



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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

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

Help ensure our sustainability.

Give to AgEcon Search

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

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

Realized Farm-Level Returns to Post-Harvest Grain Storage and Marketing



By Joseph P. Janzen

Joseph P. Janzen is an Assistant Professor in the Department of Agricultural and Consumer Economics, University of Illinois Urbana-Champaign.

Abstract

Commodity price variability is a major component of fluctuations in net farm income. Farm managers assume some of this price risk by choice when they store grain after harvest. This study estimates the realized returns from these post-harvest grain storage and marketing activities and shows that they are small on a risk-adjusted basis, particularly relative to the downside risk of negative returns. One explanation is that farm managers' use of post-harvest forward contracting is limited so they are subject to considerable flat price risk.

INTRODUCTION

Farm managers express consistent concern about grain market price risk. Surveys of farmer risk perceptions routinely rate commodity marketing as one of the most important risks faced in farm business management (Thompson, Bir, and Widmar, 2019; Atta and Micheels, 2020). In the aggregate, grain price variability is a major determinant of changes in farm profitability: elevated grain prices in 2007-2012 and 2020-2022 coincided with periods of record net farm income (USDA Economic Research Service, 2023). While these commodity price gyrations are certainly beyond the control of the farm manager, he or she

does have some choice about how much price risk the farm is exposed to with various market-based and government-backed risk management tools, including forward contracting, futures hedging, crop insurance, and commodity programs. This is especially true after harvest: post-harvest price risk is assumed voluntarily since the farm manager can transfer this risk to others by selling the crop at harvest or contracting for sale later in the marketing year. Farm managers may want to store to take advantage of seasonal price patterns—higher prices later in the marketing period compared to the harvest-time price—but doing so may be risky.

How much post-harvest price risk do farmers bear? This study assesses the post-harvest marketing performance of individual farms and quantifies the risk borne by farm managers who hold grain in storage after harvest. The analysis suggests post-harvest marketing and storage is a major component of the overall marketing strategy for corn and soybean farms in Illinois and throughout the US corn belt. I quantify the range of marketing outcomes experienced by individual farms that hold grain after harvest and compare it to realized prices received for grain sales made near to harvest. This assumes these farm-level distributions are informative about the range of returns to post-harvest marketing that farm managers may realize in the future.

Quantifying the realized range of potential post-harvest grain marketing outcomes is the major contribution of this study relative to prior research. Most previous analyses of farm marketing performance, including analyses of post-harvest marketing, use market-level data. In these studies, the returns to an assumed set of post-harvest marketing strategies are measured against a benchmark that is typically the cash price level observed during the harvest period. For instance, Edwards, et al. (2020) compare the net returns to unhedged and hedged post-harvest sales at varying storage horizons, with gains from these strategies assessed against the harvest-time cash price. Dietz, et al. (2009) conduct a similar analysis and show that different price baselines against which to compare the post-harvest price achieved by storage lead to significant differences in results. Because these studies rely on a limited set of market-level outcomes, they tend to be more prescriptive; it is unclear how they compare to actual

behavior, which is the result of a more complex set of marketing strategies. In reality, farm managers can choose to sell on any day (including all those days before harvest) for delivery on any day post-harvest. Other analyses of grain marketing performance do study farm-level marketing decisions (e.g., Anderson and Brorsen, 2005; Jacobs, Li, and Hayes, 2018) using grain purchaser records, but these data only record the interactions between farms and a single buyer.

This study differs by using farm-level data that covers the entirety of the farm manager's marketing decisions within a marketing year. Previous studies showed there are positive profits to post-harvest marketing that vary significantly across years. I show that realized marketing outcomes' returns vary more: there is substantial variation across farms even within the same year. This suggests both that farms employ more complex marketing strategies than accounted for in previous studies and that farmers may be assuming more price risk than previously thought.

This study proceeds as follows. First, I describe the typical seasonal pattern in local cash prices for corn and soybeans. I show prices typically rise about 20% between the harvest-time low and the post-harvest high. However, markets may deviate dramatically from this pattern in any given year, so holding commodity inventories after harvest does not guarantee the farm will receive higher prices. Next, using farm-level data, I show that farms generally hold significant proportions of their own production in inventory at calendar year end, which closely follows the harvest period. Only limited corn and soybean sales are realized in the near-to-harvest period between September 1 and December 31.

