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.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Potential Benefits of Water-Use Efficiency Technologies in Southeastern Wyoming

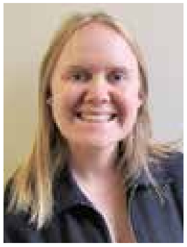
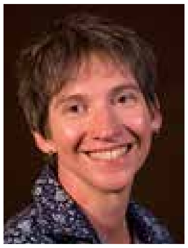
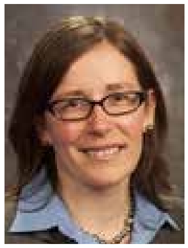
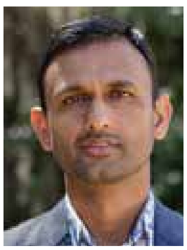

By Alicia Mattke, Kristiana Hansen, Dannele Peck, Vivek Sharma, Scott Miller, and Christopher Bastian

Alicia Mattke was a Graduate Research Assistant on this project. Kristiana Hansen is an Associate Professor and Extension Water Resource Economist, Department of Agricultural and Applied Economics, University of Wyoming. Dannele Peck is a Director and Agricultural Economist, USDA Northern Plains Climate Hub. Vivek Sharma is an Assistant Professor and Precision Water Management Specialist, Department of Agricultural and Biological Engineering, University of Florida. Scott Miller is a Professor, Department of Ecosystem Science and Management, University of Wyoming. Christopher Bastian is a Professor, Department of Agricultural and Applied Economics, University of Wyoming

Acknowledgement

This research was funded by the University of Wyoming Water Research Program. Opinions expressed in this manuscript are those of the authors and not the granting agency. The authors thank Tim Foster (University of Manchester), Jim Pike (NRCS retired District Conservationist), Jim Cochran (Laramie County Conservation District Manager, retired), Kaila Willis (Wyoming State Engineer's Office), and many Laramie County groundwater users for their time and input on this project.

Abstract

The southeastern portion of Wyoming is an agriculture-dependent area that relies heavily on groundwater from the High Plains Aquifer to grow crops. Like other states across the High Plains region, withdrawal rates in this area are higher than recharge rates, causing groundwater levels to decline. This study uses annual and intra-seasonal farm-level dynamic optimization models to determine whether water-use efficiency (WUE) technologies—specifically soil moisture sensors—can be beneficial to producers if water availability became more limited in the future. Results indicate that WUE technologies can help producers minimize financial losses that might otherwise occur from reduced water availability.

INTRODUCTION

Aquifer depletion has been a growing challenge across the United States due to changes in climate and irrigation pumping rates that exceed annual recharge. This can have negative impacts on agricultural producers in areas dependent on groundwater irrigation (Lansford et al., 1983). The Ogallala Aquifer, also known as the High Plains Aquifer, is the most intensively used aquifer in the United States (Maupin and Barber, 2005). In 2000, 23% of total groundwater withdrawals in the United States and 30% of total irrigation withdrawals were from the High Plains Aquifer (Maupin and Barber, 2005). The High Plains Aquifer provides groundwater for drinking water, livestock production, agricultural production, and mining in the region, which includes Colorado, Kansas, Nebraska, New Mexico, Oklahoma, South Dakota, Texas, and Wyoming (Figure 1). Agricultural production, specifically irrigation use, is responsible for 94% of the total withdrawal from the High Plains Aquifer (Dennehy, 2000). About 19.9 billion gallons of water are pumped from the aquifer per day (Dennehy, 2000). Pumping at these high rates exceeds the annual recharge rate in many parts of the aquifer, which is not sustainable and could lead to a decrease in groundwater availability and thus agricultural production in this region in the future.

