_T^A and p_T^B in the first and the second point in time respectively. Quality A is abundant with an inventory I^A and reflects full carrying charges:

$$E_t(p_T^A) \approx p_t^A \times (1 + r_{t,T}) + c_{t,T}, \quad (16)$$

with $r_{t,T}$ the interest charges and $c_{t,T}$ the storage cost. Quality B however is in stock-out right now, implying that its price at t is higher than the price for quality A .¹¹ Since stock-outs cannot last long, the expected price for T is relatively lower. Let us assume for convenience that:

$$E_t(p_T^B) \ll p_t^B \times (1 + r_{t,T}) + c_{t,T}, \quad (17)$$

and that the quality aggregated prices p_t^{agg} and p_T^{agg} have the following relationship as a result of the stock-out in quality B :

$$E_t(p_T^{agg}) \equiv \frac{E_t(p_T^A) + E_t(p_T^B)}{2} < p_t^{agg} \equiv \frac{p_t^A + p_t^B}{2} \times (1 + r_{t,T}) + c_{t,T}. \quad (18)$$

Equation (18) shows what happens when prices and inventory levels for quality A and B are averaged. If we had detailed information for each quality, we would know that quality A is not backwarddated and that quality B is subject to stock-out conditions. However, if the prices and the inventory levels of both qualities are averaged, the commodity seems to be backwarddated at positive inventory levels even though it is not.

We can now analyze whether the predicted effect of quality aggregation is empirically detectable. The Chicago Board of Trade microfilms contain weekly spot price information for each quality available at each point in time. Though the data are not perfect,¹² they are the best yet to test for the quality aggregation effect. We estimate the yield/inventory relation for each quality quoted using the method described in Section 1,¹³ and we count the number of successes, that is, the number of times the curve is significantly negative-sloped and convex. This is quite conservative in the sense that some rejected curves look flawless but

¹¹The price difference must be limited by the transformation cost of changing quality A into quality B since the two are close substitutes, but the sign of the price difference is quite predictable.

¹²The data are weekly averages of the highest and the lowest price of that week and are therefore not fully accurate. Moreover, we have no idea how liquid these spot prices were and whether they were deliverable on the futures contract.

¹³Qualities that were quoted only fragmentarily between 1885 and 1935 are not included.

are not significant—possibly because of the samples are now much smaller than before—while other curves are perfect except for a positively-sloped section at extremely low inventory levels (< 0.15 , normalized), plausibly reflecting underreported inventory levels. However, we prefer to be conservative in accepting a Working curve in order not to bias the results upward.

Panel A of Table 7 summarizes the results for wheat.¹⁴ There is a perfect Working curve for 69% of the tested price-normalized inventory relationships. That is, in only 31% of the tested relationships there is some flaw, however small, in the estimated curve. Obviously, then, the Working curve cannot be just the result of aggregation. Though the Working curve does not seem to be an artefact of data aggregation, we still find it interesting to test whether averaging the prices of the qualities that do not exhibit a Working curve will lead to the emergence of a Working curve for that average. We find that a Working curve does emerge. It remains unclear to what extent this result reflects the noise-reduction effect of averaging rather than its potential aggregation-bias effect. Still, from the wheat data we conclude that in aggregated data part of the evidence may reflect some quality bias.

Panel B shows the results for oats. For oats, there is a flawless Working curve in 64% of the cases analyzed. The July contract is the only maturity where the Working curve is not emerging at all. In this case, testing the relationship with the average of the prices does not lead to a Working curve either.

Panel C, lastly, presents the results for corn. A full 79% of the cases now exhibit a perfect Working curve. Interestingly, in this case the test of the relationship with the average of the prices leads to a Working curve in only one out of four cases. It must be added, in fairness, that the relationship for the price averages does become more Working curve-like. Thus, while we certainly cannot ascribe the Working effect for corn to pure aggregation bias, quality aggregation may still have contributed something.

To wrap up: although in the smaller samples that are free of aggregation bias the Working curve is not always unambiguously and perfectly present, it certainly exists for a clear majority of contract/quality combinations. We can also conclude that data aggregation may have led to a relationship that looks like the Working curve even if the underlying data do not, or not perfectly, exhibit this relationship; this may have been the reflection of noisy individual data, but we cannot rule out that part of it may be spurious.

¹⁴Regression results and graphs showing the relationship between prices and normalized inventories are available upon request.

Table 7: The Working curve from spot prices of various qualities from 1885 to 1935

| Quality | May | | Jul | | Sep | | Dec | |
|--------------------------------|-----|----|-----|----|-----|----|-----|----|
| | M1 | M2 | M1 | M2 | M1 | M2 | M1 | M2 |
| Wheat | | | | | | | | |
| No1 Dark Hard Winter | ✓ | - | ✓ | ✓ | ✓ | No | No | - |
| No1 Dark Northern Spring | ✓ | - | ✓ | ✓ | ✓ | No | No | - |
| No1 Hard Winter | ✓ | - | ✓ | ✓ | ✓ | No | No | No |
| No1 Northern Spring | ✓* | - | ✓ | ✓ | ✓ | ✓ | ✓ | No |
| No1 Red Winter | ✓ | - | ✓ | No | No | No | No | No |
| No2 Dark Hard Winter | ✓* | - | No | ✓ | No | No | No | - |
| No2 Dark Northern Spring | ✓ | - | No | ✓ | ✓ | No | ✓* | ✓* |
| No2 Hard Winter | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| No2 Northern Spring | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| No2 Red Winter | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | No |
| No2 Spring | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | - |
| No3 Hard | No | - | No | No | No | No | No | - |
| No3 Hard Winter | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | No |
| No3 Red Winter | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| No3 Spring | No | - | No | No | No | ✓ | No | ✓ |
| No3 Spring and Northern Spring | No | - | No | ✓* | ✓ | No | ✓ | ✓ |
| No4 Red Winter | ✓ | - | ✓ | ✓ | ✓ | ✓ | ✓* | ✓* |
| Quality Average | No | - | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Oats | | | | | | | | |
| No1 White | ✓ | - | - | No | ✓ | No | ✓ | No |
| No2 | ✓ | - | - | No | ✓ | ✓ | ✓ | No |
| No2 White | ✓ | - | - | ~ | ✓ | No | ✓ | ✓ |
| No3 White | ✓ | - | - | ~ | ✓ | No | ✓ | ✓ |
| No3 White Illinois | ✓ | - | - | No | ✓ | No | ✓ | No |
| No3 White Sample | ✓ | - | - | No | ✓ | ✓ | ✓ | ✓ |
| No4 White | ✓ | - | - | ~ | ✓ | No | ✓ | ✓ |
| Quality Average | - | - | - | No | - | No | - | ✓ |
| Corn | | | | | | | | |
| No2 | ✓* | ✓ | ✓* | - | No | - | ✓ | No |
| No2 Mixed | ✓ | ✓ | ✓ | - | ✓* | - | ✓ | ✓ |
| No2 Yellow | ✓* | ✓* | No | - | No | - | No | No |
| No3 | No | ✓* | No | - | No | - | No | No |
| No3 Illinois | ✓ | ✓ | No | - | ✓ | - | ✓ | ✓ |
| No3 Mixed | ✓ | ✓ | ✓ | - | ✓* | - | ✓ | ✓ |
| No3 Sample | ✓ | ✓ | No | - | ✓ | - | ✓ | ✓ |
| No3 White | ✓ | ✓ | ✓ | - | ✓* | - | ✓ | ✓ |
| No3 Yellow Illinois | ✓ | ✓ | No | - | ✓ | - | ✓ | ✓ |
| No3 Yellow Sample | ✓ | ✓ | No | - | ✓ | - | ✓ | ✓ |
| No4 | ✓ | ✓ | ✓ | - | ✓* | - | ✓ | ✓ |
| No4 Mixed | ✓ | ✓ | ✓ | - | ✓* | - | ✓ | ✓ |
| Quality Average | - | - | ✓* | - | No | - | No | No |