My main analysis calculates realized gross returns to grain storage for corn and soybeans on farms in Illinois as the difference between the price received by farm managers for deferred sales realized after January 1 and prices received for near-to-harvest sales. I show these returns are on average small and positive, which is roughly consistent with the average difference in cash prices between deferred and near-to-harvest periods of the marketing year. However, gross returns vary widely across farms in nearly all years. Observed variation across farms within each marketing year quantifies the risk to grain marketing. Using these volatility measures, I calculate risk-adjusted returns to post-harvest marketing, finding that the realized risk-adjusted returns are small and the downside risk is significant. As a group, farm managers are not choosing to capture seasonal price appreciation

through risk-minimizing marketing strategies such as forward contracting or storage hedges considered in earlier studies.

BACKGROUND

Seasonal Price Patterns

Seasonal price patterns provide incentives for farm managers and other decision-makers in the grain supply chain to store grain and make sales for delivery after harvest. Broadly, this pattern involves relatively low prices at harvest and relatively high prices later in the marketing period prior to the next harvest. Figure 1 illustrates the seasonal pattern for corn and soybeans using USDA Agricultural Marketing Service cash market price data for Central Illinois from the 2004-05 to 2019-20 marketing years. Note the marketing year for corn and soybeans runs from September 1 to August 31 of the following calendar year. The values shown in Figure 1 are deviations in a given week from the simple marketing year average price, which is the unweighted mean of daily price observations from that marketing year. These deviations remove differences in price levels across marketing years to focus on seasonal price changes within each year. The mean price series represents the typical difference between the price in a given week and the marketing year average price.

Figure 1 shows both corn and soybean prices hit seasonal lows at the beginning of October. The mean weekly deviation from the season average price (the thick line in Figure 1) is lowest at this point, coinciding with the typical harvest period in Central Illinois. Seasonal highs occur in the months of June and July, with prices tending to rise steadily between October and June. Based on the mean price pattern in Figure 1, both corn and soybean prices appreciate approximately 20% between the seasonal low and high. It is this seasonal price pattern that farm managers may seek to exploit by holding grain inventory after harvest. Note the seasonal high prices observed in June apply only to old-crop supplies and not to new-crop production that is typically planted before June and harvested in the fall; this seasonal pattern does not apply to pre-harvest marketing.

In any given year, there may be substantial differences between observed prices and the typical seasonal pattern. Figure 1 shows variation across years in prices at each week of the marketing year using the minimum and maximum deviations from the season average price observed between 2004/05 and 2019/20 and the standard deviation of these values across years for each week. The minimum values show that actual

price levels may be below the season average price at any point in the marketing year, so positive returns to storage are never guaranteed.

To evaluate marketing performance, note seasonal price patterns typically break even or equal the season average price near January 1 at the end of one calendar year and beginning of the next. (For soybeans, the mean seasonal price pattern reaches a point just below zero around January 1 and remains slightly below zero for several weeks.) Thus, farm managers would typically need to wait to sell grain until the new calendar year to receive prices above the marketing year average (assuming only cash sales are made). This calendar year end cut-off point is important for my farm-level analysis below.

Measuring Returns to Storage

Assessing profitability of farm commodity storage relies on some comparison of nearby and deferred prices. The nearby price represents in part the opportunity cost of storage, i.e., the value the farm would have received had it not stored. A common benchmark for the nearby price in many storage analyses is the cash price at harvest time. The deferred price represents the value the farm receives or expects to receive when the commodity is removed from inventory. The deferred price is typically the cash price later in the marketing year or the forward contract price offered at harvest for delivery later in the marketing year. In percentage terms, the gross return to storage is therefore

$$\text{Gross Return} = \frac{P_D - P_N}{P_N} \quad (1)$$

where P_D is the deferred price and P_N is the nearby price. This price comparison is grounded in the economic theory of commodity storage (Williams and Wright, 1991), which explains how a theoretical commodity-storing firm evaluates current and expected future prices. It is a gross return because it does not account for physical storage costs such as the handling, maintenance, and deterioration of grain inventories. It also ignores the time value of money associated with the foregone revenue from selling at the nearby price.

Accounting for storage costs at the farm level is more complex than in market-level analyses such as Edwards, et al. (2020). Market-level analysis typically assumes a physical storage cost that is a single fixed rate per month. Physical storage costs are likely to vary across farms and to vary with length of the storage

period in ways not encapsulated by a single per-month rate. In the same way, firms may have different opportunity costs of storage that depart from the time value of money given by benchmark interest rates.

For example, recent research suggests grain storage decisions may be a function of working capital and the farm's financial position, not just market-level interest rates (Janzen, Swearingen, and Yu, 2023). Given these complications, I consider gross returns only.