Our study area is eastern Laramie County, Wyoming, which includes the towns of Albin, Pine Bluffs, and Carpenter. Though considerable research on aquifer depletion exists, few studies have been done in southeastern Wyoming (Willis, 2019), an agriculture-dependent area that relies on groundwater from the High Plains Aquifer for crop production. Laramie County uses groundwater to irrigate 81% of its total irrigated acres. (Dahlgreen, 2018). In 2015, total groundwater withdrawals for irrigation across the entire state of Wyoming were 602,000 acre-feet, 120,000 acre-feet of which (20%) were withdrawn in Laramie County (Dieter, 2018). Use of irrigation has made Laramie County a top agricultural producer in Wyoming, where it ranks first out of 23 counties in the state for production of wheat for grain, third in corn for grain, and fourth in dry edible beans (USDA, 2012). The economies of Albin, Pine Bluffs, and Carpenter rely almost exclusively on agricultural production, which depends in part on groundwater resources. Area producers have expressed concern about groundwater table declines in the area, which have increased energy costs and reduced available groundwater supplies for some producers.

Producers are interested in understanding the potential economic benefits of adopting water-use efficiency (WUE) technologies compared to using current “rule of thumb” irrigation practices in the area. WUE technologies are instruments that could improve irrigation scheduling throughout the growing season (e.g., soil moisture sensors and variable frequency drives). We analyze the potential for WUE technologies—specifically soil moisture sensors—to decrease energy costs, whether electric, propane, or diesel, and groundwater use while maintaining or improving producers’ net returns. Soil moisture sensors can help producers with irrigation management by measuring how much moisture is in the soil, thus potentially reducing irrigation and improving field-level WUE. Reducing irrigation can decrease electricity costs of production associated with pumping, and can reduce fertilizer loss to runoff and leaching, potentially without reducing physical or economic production (Sharma, 2018). Past research regarding water conservation and adoption of irrigation technologies suggests that the benefits, costs, and economic feasibility of adopting measures such as WUE are likely to be highly variable across regions and crops (Guerrero et al., 2016; Young et al., 2004; Lansford et al., 1984).

Given past literature, this research seeks to answer whether implementing WUE technologies can potentially improve returns, compared to traditional irrigation practices in the area, particularly in the presence of limited water availability. We accomplish this objective by comparing a farm-level model using Discrete Stochastic Sequential Programming (DSSP) that allows for decisions to be made within the growing season in response to changing precipitation conditions versus an annual DSSP model. The annual model version does not allow for changes in irrigation in response to precipitation throughout the growing season, i.e., the model continues to irrigate the same amount throughout the season once a mix of crops has been chosen. The model does not incorporate soil moisture sensors directly—instead, results indicate a range of expected net revenue from adjusting irrigation use in response to changing precipitation, which represents the potential benefits from the adoption of soil moisture sensors or other WUE technologies for the representative farm modeled here.

Groundwater regulators, stakeholders, and producers in the study area recently held discussions to consider policy options to reduce pressure on the aquifer (Willis, 2019). One such policy option is allocation, which would limit the quantity of irrigation water applied on a per-acre basis. Producers in the study area are familiar with the concept of allocation because groundwater withdrawals in adjacent western Nebraska counties are limited by allocation (Willis, 2019), and although discussions in the study area ultimately did not result in adoption of allocation, it could still be adopted in the future. We consequently estimate these models under the full irrigation currently practiced in the region as well as under irrigation constrained by allocation to compare the relative economic benefits of WUE technologies under the two irrigation regimes.

DATA AND METHODS

This study models a representative farm (650 acres under 5 pivots) in eastern Laramie County at both the annual and intra-seasonal time scales. Crops included in the model are irrigated and dryland corn for grain, irrigated and dryland alfalfa, irrigated and dryland winter wheat, irrigated dry edible beans, and a dryland crop rotation. The model has three components: economic, agronomic, and hydrologic.

Economic Component of Model

Our study allows the representative farm to adjust irrigation at several points during the growing season in response to precipitation. It is flexible enough to reveal how intra-seasonal decision-making affects a hypothetical producer's expected profit, yield, groundwater use, and energy costs. DSSP was used for this intra-seasonal model, which allowed our intra-seasonal model to choose deficit irrigation strategies that optimize producers' expected profit, similar to the approach taken by Peck and Adams (2010).