Notes:

1. In this table, we regress the cost-adjusted basis and the discounted risk premium on a non-linear function of inventories as discussed earlier. More specifically, we regress the following:

$$\psi_{t,T} = f(x_t)(T - t) + g(x_t) \max(T - T_h, 0) + \text{seasonals}.$$

2. We use OLS corrected for Newey-West standard errors for the regression to account for the autocorrelation induced by overlapping data.

3. We then fit a curve that describes the estimated relationship between the spread and normalized inventories and make a scatter graph of this curve and normalized inventory levels. Visual inspection of the result indicates whether the typical Working curve is present or not. The quality average row indicates the results for the average price of the qualities that did not yield a Working curve.

4. M1 indicates that $T_h \ni [t, T]$ and M2 indicates that $T_h \in [t, T]$; in other words, M1 and M2 are the curves both after and before the harvest h . ✓ indicates the presence of a Working curve, ✓* indicates the presence of a Working curve when the normalized inventory is larger than 0.1 (this means we eliminate potential stock-out effects), No indicates the absence of a Working curve and - indicates insufficient observations.

4 The Imbedded-option Hypothesis Cannot Replicate the Working Curve

In futures markets, the seller of a futures contract receives a number of options that are inherent to the contract: within well-defined bounds, the seller can choose the quality, location and time of the deliveries. Consider an agent with an inventory of one unit who has to determine whether to sell his inventory now, at t , or in the future, at T . Assuming that the agent is not better informed than the market as a whole, the expected risk-adjusted future spot price, $E_t^Q(\tilde{S}_T)$, equals the current futures price $F_{t,T}$. The agent will only hold inventory if the expected value of one unit at T exceeds the spot price now, S_t , adjusted for storage costs and the opportunity cost of not selling this unit at S_t and investing the proceeds. Therefore he holds inventory if

$$F_{t,T} \geq [S_t + PV(c_{t,T})](1 + r_{t,T}), \quad (19)$$

where, as before, $PV(c_{t,T})$ denotes the present value of paying the storage cost over many periods and $r_{t,T}$ is the risk-free interest rate. As we discussed above, the systematic violation of this bound in futures markets has been explained by the existence of a convenience yield, the existence of a risk premium, or mismeasurements due to transaction costs and data aggregation. To our knowledge, the impact of the imbedded options inherent in a futures contract on the above relationship has not been considered so far. However, research has been done on the valuation of these options. Early research studied the three options either separately or additively (Boyle, 1989; Chance and Hemler, 1993; Fleming and Harvey, 1994; Hemler, 1990; Pirrong, Kormendi and Meguire, 1994; and Silk, 1997). Because the three options clearly interact with one another, Hranaiova, Tomek and Jarrow (2005) consider the value, at the first day of the delivery month, of the joint timing and location option. They do not explicitly consider the quality option because it can be valued as a location option - the two options being qualitatively equal.¹⁵ They show that the joint option can be valued and that its value depends on the maturity month. They also prove that ignoring the interaction effect between timing and location can lead to misleading estimates of the imbedded options: the two options are not independent and thus cannot be priced separately and then added up. To analyze how these imbedded options affect Equation (19), we set up a simple model of inventory holding. The focus of this section is not to value options, but to analyze whether the delivery options

¹⁵Indeed, the quality option just adds choices for the holder of a futures contract. Now the option can be considered as the paired choice (quality, location) instead of location only. The option type remains the same and thus the option can be valued in the same way.

could account for backwardation and explain the Working curve. Therefore, we are interested in the determinants of the delivery options value. Before we can do that however, we need to have a look at how the delivery options add value to the option holder.