DATA

To measure post-harvest marketing performance at the farm level, I use data on corn and soybean production, sales, and inventories from Illinois Farm Business Farm Management (FBFM). Illinois FBFM is a cooperative association of more than 5,000 farmer cooperators who work with association field staff to collect financial and agronomic data for tax filing, financial statement preparation, and business benchmarking. FBFM partners with the Department of Agricultural and Consumer Economics at the University of Illinois, Urbana-Champaign, and its *farmdoc* extension project team to make data available for use in extension and research activities.

Illinois FBFM cooperators cover all regions of the state and represent approximately 25% of Illinois farmland acreage. FBFM data are used to develop University of Illinois crop budgets, which are based on audit-quality financial records from more than 1,000 farms each year. Note that an FBFM farm may include multiple farmer cooperators whose farm operations are joint. Since FBFM is a voluntary association, its records by design do not constitute a statistically representative sample of Illinois farms. However, recent comparisons of summary statistics and demographic measures for FBFM and those from the USDA's Agricultural Resource Management Survey (ARMS) indicate good representation of commercial-scale Illinois crop farms in the FBFM data (Kuethe et al., 2014).

Observing farms for multiple years (and thus across varied market and agronomic conditions) is one major advantage of the FBFM data. For this analysis, the data include records for all member farms in the period 2004 to 2020. I compile a sample of farm financial records certified as useable for research and benchmarking purposes by FBFM field staff, including only grain farms that grew corn or soybeans each year and excluding farms with zero operated acres and zero tillable acres. Most farms record production of both corn and soybeans each year. However, since farms participate in FBFM voluntarily, I do not observe all farms or commodities every year. For this analysis,

the data are an unbalanced panel of 31,111 farm-commodity-year observations from 17 calendar years. Each observation is specific to a farm, commodity (corn or soybeans), and calendar year. While farms almost always report both corn and soybean production and sales in a given year, farms generally are not observed in the data for all 17 years; the mean length of time a farm remains in the data is just less than six years.

I quantify the importance of post-harvest storage and marketing to these Illinois grain farms by extracting relevant quantities from farm financial statements. The balance sheet shows the level of key state variables, principally the quantity of inventories, at the end of each calendar year, and the income statement includes measures of important commodity flows during the calendar year. The basic accounting identity that describes inventory quantity dynamics is

$$Inventory_{t-1} + Production_t - Sales_t = Inventory_t \quad (2)$$

That is, inventory for the current year (t) must equal inventory for the previous year ($t - 1$) plus current year production less current year sales. I observe both the quantity in bushels and value in nominal dollars for both inventory and sales, then infer an implicit average price by dividing value by quantity. For example, I can calculate the farm-specific price received for all sales made within a calendar year by dividing the value of sales recorded on the farm income statement by the quantity of sales for a given commodity.

Observing calendar year sales quantity and value alone would be insufficient to evaluate post-harvest marketing performance because the inferred average sales price includes sales of both old-crop inventory carried into the calendar year ($Inventory_{t-1}$) and new-crop production ($Production_t$). However, Illinois FBFM records the quantity and value of what it calls old-crop and new-crop sales. New-crop sales are sales of current calendar year production realized prior to the end of the calendar year; I call these near-to-harvest sales and denote them as $Sales_t^N$. Near-to-harvest sales are realized in the sense that delivery is made and revenues are received before January 1. Commodities held in on-farm storage and unsold, those delivered into commercial storage where ownership is retained, and those held in any location but forward contracted for delivery and transfer of ownership on or after January 1 are old-crop sales for the next calendar year. I refer to these old-crop sales as deferred sales and denote them as $Sales_{t+1}^D$ to make clear these sales are realized in calendar year and accounting period ($t + 1$).

The quantities $Sales_t^N$ and $Production_t$ allow us to assess the importance of near-to-harvest sales to each farm. Figure 2 describes the distribution across farms of near-to-harvest sales realized prior to January 1 as a proportion of calendar year crop production by crop for FBFM farms in the period 2004 to 2019, i.e., $Sales_t^N / Production_t$. Farms with a zero share of near-to-harvest sales have realized no sales of new-crop production before January 1. These farms may have made forward sales of current production, but such sales are not yet realized prior to the new calendar year; new-crop production remains in inventory. Farms with 100% near-to-harvest sales have sold their entire calendar year production by January 1 and hold no commodity inventories as of year-end.

Figure 2 shows that although I observe farms at all points in the distribution, small shares of new-crop production sold before January 1 are much more common. Most notable is the share of farms that have realized zero or near-zero sales of crops produced in a given calendar year. For both corn and soybeans, more than 40% of farms have no sales of near-to-harvest sales. This holds across all years, and it is also true in specific years. While the specific share of farms with zero near-to-harvest sales fluctuates from year to year, it typically ranges between 30% and 50%.