Expected profit ($E\pi$) was determined by using the probability of precipitation occurring at above, near normal, or below levels in each stage (*S1*, *S2*, and *S3*) of the growing season:

$$E\pi = \sum_R(PS1_R) * (NREVS1_R) + \sum_R(PS2_R) * (NREVS2_R) + \sum_R(PS3_R) * (NREVS3_R) \quad (1)$$

where $PS1_R$ is a vector of precipitation probabilities for the first stage (*S1*) representing above, near normal, and below precipitation levels. The parameters $PS2_R$ and $PS3_R$ were similarly constructed for the second and third stages (*S2* and *S3*). We assumed precipitation in a given stage is independent of precipitation in the other stages, therefore, the joint probability of a sequence of precipitation events across the season is simply the product of their independent probabilities. The use of sequential decision variables within the growing season was informed by Houk, Taylor, and Frasier (2000). Our decision variables were as follows: 1) *X1*, the producer's cropping decisions at the beginning of the season; 2) *W2*, the first decision on how much to irrigate after stage 1 (*S1*) precipitation is revealed; 3) *W3*, the second decision on how much to irrigate after stage 2 (*S2*) precipitation is revealed; and 4) *W4*, the third decision on how much to irrigate after stage 3 (*S3*) precipitation is revealed (Figure 2).

$E\pi$ is a function of net revenue under the three possible precipitation realizations in each stage, where net revenue is the revenue (price multiplied by yield) minus the variable costs (net revenue is also known as Returns Over Variable Costs (ROVC)). In this model, the producer's variable costs were broken into five components: 1) total electricity costs, 2) seed costs, 3) water costs, 4) all other variable costs (including wage labor), and 5) irrigation technology costs. Land and management costs were not included in the model. Profit is expected to vary with net revenue, assuming constant fixed costs. Net revenue varies by precipitation realization (*R*) in each season due to costs (e.g., total electricity costs) that vary with irrigation decisions, so the model includes four equations that ensure irrigation decisions are consistent with decisions that have been made at previous stages, such that the model cannot switch mid-season from one crop to another on a pivot-section. Six additional equations served as water balance equations to ensure that more water was not used than allowed on each pivot and for the whole farm. There were also four rotational constraints in the model to ensure that a single crop (i.e., monoculture) was not grown on every pivot section and instead reflect crop rotations common for the study area.

It should be noted that while crop insurance could be used to mitigate short-term risk associated with yield loss from drought for these crops, we did not include crop insurance payments in this model. To include crop insurance, we would have to decide on and use appropriate coverage levels for the area and related insurance costs, then calculate trigger levels and payouts across all scenarios. It was also expected that if long-term depletion and related allocation from the aquifer occurred, crop insurance rates and payoffs for the region would be adjusted as well, but we have no way of knowing what that insurance response might be. Overall, for these reasons, it was felt that addressing crop insurance in the model would detract from the primary objective of this research, which is to understand how potential changes in WUE could affect irrigation decisions and ultimately returns from crop production.

Crop and Price Data

We used data from Willis (2019), who constructed individual enterprise budgets for each crop, in each community, to estimate the costs associated with production. Willis used budgets developed by Klein et al. (2018) as a starting point. Albin, Carpenter, and Pine Bluffs producers confirmed that the modified budgets used by Willis (2019) were comparable enough to use as a foundation for the analysis. In our current study, we used the crop data collected by Willis (2019) for the Pine Bluffs community. Output prices in our model were assumed to be the 15-year (2002-2016) average price for each crop in Wyoming deflated to the same year as the crop budgets used by Willis (2019), as reported by USDA NASS (2017).

Agronomic Component of Model

In the intra-seasonal version of the model, crop yields are a function of precipitation and water applied at different points during the growing season. Precipitation occurring in each stage informs how much a producer chooses to irrigate at *W2*, *W3,* and *W4,* respectively. *S1* includes precipitation from May 7 through June 30, *S2* includes precipitation from July 1 through August 23, and *S3* includes precipitation from August 24 through October 1. Alfalfa has a different planting date (04/01) to account for the precipitation that occurs between 04/01 and 05/07. These dates were chosen based on corn and dry beans planting and harvesting dates and when their growth stages start and end. The USDA has a field crops handbook (USDA, 2010) that outlines the planting and harvesting dates for crops grown in all 50 states, which we used to decide the planting and harvesting dates for corn and dry beans.