4.1 Setting the stage

We adopt a homogeneous-expectations setting with perfect and complete financial markets. Trading in spot and futures markets is continuous, but we assume that there are only two decision windows: t , the current period, and $T - \Delta$ to T , a future period corresponding to the delivery month of a traded futures contract with maturity date T . Consider the following problem. An inventory holder specializes in storing one commodity that comes in two qualities (A and B), each of which are deliverable on a futures contract. The inventory holder can choose between selling one unit of quality B, the par quality, spot at S_t^B or forward at $F_{t,T}$. Following the traditional approach, the discounted period- t cashflow that accrues to the inventory holder when he chooses to keep the par commodity B and sell it forward is equal to:

$$\frac{F_{t,T}}{1 + r_{t,T}} - PV(c_{t,T}), \quad (20)$$

where $PV(c_{t,T})$ denotes the present value of paying the storage cost over many periods (we assume that unit storage costs are independent of the quality stored). If he chooses to sell now, his income is S_t^B . As a result, a rational agent will hold inventory only if:

$$F_{t,T} \geq [S_t^B + PV(c_{t,T})] \times (1 + r_{t,T}), \quad (21)$$

i.e. when there is no backwardation. However, the traditional result given in Equation (21) underestimates the benefits that accrue to the inventory holder from selling the commodity forward. Indeed, the above discussion ignores the value of the imbedded delivery options *re* timing, location and quality. Assume the three options have a combined value of $O_{t,T|F_{t,T}}$ in period t for a futures contract initiated at time t_0 and maturing at T . Then the seller of the futures contract - such as our inventory holder - not only receives the futures price at maturity but also implicitly obtains the compound option with a positive expected cashflow. As a result, the supply-of-storage equation given in Equation (21) should be adjusted to include this option income, $O_{t,T|F_{t,T}}$. A rational agent holds inventory only if:

$$\frac{F_{t,T}}{1 + r_{t,T}} + O_{t,T|F_{t,T}} \geq S_t^B + PV(c_{t,T}). \quad (22)$$

Equation (22) shows that if this option income is ignored, it may look like prices are backwardated while in fact they reflect an additional income for the inventory holder that has been

ignored so far. The question still remains however what drives the value of these three options and whether they can explain the Working curve.

Consider a futures contract on a commodity deliverable for quality A and B at location 1 and 2. Assume that each quality is available at each location. The par price $F_{t_0,T}$ is quoted for quality B at location 1. To simplify notation, we drop the mention that the option value is for the above futures contract initiated at time t_0 . If quality A is delivered on the contract, a fixed premium p has to be paid to the seller;¹⁶ if location 2 is chosen, the price received by the seller is decreased by the fixed amount d . We look at the option cashflows recursively. At T , the maturity date, the seller has to deliver the commodity. He can choose between delivery of quality A or B at location 1 or 2. This yields the following cashflows: $\{F_{t_0,T} + p - S_T^{A,1}; F_{t_0,T} + p - d - S_T^{A,2}; F_{t_0,T} - S_T^{B,1}; F_{t_0,T} - d - S_T^{B,2}\}$. To maximize his profits, the seller chooses to deliver at the quality and location that minimizes the spot price at T adjusted for the premium or discount. Therefore he chooses $\min(S_T^{A,1} - p; S_T^{A,2} - p + d; S_T^{B,1}; S_T^{B,2} + d)$. If there was no option, the seller would receive $F_{t_0,T} - S_T^{B,1}$. Therefore, the option payoff at T , also the option value, is

$$O_{T,T} = P_{T,T} = \max \left\{ 0; S_T^{B,1} - (S_T^{A,1} - p); S_T^{B,1} - (S_T^{A,2} - p + d); S_T^{B,1} - (S_T^{B,2} + d) \right\}, \quad (23)$$

where $P_{i,T}$ is the payoff of an option with maturity T if it is exercised at time i . This is like a put option with the par price (quality B at location 1) as the strike price. At $T - 1$, the seller has a similar option, but he also has the option to wait with delivery for one more day. If he delivers at $T - 1$, the option payoff is

$$P_{T-1,T} = \max \left\{ 0; S_{T-1}^{B,1} - (S_{T-1}^{A,1} - p); S_{T-1}^{B,1} - (S_{T-1}^{A,2} - p + d); S_{T-1}^{B,1} - (S_{T-1}^{B,2} + d) \right\}. \quad (24)$$

So if he waits, he gains $PV(E_{T-1}[O_{T,T}])$, the present value of the expected value at $T - 1$ of the option at maturity. The seller will only exercise the option at $T - 1$ if $P_{T-1,T} \geq PV(E_{T-1}[O_{T,T}])$. Thus, the option value at $T - 1$ is equal to

$$O_{T-1,T} = \max \{ P_{T-1,T}, PV(E_{T-1}[O_{T,T}]) \}. \quad (25)$$

Similarly, on the first day of the delivery month, the day at which the option turns American, the option value is

$$O_{T-\Delta,T} = \max \{ P_{T-\Delta,T}, PV(E_{t-\Delta}[O_{T-\Delta+1,T}]) \}. \quad (26)$$

¹⁶Note that we assume that quality A is of higher grade than quality B

In a nutshell, the option imbedded in a futures contract is essentially a Bermuda-style option: it can only be exercised on certain specified days during the life of the option, here the delivery month. More specifically, the imbedded option can be considered as a Bermuda-style option to deliver the worst of n assets.

4.2 Determinants of the imbedded-option price

From the above discussion, we can infer the factors that affect the value of the delivery options during their life. The value of an option on the higher of three variables depends on time to maturity, the risk-free rate, the (near-constant) premium or discount for delivering a non-par quality, the convenience yield (which plays a similar role as the dividend in American-style stock options), and the variance-covariance matrix, both across qualities and over the delivery dates. The interest rate probably has a small impact, given how short-lived the options are. We control for time to maturity by always studying the conventionally estimated yield (including the hypothesized option value) at the same stage in its life, the beginning of the delivery month, denoted as $T - \Delta$. The remaining key variables are the variance-covariance matrix and the expected convenience yields, if any, at the beginning of the delivery month, which may trigger early exercise. If current inventory is correlated with the variances and covariances or with convenience yield at $T - \Delta$, the delivery options can create a spurious relation between the conventionally estimated convenience yield (which includes the option value) and current inventory.