The large proportion of farms that have little or no grain sold by January 1 of a given marketing year suggests these farms may face substantial price risk on the inventories they hold into the new calendar year. As noted above, these farms do have access to a wide array of price risk management tools. If these farms are proactively using these tools, then farms face less price risk, and marketing outcomes may not vary across farms or vary with post-January 1 market-level price changes. I use data on realized marketing performance on Illinois FBFM farms to assess these conditions.

RESULTS

To describe the post-harvest marketing performance of farms in my data, I estimate annual gross returns to storage for both corn and soybean farms in the Illinois FBFM. I use the panel structure of the data to calculate the prices received by farms for both near-to-harvest and deferred sales in each year. The average price received for each type of sales i in dollars per bushel, P_t^i , is the value of sales in dollars divided by the sales quantity in bushels. Near-to-harvest sales realized prior to January 1 represent the opportunity cost incurred by holding the commodity in storage and realizing sales later in the marketing year. Gross returns from deferred

sales relative to this near-to-harvest benchmark assume sales made after January 1 could have been sold at the price realized for the farm's earlier near-to-harvest sales. The near-to-harvest price benchmark is the amalgam of all pre-harvest marketing actions, which include the harvest cash sales used as a benchmark in previous studies as well as forward contracts delivered at harvest and other pre-harvest marketing strategies.

The gross return to storage is a summary measure of the profitability of post-harvest marketing that requires data from separate calendar years, i.e., the near-to-harvest price at year t and the deferred price from year $t + 1$. I calculate the gross percentage return for the marketing year that spans calendar years t and $t + 1$ as

$$\text{Gross Return} = \frac{P_{t+1}^D - P_t^N}{P_t^N} \quad (3)$$

This calculation limits the number of observations available for two reasons. First, many farms in my sample have no near-to-harvest sales ($\text{Sales}_t^N = 0$), so I cannot calculate P_t^N for these farms. As shown in Figure 2, this is more than 40% of the sample. Second, I can only calculate gross returns for farms for which I have data in consecutive calendar years. I therefore lose at least one observation per farm and commodity, including all those near-to-harvest sales observed in calendar year 2020. These limitations reduce sample size to 11,874 farm-commodity-year observations.

The gross return measure is particularly informative because it represents an individual-adjusted measure of marketing performance. It is specific to the farm's location and the set of local markets to which it can deliver, which do not change substantially between the pre-January 1 and post-January 1 periods. It is also specific to the quality of grain produced by the farm in that calendar period. In general, this comparison adjusts for many farm-specific factors that affect marketing performance and do not vary over time. These include farm manager ability, education, and risk preferences as well as other relevant aspects of the farm's business operations and financial capacity.

Variation Across Farms in Returns to Storage

I plot the distribution of gross returns from commodity storage and marketing for each commodity and marketing year in my sample period in Figure 3. Note these distributions weight each farm-level observation

equally; outcomes for large farms (which market more bushels) are treated as equally likely as outcomes for small farms. These distributions also do not account for the proportion of near-to-harvest and deferred sales on each farm. Gross returns only represent the raw price difference between realized near-to-harvest and deferred sales. Extreme values are also replaced (winsorized) at the top and bottom 0.5% of the entire distribution to reduce the impact of outliers.

Figure 3 shows that gross returns to commodity storage vary widely across farms producing corn and soybeans in Illinois. The range across all years runs from roughly -40% to +50%. For individual marketing years, the range of returns for the bulk of the distribution is at least 10 percentage points, but it is often much greater, and in some extreme cases, significant numbers of farms are receiving returns to storage that are 20 to 30 percentage points below the top performing farms. Both crops experience similar levels of cross-farm variation in returns.

Negative gross returns to storage are surprisingly common. A gross return of zero indicates the price received for near-to-harvest and deferred sales is equal, so there was no realized benefit to holding grain in storage. The marker below each distribution in Figure 3 indicates the median value in that year. Median returns are below zero in 6 of 16 marketing years for both corn and soybeans. Substantial portions of the mass of the distribution of returns are below zero every year, even in years like 2006/07 and 2007/08 when deferred marketing was exceptionally profitable. The common presence of negative returns suggests that farm managers realize more downside risk from storage than one might think, given the typical seasonal price pattern observed in Figure 1.