Precipitation data are from area weather stations and span the years 1902-2015. However, we used the most recent 30 years of this historical precipitation data (1986-2015), which is standard for this type of research. From this data, we developed a set of precipitation probabilities for each stage, where an individual set reports the probability of each state of nature occurring within each stage. Each stage had three probabilities: the probability that precipitation was above normal (PA), near normal (PN), or below normal (PB). For *S1*, PA = 0.36, PN = 0.27, and PB = 0.37. For *S2*, PA = 0.34, PN = 0.39, and PB = 0.28. For *S3*, PA = 0.40, PN = 0.32, and PB = 0.28. These probabilities inform the calculation of $E\pi$ in Equation 1.

We used AquaCrop to determine the yield responses for our irrigated row crops (corn and dry beans) because of its ability to simulate yield responses in situations of deficit irrigation (Steduto et al., 2009; Steduto et al., 2012). The required inputs for AquaCrop include weather data, crop characteristics, soil profile characteristics, characteristics of the groundwater table, and irrigation and field management practices (Steduto et al., 2012). AquaCrop has default files provided for some crops, soil profiles, groundwater table levels, and irrigation and field management practices. Thus, the minimum observed data needed to parameterize AquaCrop for southeastern Wyoming is climate data.

Climate data came from area weather stations and spanned the years 1957-2015. We used 30 years of this climate data to match the 30 years of precipitation data described earlier. The climate data included maximum temperature, minimum temperature, precipitation, relative humidity (RH), wind speed, and solar radiation. These data were used to calculate reference evapotranspiration (ETo), using the Penman-Monteith conversion equation. The weather data helped to calibrate the AquaCrop model to reflect the climate of Laramie County, Wyoming.

We initially assumed the default crop characteristics provided in AquaCrop for corn and dry beans. The output from running these default parameters showed, however, that some of the crop parameters for both crops needed to be adjusted to reflect typical southeastern Wyoming yields and water application amounts.

Several corn parameters were changed to reflect the High Plains region, including the response to water stress parameters and days between growth stages. These parameters were changed based on parameter values provided in Araya et al. (2017) and Abedinpour et al. (2012). Several dry bean parameters were also changed to reflect the High Plains region. Crop-stage length and growing-season length were provided based on field trials conducted at the University of Wyoming agricultural research station in Powell, WY. Other parameters (e.g., crop response factors) were informed by Espadafor et al. (2017). These region-specific parameters improved AquaCrop's ability to replicate yield and water application levels known to exist in Wyoming, which provided greater assurance that the generated functions give reasonable estimates of the yield-water application relationship for water application levels not generally observed in Wyoming.

The soil profile characteristics for eastern Laramie County were retrieved from the NRCS SSURGO database. The majority of southeastern Laramie County has sandy loam soil, which helped develop specific soil-type characteristics such as soil hydraulic properties, total thickness of soil compartments, total number of soil layers, readily evaporable water, percent sand, percent clay, organic matter, penetrability, saturation, field capacity, wilting point, and saturated hydraulic conductivity (*Ksat*). The default soil file for sandy loam in AquaCrop was used for the simulations.

AquaCrop was not used to estimate the yield-water relationship for alfalfa because, at the time of this research, AquaCrop did not yet have default files available for alfalfa. AquaCrop was also not used to estimate the yield-water relationship for winter wheat due to time constraints. Instead, we used an equation from FAO 33 (Doorenbos and Kassam, 1979) and Bernardo et al. (1987) to simulate the yield response to intra-seasonal irrigation decision-making for alfalfa and winter wheat. This equation indicated that yield

(Ya is actual yield, and Ym is maximum yield expressed in units per area of land such as kg/ha) is a function of crop coefficients, K_{yi}, actual evapotranspiration, Et_{ai}, and potential evapotranspiration, Et_{pi}. The i subscript indicates different stages within the growing season. The equation is as follows:

$$\frac{Ya}{Ym} = \Pi_i^3\left[1 - K_{yi}\left(1 - \frac{Et_{ai}}{Et_{pi}}\right)\right] \quad (2)$$

Initial K_{yi} values came from FAO 56 and were adjusted to reflect local crop stress conditions in Wyoming. The Et_{pi} values came from the observed weather data collected in Cheyenne, WY. Et_{ai} was calculated by summing precipitation, irrigation, and soil moisture contributions. In this study, we assumed that soil moisture contribution is the same throughout the growing season. For the equation to be intra-seasonal, K_{yi}, Et_a, and Et_p varied throughout the growing season.