Variance at $T - \Delta$ should be higher the lower the inventories are, because the low inventories limit the supply sides' elasticity, meaning that any demand fluctuation triggers a heftier price rise (see Appendix A). Since inventories are strongly seasonal, variance should be seasonal too; the question is whether past inventories allow us to improve upon the variance predicted by a pure-seasonal model. Thus, we estimate two competing models for each good i , without and with lagged inventories:

$$[r_t^i - E(r_t^i)]^2 = \alpha + \sum_{j=1}^{12} \beta_j D_{j(t)} + \varepsilon_t, \quad (27)$$

$$[r_t^i - E(r_t^i)]^2 = \alpha_n + \sum_{j=1}^{12} \beta_{n,j} D_{j(t)} + \gamma_n x_{t-n}^i + \varepsilon'_t, \quad (28)$$

where r_t^i is the daily return at t for commodity i and $D_{j(t)}$ is a dummy variable indicating whether the t -th observation pertains to month j . Similarly, the covariance between two prices of closely related goods at $T - \Delta$ should increase the lower the inventories are, because any

sharp price change spills over into the market of the substitute and triggers a similar price change. Again, this logic predicts a significant seasonal in the covariances, and the question is whether past inventories allow us to improve upon the covariance predicted by a pure-seasonal model. The equations now are

$$\left[r_t^{i,q} - E\left(r_t^{i,q}\right) \right] \left[r_t^{i,k} - E\left(r_t^{i,k}\right) \right] = \alpha' + \sum_{j=1}^{12} \beta'_j D_{j(t)} + \varepsilon_t, \quad (29)$$

$$\left[r_t^{i,q} - E\left(r_t^{i,q}\right) \right] \left[r_t^{i,k} - E\left(r_t^{i,k}\right) \right] = \alpha'_n + \sum_{j=1}^{12} \beta'_{n,j} D_{j(t)} + \gamma'_n x_{t-n}^i + \varepsilon'_t, \quad (30)$$

where $r_t^{i,q}$ and $r_t^{i,k}$ are the weekly returns at t for quality q and quality k of commodity i and $D_{j(t)}$ is again a dummy variable for month j

We verify the predictions regarding seasonality in the (co)variances by running the above regressions on commodity return volatilities using the data on quality spot prices quoted on the Chicago Board of Trade between 1885 and 1935 for wheat, corn and oats. We have daily spot data for the par quality (the quality traded at par on the futures contract) from which we can infer the variance seasonality; and we have weekly spot price data for various qualities available simultaneously on the spot market. Figure 4 displays the seasonals of the variance of the returns, the inventories and the of pairwise covariances. We see that the seasonality patterns are clearly related and conform to our expectation. In months where inventory is typically low, the variability of returns is typically high and vice versa. When inventory is low, the pairwise correlation of returns is also high and vice versa.

We still have to confirm that the variance and pairwise correlation are not related to past inventories. We estimate Equations (28) and (30), which model the variance and pairwise correlation of wheat, corn and oats on the basis of seasonal dummies and past inventories. We find that past inventory only has a limited effect on current variance and pairwise correlation: for wheat a few lags are significant; for corn, the number of significant effects is hardly above what one would expect to get by chance; while for oats past inventory is even not significant at all (see results in Table 8).

Lastly, we discuss whether the ‘dividend’ at $T - \Delta$, the moment the options become exercisable, can be proxied for by earlier inventory data, dated $T - \Delta - n$ and thus induce an option value that consistently explains the Working curve. If there is a convenience yield (the ‘dividend’) at the maturity month, then the question is actually irrelevant. First, if there is a convenience yield at maturity that impacts the option value, then the convenience yield hypothesis cannot be rejected for any maturity: either there is a convenience yield or there isn’t.

Figure 4: Seasonality patterns of return volatility, inventories and pairwise correlation

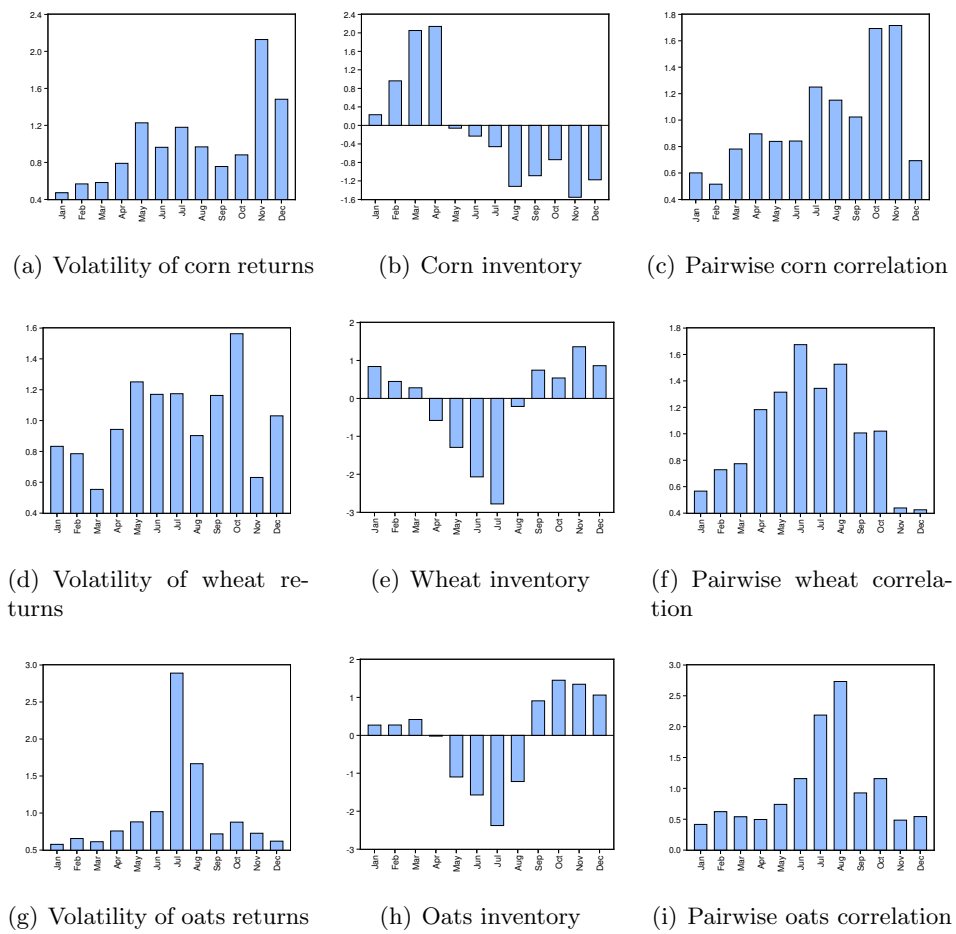


Table 8: Are variance and pairwise correlation affected by past inventories? Evidence for 1885-1935

| Variable | Wheat | | | | Corn | | | | Oats | | | |
|------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| | Var | | Corr | | Var | | Corr | | Var | | Corr | |
| | M_1 | M_2 | M_1 | M_2 | M_1 | M_2 | M_1 | M_2 | M_1 | M_2 | M_1 | M_2 |
| x_{t-4} | ns | | ** | *** | ns | | ns | | ns | | ns | |
| x_{t-8} | ns | | ns | | ns | | ns | | ns | | ns | |
| x_{t-12} | ns | | ns | *** | ns | | ns | | ns | | ns | |
| x_{t-16} | ns | | ns | | ns | * | ns | | ns | | ns | |
| x_{t-20} | ns | * | ns | *** | ns | | ns | ** | ns | | ns | |
| x_{t-24} | * | ** | ns | | ns | | ns | | ns | * | ns | |
| x_{t-28} | ns | | ns | *** | ns | | ns | | ns | | ns | |
| x_{t-32} | ns | | ns | | ns | | ns | | ns | | ns | |