Marketing years with strongly positive returns tend to be those where cash prices rose a lot after harvest, such as 2006/07, 2007/08, and 2010/11. Farm managers realize that much of this upside price movement is suggestive (but not definitive) evidence that significant portions of the stored corn and soybean crops are uncontracted and/or unpriced until late in the marketing year. The converse, that farm managers realize negative returns when prices fall after harvest, is less clear. Years with negative returns such as 2008/09, 2014/15, and 2019/20 featured periods of modest price declines below the marketing year average during the deferred January 1 to August 31 marketing period. A full assessment of the relationship between market-level price outcomes and farm-level marketing outcomes is beyond the scope of this analysis and left for future study.

Estimating Aggregate Risk-Adjusted Returns

To summarize findings on the realized returns to commodity storage for Illinois farms, summary statistics for the gross returns data visualized in Figure 3 are presented in Table 1. First, I calculate the mean and median returns across all farms and years for each crop. I find the average gross return is 6.5% for corn and 5.0% for soybeans, with the median gross return 3.4% for corn and 2.6% for soybeans. Mean values above the median indicate a slight right skew in the distribution of returns across all years driven in part by those high return years where price appreciation led nearly all farms to experience positive gross returns.

Mean gross returns are remarkably similar to the seasonal market-level price changes found in Figure 1. Recall the low-to-high average seasonal price increase is roughly 20%, but it is unreasonable to expect farms to time their sales on these exact dates. The difference between the average weekly price between January 1 and August 31 (the deferred period of the marketing year) and the average weekly price between September 1 and December 31 (the near-to-harvest sales period) is 7.2% for corn and 7.0% for soybeans. This level of price appreciation between the two periods of the marketing year is similar in magnitude to the mean gross returns to post-harvest storage and marketing of 6.5% and 5.0%, respectively, suggesting that cash price variability and the timing of cash market sales may drive much of the variation in farm marketing performance.

To assess the variability of storage returns, I calculate two standard deviations in Table 1. The first unconditional standard deviation includes variation across all farms and years for each crop. The second standard deviation assesses variability across farms within years by subtracting the year-specific mean return from each gross return observation prior to calculating the standard deviation. The within-year standard deviation is slightly smaller since it does not include variability in returns across years; the second measure is the preferred estimate of the realized risk borne by farm managers in post-harvest grain marketing. Supposing a given farm manager is not predisposed to overperform or underperform his or her peers, he or she can expect the distribution of gross return outcomes in any given year to reflect the expected probability of his or her own returns. Then the within-year standard deviation accurately describes the risk he or she should expect to face.

I find the volatility or standard deviation of within-year gross returns is 14.9% for corn and 12.9% for soybeans. To place this risk measure in context, I employ a standard measure of risk-adjusted return, i.e., the Sharpe Ratio, a unitless measure of reward relative to variability. The Sharpe Ratio is calculated as the expected excess return relative to a risk-free asset divided by the standard deviation of the excess return (Sharpe, 1994). In this case, gross returns to storage are a form of excess return since the price of near-to-harvest sales represents foregone revenue from potential sales that would not be subject to post-harvest price variability (and thus could be considered “risk free” relative to the option to retain ownership and market grain after harvest).

In Table 1, I calculate the Sharpe Ratio as the mean gross return divided by the within-year standard deviation. I find Sharpe Ratios of approximately 0.43 for corn and 0.38 for soybeans, suggesting both crops exhibit similar returns to storage on a risk-adjusted basis. These levels would be considered low in the context of most investment/portfolio analysis. They are similar to Sharpe Ratios calculated for farm-level returns from all farm operations, not just storage (Langemeier and Yeager, 2021). More generally, my results suggest the risk-adjusted returns for post-harvest grain marketing are not large. Grain storage returns are unlikely to feature returns much larger than other aspects of the farm operation. These risk-adjusted returns are also likely smaller than the returns from readily available off-farm investments in public equity and bond markets. However, this does not necessarily imply grain storage is a “bad” investment as there may be other returns to holding commodity inventories. I discuss this possibility in the conclusion.

One limitation of the Sharpe Ratio is that it views upside and downside risk equally. In the context of post-harvest grain marketing, upside risk may be viewed as a benefit of holding unpriced inventory rather than selling. Farms may instead assess risk-adjusted returns relative only to downside risk or the probability that returns fall below some minimum acceptable level. While this level is unobservable, I consider gross returns below zero as exhibiting considerable downside for the farm. There are a significant number of instances of negative gross returns to storage, i.e., 39.6% of farm-year observations for corn and 38.4% of farm-year observations for soybeans.