Hydrologic Component of Model

The hydrologic component of our intra-seasonal model consisted of equations governing lift and pumping costs. $Lift$ is the depth to water, in feet, and helps determine how much pumping water from the aquifer will cost a producer. It was calculated by:

$$Lift_t = Lift_{t-1} + CalRatio * WatUseDepth_{t-1} - Recharge \quad (3)$$

where $WatUseDepth$ represents the irrigation water applied converted to feet, and $CalRatio$ and $Recharge$ are used to calibrate the aquifer to status quo. Status quo represents the aquifer if no changes are made to reduce groundwater use.

Pumping costs were calculated by using the four-step approach from Black and Rogers (1993), who used lift, well pressure, pumping capacity, and pumping hours to determine total electricity costs per pivot section. (Please see Willis (2019) and Grahmann (2020) for details.) Irrigation in the study area is primarily powered by electricity, which we therefore assumed in our model. If other, more costly, energy options had to be employed, the benefits of implementing WUE technologies would be even greater than our estimates indicate.

Annual versus Intra-Seasonal Versions of the Model

The annual version of the model was identical to the intra-seasonal version described above except that the producer no longer had the option of making mid-season changes to irrigation management in response to precipitation. Thus, the only decision variable was the decision of what crops to plant on each pivot section at the start of each season. Crops planted were either fully irrigated throughout the season (D1) or dryland (D3). The producer has no ability to switch to deficit irrigation (D2) at later stages of the season in response to precipitation. This annual version of the model is similar to most studies that have been done on water use in groundwater-dependent agricultural areas (Golden and Johnson, 2013; Brozovic and Islam, 2010; Golden and Guerrero, 2017). The only exceptions of which we are aware are Foster, Brozovic, and Butler (2015) and Hrozencik et al. (2017).

These changes simplify equation (1) by removing the expectation operator and the indices representing mid-season precipitation realizations and decisions:

$$NETREV = \sum_{PS,C}^{1}\left[\left(O_C * Yld_{X,PS,C}\right) - \left(ElectT + WaterRt_C * OtherWaterCost + SeedRt_C * SeedCost_C + OtherCost_C\right)\right] \quad (4)$$

Everything else about the economic component of the model remained the same as it was in the intra-seasonal version. The only impact of the annual version of the model on the agronomic component was that any permutations of precipitation and yield that involved deficit irrigation (D2) were not considered. This reduced the number of permutations from 216 to 81.

The hydrology component of the model was unchanged from the intra-seasonal version described above. Regardless of model version, the hydrology component was annual in the sense that depth to water did not increase over the course of the season in response to pumping. If depth to water were to increase over the course of the season in response to pumping, the additional pumping cost associated with increased depth to water could influence producers to pump less water, depending on aquifer conditions and pumping costs.

Baseline versus Allocation Scenarios

In the Baseline scenario, the farm had 12,000 ac-in of water available (2,400 ac-in per pivot, or approximately 18 ac-in per acre on average), which is more than enough to grow any fully irrigated crop. For example, fully irrigated alfalfa is the thirstiest crop, and 18 ac-in per acre is more than sufficient to grow fully irrigated alfalfa on all pivot sections. In the Allocation scenario, the farm had 7,800 ac-in of water available (1,560 ac-in per pivot, or approximately 12 ac-in per acre on average).