Notes:

1. In this table, we regress the variance and the pairwise correlations linear function of normalized inventories, seasonal dummies and their own first lag.

2. We use a pooled regression for the pairwise correlations of wheat, corn and oats and a normal OLS regression with Newey-West correction for the variance. We report the results for the past inventory terms only so as not to clutter the table.

3. *, ** and *** denote significance at the 10%, 5% and 1% level respectively; ns means the coefficient is not significant at the 10% level; M_1 shows the regression results when all terms are included; M_2 gives the regression results after the insignificant inventory lags have been deleted.

Second, Carbonez, Nguyen and Sercu (2008) show theoretically that the convenience yield at t is impacted by future period-by-period convenience yields. As a result, convenience yield at t is related to convenience yield at the maturity month. However, again, this does not preclude the existence of a convenience yield.

To conclude this section, we can state that the realized value of the imbedded options seems to depend mainly on concurrent inventory levels for the maturity month. There may be some feedback from past inventories, but not consistently across maturities and commodities. That is, during the contract's life, the expected future option value is hardly related inventory levels at the time of setting the price. As a result, the imbedded options inherent in an agricultural futures contract cannot explain the Working curve: another mechanism must account for the bulk of that relation.

5 Conclusion

In this paper, we analyzed whether a Working curve is present in storable commodity markets in Chicago, more specifically for wheat, corn and oats and whether the risk premium hypothesis, the mismeasurement hypothesis and the imbedded-option hypothesis can account for the shape of the relationship between the cost-adjusted basis and the level of inventories. We show that neither of these hypotheses can explain much of the Working curve and that therefore the convenience yield hypothesis cannot be rejected.

First, we find that the risk premium may indirectly play a role as it is weakly related to inventory levels. Although the risk premium may be a small part of the cost-adjusted basis, it cannot explain the particular shape and the explanatory power of the relationship between the cost-adjusted basis and inventories.

Second, we show that mismeasurement may account for the Working curve in some situations. In particular, data aggregation occasionally leads to a Working curve even if the disaggregated data show no clear evidence of this relationship. However, the Working curve is also present when mismeasurements issues are avoided or accounted for. In sum, the evidence is mixed. We surely cannot discard the Working curve but neither can we deny an additional impact from data aggregation.

Third, we find that the imbedded options certainly explain part of the cost-adjusted basis, but that their value is not related with inventory levels at time t . The option value is likely to depend on inventory levels during the delivery month, but this is not what the Working curve is about.

As a result, we can conclude that the risk premium, data aggregation and the imbedded options can be discredited when it comes to explaining the Working curve found empirically. This leads to the conclusion that something else must explain the relationship between the cost-adjusted basis and inventory levels. The well-documented convenience yield is the most likely candidate. Though it is an unobservable variable, it has been well-documented theoretically and has shown its merits in empirical continuous finance work in the past.

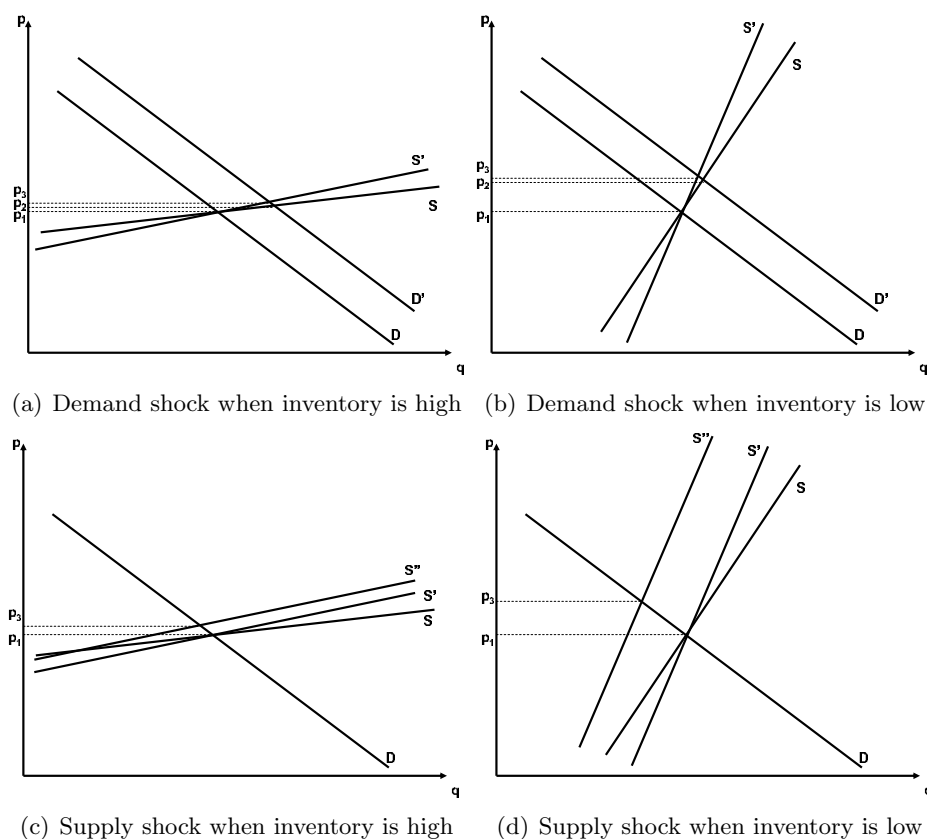
Appendix A: Inventories, Volatility and Correlations

Typically, asset prices in a worst-of- n -assets option are assumed to be driven by a set of geometric brownian motions with constant second moments. However, for commodity prices, we can do better than that.