To calculate a measure of returns adjusted for downside risk only, I calculate the Sortino Ratio, which

replaces the standard deviation term in the Sharpe Ratio with the target semi-deviation, i.e., the standard deviation of excess returns below a target (Sortino and Price, 1994). In this case, the target is positive gross returns. Table 1 shows that the Sortino Ratio in my sample period is 0.73 for corn and 0.64 for soybeans. While there is no objective threshold below which the Sortino Ratio is too low, these levels are concerning as they are well below the levels observed in Langemeier and Yeager (2021) for farm-level returns from all operations. The Sortino Ratio suggests the downside risk from post-harvest grain marketing is economically significant.

CONCLUSIONS

This study uses farm-level data to estimate the price risk assumed by farm managers who store grain after harvest. Post-harvest grain marketing is a strategy designed to profit from seasonal price appreciation typical in many agricultural commodity markets. I show that many farms employ this strategy, but near-to-harvest sales are typically a small share of production, and many farms realize no sales of new-crop corn or soybeans before January 1 following a fall harvest. I show the returns to post-harvest marketing are on average small and positive, approximately equivalent to the average difference in cash prices between deferred and near-to-harvest periods of the marketing year. However, farm-specific gross returns differ dramatically. I use observed variation across farms within each marketing year to quantify the risk to grain marketing and calculate risk-adjusted measures of the returns to post-harvest marketing. I find realized risk-adjusted returns are small and the downside risk is significant.

There are at least two caveats to the analysis of gross returns to post-harvest storage and marketing presented above. First, my gross return measure does not account for all benefits and costs of commodity storage incurred by deferred sales. Waiting to realize sales until after calendar year end entails additional

physical and opportunity costs of storage. If these are the only excluded benefits or costs to post-harvest marketing, then realized farm-level returns from storage are even poorer than indicated here. However, there may be additional unobserved benefits to storage beyond the price improvement realized on deferred sales. First, storage may be used to facilitate other aspects of farm operations that may provide significant benefits to farm managers. For example, on-farm storage may be used to speed harvest progress when local grain elevators or processing plants experience harvest-time congestion. Second, deferring grain sales until a new calendar year may provide income tax benefits to farm businesses that are difficult to quantify. By smoothing revenue across tax periods, deferred grain sales can reduce income tax liabilities for farms (Davenport, Boehlje, and Martin, 1982; McNew and Gardner, 1999). A second caveat is that this analysis ignores quantities when calculating the gross returns to storage. Farm managers who defer a large portion of sales to a subsequent calendar year assume more risk than those who realize large sales near to harvest. If farms that have a small share of near-to-harvest sales (those to the left of the distributions in Figure 2) are systematically able to realize better prices on deferred sales, then my results understate the risk-adjusted returns to post-harvest marketing. However, I have no evidence to indicate this is the case.

This study suggests that farm managers should carefully weigh the risks of deferring grain sales until later in the marketing year, especially unhedged sales. Although farmers do realize profits in the aggregate from selling later, the wide variety of outcomes from deferred sales shows that the downside risk of losing money on stored grain is substantial. Farmers can manage this risk and secure gains from deferred sales through forward contracting. Seasonal price appreciation is real but cannot be realized with certainty unless farmers use forward sales to capture those gains.