RESULTS

Figure 3 indicates crop mix and irrigation levels by model version and scenario. The annual version of the model, Baseline, replicated the typical crop mix observed in the study area: four half-pivots of alfalfa, two each of corn, dry edible beans, and winter wheat, all fully irrigated (Figure 3, column a). In the annual version, Allocation, two half-pivots of alfalfa were converted to the dryland crop rotation; the other half-pivots remained fully irrigated (Figure 3, column b). We assumed that a producer would never choose to deficit irrigate throughout the whole season, based on conversations with area producers and yield results for the area from AquaCrop. Thus deficit irrigation was not included as an option in the annual version of the model. Expected profit in the annual model decreased by 20.98% ($50,798) between the Baseline and Allocation scenarios, and water use decreased by 26.82% (3.46 ac-in/ac) (Table 1).

In the intra-seasonal model version, deficit irrigation can be used in any season, and the model allows for irrigation adjustments in response to within-season precipitation. In the intra-seasonal model, Baseline, the crop mix is the same as it was in the annual model (Figure 3, column c). However, deficit irrigation was used for irrigated alfalfa, corn, dry edible beans, and winter wheat. (A crop appears as deficit-irrigated in Figure 3 if deficit irrigation was used on the crop in at least one stage, under at least one type of precipitation.) In the Allocation scenario, the crop mix is the same as it was in the annual model except that now, one half-pivot is planted to the dryland crop rotation instead of two (Figure 3, column d). Deficit irrigation is once again used for irrigated alfalfa, corn, dry edible beans, and winter wheat in at least one stage, under at least one type of precipitation.

When intra-seasonal Allocation is compared to intra-seasonal Baseline, there is a 9.79% ($23,995) decrease in expected profits and a 19.44% (1,551.72 ac-in) decrease in water use. Table 1 shows the differences between the two scenarios. There is a larger drop in water use than there is in expected profits on a percentage basis, which is useful information for policy makers. Producers would need to receive approximately $15.69 per ac-in. to consider participating in water use reduction programs.

In the Baseline scenario, expected profits increased by 1.27% ($3,069) in the intra-seasonal model relative to the annual model. Water use declined by 4.81% (0.62 ac-in/ac) in the intra-seasonal model relative to the annual model. This was because the producer found it optimal to deficit irrigate even when water was plentiful (i.e., in the Baseline scenario) to avoid pumping costs. Although there was an increase in expected profits and a decrease in water use in the intra-seasonal version relative to the annual version, this was not a large difference, suggesting that the choice of adopting water saving technology is less impactful when water is plentiful.

The difference in results between the annual and intra-seasonal versions of the model were more substantial for the Allocation scenario. Under the constraint of allocation, expected profits in the intra-seasonal model were 15.62% ($29,872) higher relative to the annual model. Water use in the intra-seasonal allocation scenario was 4.77% (0.45 ac-in/acre) higher relative to the annual allocation scenario. To clarify, allocation reduced expected profit in both models, by $50,798 in the annual model, and by $23,995 in the intra-seasonal model. But the reduction in expected profit was smaller in the intra-seasonal model, thanks to the added flexibility of intra-seasonal deficit irrigation decisions informed by precipitation in each stage.

The difference in profitability between the annual and intra-seasonal models under allocation can help inform producers trying to decide whether it would be feasible to include WUE technologies to help manage irrigation scheduling. Given the crops grown and climate conditions prevalent in eastern Laramie County, expected profits were almost $30,000 (16%) higher for a five-pivot farm when irrigation decisions were updated throughout the season based on precipitation amounts rather than just one decision at the beginning of the year when groundwater was limited. These results suggest that a producer whose operation's characteristics are similar to those modeled here would find it beneficial to spend up to $30,000 on soil moisture sensors or other WUE technologies that could help them adjust irrigation decisions in response to within-season precipitation.

CONCLUSION

This study compared the expected profit and water use between annual and intra-seasonal versions of a farm-level dynamic optimization model. We used the two model versions to simulate whether WUE technologies could be economically beneficial to producers. To do this, we ran two scenarios: a Baseline scenario (i.e., assuming plentiful irrigation water) and an Allocation scenario in which irrigation water is limited. The results of those two scenarios (Baseline versus Allocation) were then compared between the two model versions (annual versus intra-seasonal).