From the literature (among others Schwartz (1997) and Sorensen (2005)), we know that agricultural prices and their volatilities are seasonal and mean-reverting. Harvest seasonality induces seasonality in commodity prices. However, inventory dampens this seasonality effect by acting as a transmission mechanism between periods of high supply and periods of low supply thus inducing mean-reversion. We look at the price reaction as the result of a demand and supply shock when inventories are low and high as an example of this seasonality effect.


Consider demand and supply shocks in period t . A positive demand shock has the fol-

Figure 5: The price reaction of a demand shock and a supply shock when inventories are low and when inventories are high



lowing effects. Tautologically, we observe a rightward shift in the demand curve. As a result, increased demand lowers inventory levels. This leads to more constrained supply conditions which reduces the elasticity of supply. A supply shortfall has similar effects. We now observe a leftward shift in the supply curve and supply conditions become more constrained, reducing the elasticity of supply. As a result, commodity price volatility will vary with the level of inventories. When inventory is high, the supply curve is very elastic. In that case, a demand or supply shock will only marginally influence spot returns. Moreover, since there is still ample inventory on hand, the elasticity of the supply curve will remain relatively high. Thus, as shown in Figure 5(a) and Figure 5(c) the price will not change too much. When inventory is low on the other hand, the supply curve is inelastic. As a result, a demand or supply shock will have a bigger impact on returns. Second, since storage is scarce, the elasticity of supply will be even more reduced which reinforces the impact on the return. Thus, as shown in Figure 5(b) and Figure 5(d) commodity prices will react more strongly. These effects will be even stronger if demand and supply is non-linear (see Figure 6).

Figure 6: The price reaction of a demand shock when demand and supply are non-linear



A similar pattern exists for the volatility of commodity prices. As Ng and Pirrong (1994) argue, prices are more volatile when supply is less elastic for a given distribution of demand shocks. As a result, commodity prices and commodity returns will be more volatile when inventory is low than when inventory is high. Since inventory levels are seasonal, returns and their volatility should be seasonal as well.

Up to this point, we looked at the commodity markets as if there were no substitutes. In reality, a number of different qualities of the same commodity are available in the commodity market. These are obviously close substitutes to each other but might require some additional investment (in terms of money or work hours) to be fully interchangeable. How will this affect the above discussion? More specifically, is the correlation between the returns of these close substitutes seasonal or not? Consider a representative producer who can substitute one quality for the other in his production process at a cost equal to C . This producer will only substitute if the price difference between the two qualities exceeds this level C . When inventories are high, the price reaction to a demand or supply shock is relatively low. As long as the price reaction leads to a price differential that is lower than C , the producer will still buy his preferred quality. As a result, when there is ample inventory, the prices of the two qualities can temporarily move independently from each other until their difference reaches C , after which demand shifts from the producer will force prices back inside the price band. When inventories are low however, a demand or supply shock can have a significant price impact for the preferred quality. If the price impact is larger than C , the producer will shift his demand to the other quality, thus inducing a related demand shock in the substitute. Thus, when inventories are scarce, demand shifts from the producer will occur much faster leading to more correlated prices than when inventories are high.

Appendix B: Risk premium tables for individual contracts

Table 9: Does the risk premium depend on inventory levels as the cost-adjusted basis? Evidence for wheat, corn and oats, sample:1885-1935