REFERENCES

- Anderson, K.B., and B.W. Brorsen. 2005. "Marketing Performance of Oklahoma Farmers." *American Journal of Agricultural Economics* 87 (5): 1265–70.
- Atta, C., and E.T. Micheels. 2020. "Identifying Risk in Production Agriculture: An Application of Best-Worst Scaling." *International Food and Agribusiness Management Review* 23 (2): 283–300. <https://doi.org/10.22434/IFAMR2019.0133>.
- Davenport, C., M.D. Boehlje, and D.B.H. Martin. 1982. "The Effects of Tax Policy on American Agriculture." Agricultural Economic Report No. 480. USDA Economic Research Service. <https://doi.org/10.22004/ag.econ.307912>.
- Dietz, S.N., et al. 2009. "The Marketing Performance of Illinois and Kansas Wheat Farmers." *Journal of Agricultural and Applied Economics* 41 (1): 177–91. <https://doi.org/10.1017/S1074070800002625>.
- Edwards, A.J., et al. 2020. "Indiana Corn and Soybean Returns to Storage." *Journal of ASFMRA* : 30–40.
- Jacobs, K.L., Z. Li, and D.J. Hayes. 2018. "Reference-Dependent Hedging: Theory and Evidence from Iowa Corn Producers." *American Journal of Agricultural Economics* 100 (5): 1450–68. <https://doi.org/10.1093/ajae/aay035>.
- Janzen, J.P., B. Swearingen, and J. Yu. 2023. "Buying Time: The Effect of Market Facilitation Program Payments on the Supply of Grain Storage." *Journal of the Agricultural and Applied Economics Association* 2 (3): 370–385. <https://doi.org/10.1002/jaa2.67>.
- Kueth, T.H., et al. 2014. "A Comparison of Data Collected through Farm Management Associations and the Agricultural Resource Management Survey." *Agricultural Finance Review* 74 (4): 492–500. <https://doi.org/10.1108/AFR-09-2014-0023>.
- Langemeier, M., and E. Yeager. 2021. "Examining Differences in Risk Adjusted Returns Among Farms." *Journal of ASFMRA* : 66–70.
- McNew, K., and B. Gardner. 1999. "Income Taxes and Price Variability in Storable Commodity Markets." *American Journal of Agricultural Economics* 8 (3): 544–52. <https://doi.org/10.2307/1244014>.
- Sharpe, W.F. 1994. "The Sharpe Ratio." *Journal of Portfolio Management* 21 (1): 49–58.
- Sortino, F.A., and L.N. Price. 1994. "Performance Measurement in a Downside Risk Framework." *The Journal of Investing* 3 (3): 59–64. <https://doi.org/10.3905/joi.3.3.59>.
- Thompson, N.M., C. Bir, and N.J. Olynk Widmar. 2019. "Farmer Perceptions of Risk in 2017." *Agribusiness* 35 (2): 182–99. <https://doi.org/10.1002/agr.21566>.
- USDA Economic Research Service. 2023. "Farm Income and Wealth Statistics." <https://www.ers.usda.gov/data-products/farm-income-and-wealth-statistics/>.
- Williams, J.C., and B.D. Wright. 1991. *Storage and Commodity Markets*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511571855>.

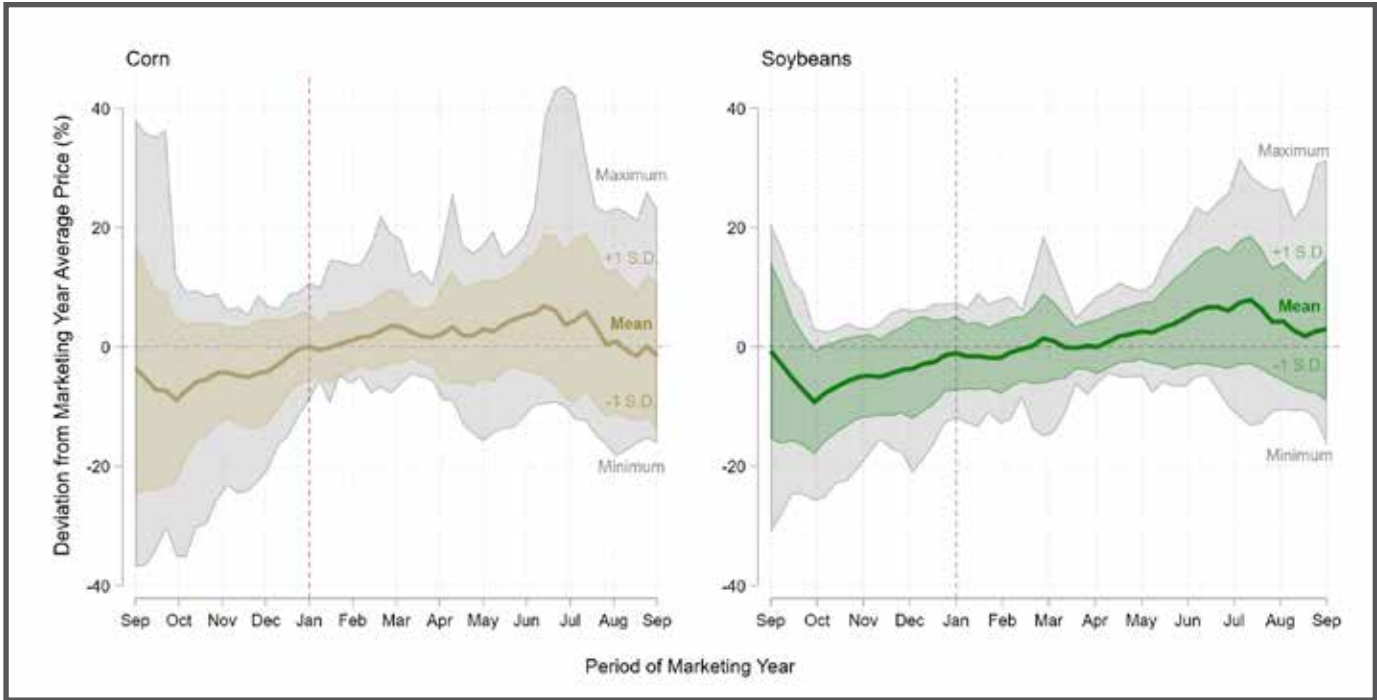


Figure 1. Typical and extreme seasonal variation in corn and soybean prices in Central Illinois, 2004–2020 (Mean (thick) lines indicate the average across all years of the weekly deviation from the marketing year average price, shaded areas indicate the range given by one standard deviation above and below the mean value (colored) and the maximum and minimum deviations observed (gray))

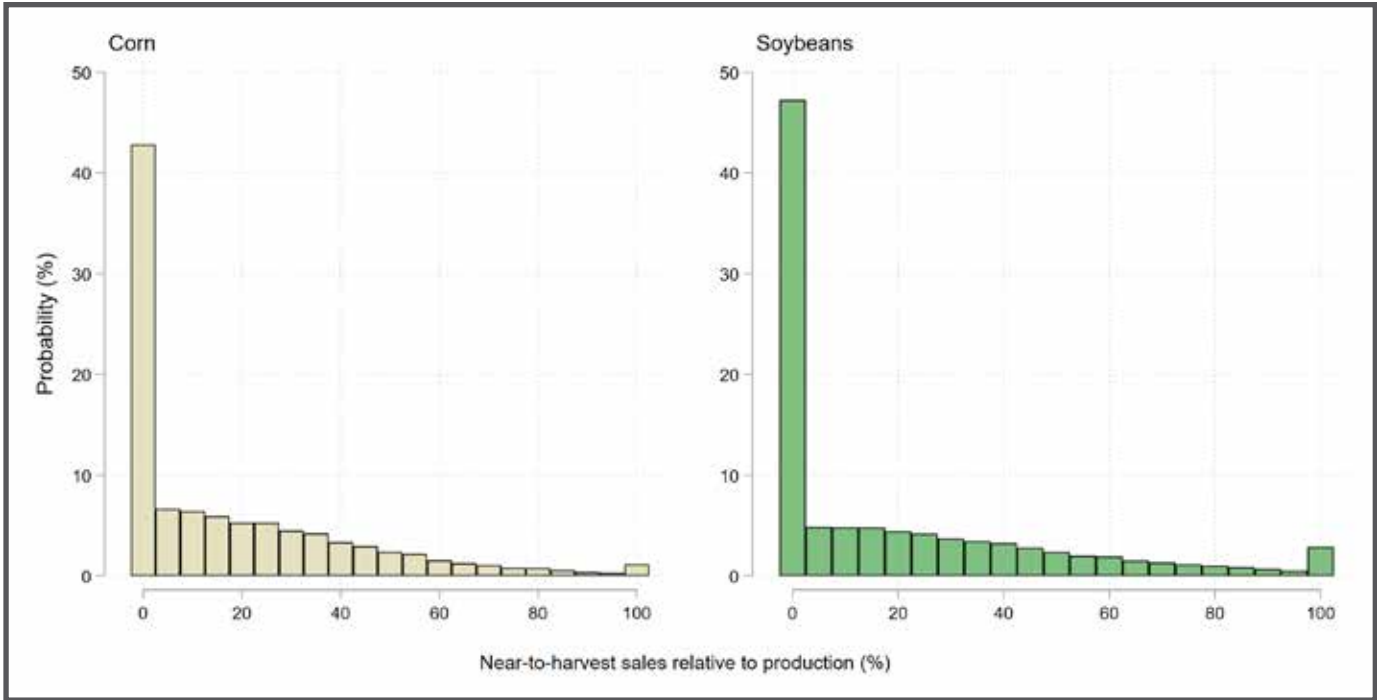


Figure 2. Distribution across farms and years of the proportion of corn and soybean sales made near to harvest




Figure 3. Distribution of gross returns from deferred (post-January 1) crop sales by commodity and marketing year, 2004/05 to 2019/20 (Markers below each distribution function indicate the median gross return in that marketing year)

Table 1. Summary Statistics for Gross Returns to Post-harvest Storage and Marketing by Crop*		
Gross Return	Corn	Soybeans
Mean	6.5%	5.0%
Median	3.4%	2.6%
Standard Deviation (Unconditional)	19.2%	15.1%
Standard Deviation (Within-Year)	14.9%	12.9%
Prob()	39.6%	38.4%
Target Semi-Deviation	8.8%	7.8%
Sharpe Ratio	0.43	0.38
Sortino Ratio	0.73	0.64

*Summary statistics are calculated from 11,874 farm-commodity-year observations for 17 marketing years from 2004–05 to 2019–20