| $T - t$ | Pool | | May | | Jul | | Sep | | RP | | Dec | |
|----------------|------------|----------------|------------|----------------|------------|----------------|------------|----------------|------------|----------------|------------|----------------|
| | RP coef | ψ coef | RP coef | ψ coef | RP coef | ψ coef | RP coef | ψ coef | RP coef | ψ coef | RP coef | ψ coef |
| Panel A: Wheat | | | | | | | | | | | | |
| 0w | 0.44*** | 0.47*** | 1.01*** | 0.88*** | 0.62*** | 0.67*** | 0.14* | 0.19*** | 0.28 | 0.02 | 0.28 | 0.39 |
| 2w | 0.22 | 0.74*** | 1.15*** | 0.74*** | 0.64*** | 1.06*** | 0.08 | 0.58*** | -0.30 | 0.01 | 0.44 | 0.6 |
| 4w | -0.02 | 0.28*** | 0.50 | 0.64** | 0.62*** | 0.91*** | -0.04** | 0.20*** | -0.91 | 0.05 | 0.32 | 0.2 |
| 6w | 0.03 | 0.42** | -0.02 | 1.09*** | 0.63 | 2.03*** | 0.01 | 0.33*** | -0.85* | 0.06 | 0.61* | 0.1 |
| 8w | 0.29 | 1.02** | 0.27 | 0.87*** | 2.24** | 3.24*** | 0.26 | 0.85*** | -0.01 | 0.00 | 0.62** | 0.15 |
| 10w | 0.03 | 1.33** | -0.26 | 1.04*** | 2.9 | 4.48*** | 0.30 | 1.44*** | 0.24 | 0.01 | 0.45** | 0.20 |
| 12w | 0.63** | 0.4 | 0.30 | 1.20*** | 2.2 | 3.36*** | 0.31 | 2.42*** | 0.34 | 0.03 | 0.65** | 0.23 |
| 14w | 0.22 | 1.65*** | 0.43 | 1.63*** | 0.67 | 3.53*** | 0.29 | 3.02*** | 0.12 | 0.00 | 0.63** | 0.34 |
| 16w | 0.73 | 2.52*** | 0.32 | 2.08** | 1.39 | 3.00*** | 2.5 | 5.23*** | 0.47 | 0.02 | 1.34*** | 0.44 |
| 18w | 0.03 | 0.75*** | 2.64 | 1.16 | 0.5 | 3.37*** | 3.3 | 4.65*** | 0.00 | 0.00 | 0.54** | 0.28 |
| 20w | -0.02 | 3.45*** | -0.16 | 1.39* | 0.8 | 3.20*** | 2.8 | 4.56*** | -0.02 | 0.00 | 3.99*** | 0.33 |
| 22w | 0.31 | 3.13*** | 3.8 | 1.65*** | 1.5 | 2.62*** | 1.1 | 3.73*** | 0.24 | 0.02 | 4.11*** | 0.37 |
| 24w | 0.18 | 3.49*** | 3.7 | 1.44** | 1.1 | 2.70*** | 0.9 | 4.21*** | 0.34 | 0.00 | 5.14*** | 0.82 |
| 26w | 0.56 | 5.44*** | 3.5 | 1.30* | 1.1 | 2.90*** | 0.9 | 2.42** | 2.63 | 0.07 | 18.31*** | 0.8 |
| 28w | 0.29 | 4.08*** | 4.4 | 0.94 | 0.7 | 2.84** | 1.7 | 3.86** | 2.69 | 0.07 | 8.45*** | 0.8 |
| 30w | 0.07 | 2.05*** | 4.2 | 0.88* | 0.7 | 3.69*** | 3.0 | 2.11 | 2.02 | 0.02 | 3.40*** | 0.52 |
| 32w | 0.12 | 1.23*** | 4.8 | 0.75** | 2.1 | 4.61*** | 3.3 | 2.10 | 0.66 | 0.00 | 2.45*** | 0.61 |
| 34w | 0.15 | 0.74*** | 4.3 | 0.52** | 2.2 | 4.54*** | 2.1 | 3.78 | 10.00 | 0.04 | 2.24* | 0.48 |
| Panel B: Corn | | | | | | | | | | | | |
| 0w | 0.06 | 0.05 | -0.26 | 0.11 | 0.24* | 0.16 | 0.06 | 0.07 | 0.03 | 0.00 | 0.02 | 0.0 |
| 2w | 0.05 | 0.18*** | -0.15*** | 0.08 | 0.55 | 0.9 | 0.42* | 0.12 | -0.03 | 0.00 | 0.19*** | 0.24 |
| 4w | 0.23*** | 0.52*** | -0.23* | 0.17*** | 0.78*** | 2.2 | 0.62** | 0.34 | -0.06 | 0.00 | 0.34*** | 0.31 |
| 6w | 0.14 | 0.43** | 0.30*** | 0.01 | 0.39* | 0.4 | 1.24*** | 0.81** | -0.05* | 0.00 | 0.41*** | 0.52 |
| 8w | 0.31* | 0.61*** | 0.49 | 0.08** | 0.81** | 1.45*** | 2.7 | 1.26** | 0.22 | 0.02 | 0.72*** | 0.33 |
| 10w | 0.06 | 0.1 | 0.27*** | 0.04 | -0.74 | 0.3 | 1.12*** | 0.66* | -0.04 | 0.00 | 0.21 | 0.06 |
| 12w | 0.06 | 0.2 | 0.91*** | 0.07 | -0.47** | 0.2 | 0.49** | 0.85*** | 0.46 | 0.05 | 0.20 | 0.03 |
| 14w | 0.40** | 0.60*** | 4.8 | 1.03*** | -0.60 | 0.1 | 0.54** | 0.40 | 0.46 | 0.07 | 0.72** | 0.27 |
| 16w | 0.08 | 0.32*** | 4.4 | 0.75*** | 0.02 | 0.0 | 0.30** | 1.3 | -0.24 | 0.01 | 0.38 | 0.04 |
| 18w | 0.66*** | 0.70*** | 5.4 | 0.44*** | 0.20* | 0.1 | 0.18*** | 2.58*** | 0.90*** | 0.32 | 1.16*** | 0.55 |
| 20w | 0.35*** | 0.22** | 3.4 | 0.22** | 0.40 | 0.21 | 0.21*** | 4.08*** | 1.74** | 0.18 | 1.81*** | 0.26 |
| 22w | 0.20 | 0.3 | 0.14* | 0.19* | 0.2 | 0.48 | 0.19** | 1.82*** | 0.15 | 0.00 | 0.30 | 0.04 |
| 24w | 0.41 | 0.43*** | 4.6 | 0.30* | 1.06 | 0.2 | 0.72** | 2.18** | 1.62 | 0.09 | 1.49 | 0.18 |
| 26w | 0.58** | 0.80*** | 4.6 | 0.63*** | 1.28 | 0.3 | 0.32* | 1.87* | 4.25*** | 0.62 | 2.63*** | 0.51 |
| 28w | 0.23 | 0.32** | 4.0 | 0.77*** | 3.3 | 1.28 | 0.30 | 0.64** | 0.30 | 0.01 | 0.59* | 0.05 |
| 30w | 0.24 | 0.10 | 2.8 | 0.49 | 0.4 | 0.26** | 0.7 | 0.17 | 5.35 | 0.14 | 1.55 | 0.03 |
| 32w | 0.21 | 0.18* | 2.0 | 0.13 | 0.2 | 0.36* | 1.2 | 2.56* | -3.23 | 0.12 | 3.57 | 0.17 |
| 34w | 0.34* | 0.30* | 1.6 | 0.28 | 0.72*** | 3.1 | 0.33** | 5.01* | -1.62 | 0.02 | 5.74*** | 0.35 |
| Panel C: Oats | | | | | | | | | | | | |
| 0w | -0.00 | -0.00 | 0.01 | -0.07 | 0.33* | 0.16 | -0.00 | -0.00 | 0.19 | 0.06 | -0.02 | 0.0 |
| 2w | 0.02* | -0.00 | -0.01 | -0.06 | 0.51*** | 0.8 | 0.31** | 0.07 | 0.31** | 0.08 | 0.03 | 0.0 |
| 4w | 0.02 | 0.17*** | 0.22 | -0.10 | 2.60*** | 2.1 | 0.80* | 0.17* | -0.17 | 0.08 | -0.16 | 0.02 |
| 6w | -0.00 | 0.07*** | 0.48 | -0.17 | 0.56 | 0.1 | 0.39 | 0.38 | -0.27*** | 0.06 | 0.01 | 0.0 |
| 8w | -0.02 | 0.49*** | 3.4 | -0.13 | 0.86 | 0.1 | 0.75* | 0.6 | -0.13** | 0.01 | -0.05 | 0.1 |
| 10w | -0.08 | 0.23*** | 3.5 | -0.11 | 0.43 | 0.0 | 0.84** | 1.2 | -0.26 | 0.02 | 0.06** | 0.0 |
| 12w | -0.05 | 0.4 | 3.4 | 0.04 | 0.94 | 0.2 | 0.95*** | 0.87 | -0.06** | 0.05 | 0.00 | 0.0 |
| 14w | -0.00 | 0.4 | 4.1 | 0.24 | 0.49 | 0.0 | 0.95** | 1.0 | -0.06** | 0.02 | -0.00 | 0.0 |
| 16w | -0.01 | 0.08** | 4.1 | 0.90* | 0.02 | 0.0 | 1.14** | 0.1 | 0.06 | 0.01 | 0.08 | 0.15 |
| 18w | 0.01 | 0.03*** | 2.9 | 0.73* | -0.56 | 0.1 | 0.61** | 0.2 | 0.01** | 0.01 | 0.03*** | 0.13 |
| 20w | 0.01 | 0.51** | 3.1 | 0.40 | 0.14 | 0.0 | 0.72** | 0.96 | 0.00 | 0.00 | 1.85** | 0.10 |
| 22w | -0.02 | 0.38 | 3.3 | -0.01 | 0.46 | 0.0 | 1.07*** | 2.0 | -0.37** | 0.04 | 0.78 | 0.02 |
| 24w | 0.28 | 0.91*** | 3.9 | 1.23** | 0.1 | 0.55 | 0.88** | 1.4 | -0.59 | 0.01 | 1.92*** | 0.06 |
| 26w | 0.45 | 1.0 | 4.2 | 0.63*** | 1.26 | 0.5 | 0.79** | 1.7 | -1.04 | 0.03 | 49.48** | 0.58 |
| 28w | 0.07 | 0.17 | 4.2 | 0.06 | 1.19 | 0.5 | 0.84*** | 1.5 | -1.58 | 0.05 | 1.99** | 0.1 |
| 30w | 0.18 | 0.09 | 4.9 | 0.20* | 1.29 | 0.1 | 0.48 | 3.02* | 0.3 | 0.89 | 2.4 | 0.1 |
| 32w | 0.12 | 0.01 | 4.9 | -0.00 | 1.44* | 1.4 | 0.36 | -3.24 | na | na | 5.26 | 0.62 |
| 34w | 0.00 | 0.00 | 5.9 | 0.00 | 0.48 | 0.4 | -0.12 | 4.73*** | na | na | na | na |

Notes: see note to Table 4.

References

Benirschka, M. and Binkley, J.K. (1995) 'Optimal Storage and Marketing over Space and Time', *American Journal of the Agricultural Economics Association*, 77: 512-524.

Boyle, P.P. (1989) 'The Quality Option and Timing Option in Futures Contracts', *Journal of Finance*, 44(1): 101-113.

Brennan, D., Williams, J.C. and Wright, B.D. (1997) 'Convenience Yield without the Convenience: A Spatial-temporal Interpretation of Storage under Backwardation', *Economics Journal*, 107: 1009-1022.

Brennan, M.J. (1958) 'The Supply of Storage', *American Economics Review*, 47: 50-72.

Carbonez, K.A.E., Nguyen, Thi Tuong Van, and Sercu, P. (2008) 'Modeling the Net Convenience Premium: Exeunt Kaldor & Working, Enter Von Hayek?' Working paper.

Carter, C.A. and Giha, C.L.R. (2007) 'The Working Curve and Commodity Storage under Backwardation', *American Journal of Agricultural Economics*, 89(4): 864-872.

Chance, D. M. and Hemler, M.L. (1993) 'The Impact of Delivery Options on Futures Prices: A survey', *Journal of Futures Markets*, 13: 127-155.

Chavas, J-P., Despins, P.M. and Fortenbery, T.R. (2000) 'Inventory Dynamics under Transaction Costs', *American Journal of Agricultural Economics*, 82: 260-273.

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

Fama, E.F. and French, K.R. (1988) 'Business Cycles and the Behavior of Metal Prices', *Journal of Finance*, 43(5): 1075-1093.

Frechette, D.L. and Fackler, P.L. (1999) 'What Causes Commodity Price Backwardation?', *American Journal of Agricultural Economics*, 81: 761-771.

Gorton, G.B., Hayashi, F. and Rouwenhorst, G.K. (2007) 'The Fundamentals of Commodity Futures Returns', NBER Working Paper No. W13249.

Hemler, M. L. (1990) 'The Quality Delivery Option in Treasury Bond Futures Contracts', *Journal of Finance*, 45: 1565-1586.

Hicks, J.R. (1939) *Value and Capital: An Inquiry into some Fundamental Principles of Economic Theory*, Oxford: Clarendon Press, 331p.

Hranaiova, J., Jarrow, R.A. and Tomek, W.G. (2005) 'Estimating the Value of Delivery Options in Futures Contracts', *Journal of Financial Research*, 28(3): 363-383.

Kaldor, N. (1939) 'Speculation and Economic Stability', *Review of Economic Studies*, 1-27.

Keynes, J.M. (1930) *A Treatise on Money, Vol2: The Applied Theory of Money*, New York: Harcourt, 424p.

Liu, C.Y. and He, J.H. (1991) 'A Variance-Ratio Test of Random Walks in Foreign Exchange Rates', *Journal of Finance*, 46(2): 773-785.

Pirrong, S. C., Kormendi, R. and Meguire, P. (1994) 'Multiple Delivery Points, Pricing Dynamics, and Hedging Effectiveness in Futures Markets for Spatial Commodities', *Journal of Futures Markets*, 14: 545-573.

Silk, R. D. (1988) *Implicit delivery options in futures contracts and optimal exercise strategy*, Ph.D. dissertation, Stanford University.

Telser, L. (1958) 'Futures Trading and the Storage of Cotton and Wheat', *Journal of Political Economy*, 66: 233-255.

Working, H. (1948) 'Theory of the Inverse Carrying Charge in Futures Markets', *Journal of Farm Economics*, 30: 1-28.

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

Williams, J.C. and Wright, B.D. (1989) 'A Theory of Negative Prices for Storage', *The Journal of Futures Markets*, 9(1): 1-13.