The intra-seasonal Baseline scenario increased expected profits by $3,069 (1.27%) and decreased water use by 0.62 ac-in (4.81%) relative to the annual Baseline scenario. Incorporating intra-seasonal decision-making into the model had minimal impact on expected profits and water use when water was sufficiently available (i.e., in the Baseline scenario). However, incorporating intra-seasonal decision-making had a greater impact on expected profits and water use when water availability was restricted (i.e., in the Allocation scenario). The intra-seasonal Allocation scenario increased expected profits by $29,872 (15.62%) but also increased water use by 0.45 ac-in (4.77%), relative to the annual Allocation scenario.

If we focus on the more realistic intra-seasonal model alone, we see that allocation decreased expected profits by $23,995 (9.79%) and water use by 2.39 ac-in (19.46%) relative to the baseline. These changes in expected profits and water use were impactful for a producer, yet smaller than what was observed when allocation was implemented in the annual model. In the annual model, allocation decreased expected profits by $50,798 (20.98%) and water use by 3.46 ac-in (26.82%). This suggests that implementing WUE technologies (i.e., adjusting irrigation within the growing season in response to precipitation events) can help producers mitigate the negative economic impacts associated with a reduction in available water supplies. If water availability becomes limited or restricted in the future, our results suggest that producers might consider turning to WUE technologies. Soil moisture sensors and other WUE technologies do not explicitly enter the model, but incorporating mid-season irrigation adjustments could generate increases in expected profits, some of which could be used to implement the soil moisture sensors or other WUE technology that could facilitate these profit increases in the first place. If the net benefits of WUE technologies are positive, as our estimates indicate for the representative farm modeled here, WUE technologies could help producers determine whether a crop needs to be irrigated during different parts of the growing season and reduce (but not eliminate) the economic pain of reduced water availability.

We recognize that producers likely already make some adjustments to their irrigation plans following precipitation events, for example, they likely reduce irrigation after a heavy rainstorm, but we have no data to quantify whether this is a general practice in the area. The question is how close to the hypothetical outcomes of the intra-seasonal model might real-world producers be able to get using WUE technologies? Their actual changes in decisions and outcomes would reveal the true value of WUE technologies, as opposed to the full difference between the results of the hypothetical annual and intra-seasonal models presented in this study.

REFERENCES

Abedinpour, M., et al. 2012. "Performance Evaluation of AquaCrop Model for Maize Crop in a Semi-Arid Environment." *Agricultural Water Management* 110: 55–66.

Araya, A., et al. 2017. "Evaluating Optimum Limited Irrigation Management Strategies for Corn Production in the Ogallala Aquifer Region." *Journal of Irrigation and Drainage Engineering* 143 (10): 1–10.

Bernardo, D.J., et al. 1987. "An Irrigation Model for Management of Limited Water Supplies." *Western Journal of Agricultural Economics* 12 (2): 164–173.

Black, R.D., and D.H. Rogers. 1993. "Evaluating Pumping Plant Efficiency Using On-Farm Fuel Bills." Kansas State University Cooperative Extension Service, Manhattan, Kansas.

Brozovic, N., and S. Islam. 2010. "Estimating the Value of Groundwater in Irrigation." Selected paper for presentation at the 2010 Agricultural and Applied Economic Association. https://ageconsearch.umn.edu/record/61337?ln=en.

Dahlgreen Consulting, Inc. 2018. "Final Report: South Platte River Watershed Level I Study." Wyoming Water Development Commission.

Dennehy, K.F. 2000. "High Plains Regional Ground-Water Study." *US Geological Survey No. 091-00.*

Dieter, C.A., et al. 2018. "Estimated Use of Water in the United States in 2015." *US Geological Survey No. 1441.*

Doorenbos, J., and A.H. Kassam. 1979. "Yield Response to Water." *Irrigation and Drainage* 33: 257.

Espadafor, M., et al. 2017. "Simulation of the Responses of Dry Beans (*phaseolus vulgaris l.*) to Irrigation." *American Society of Agricultural and Biological Engineers* 60 (6): 1983–1994.

Foster, T., N. Brozovic, and A.P. Butler. 2015. "Analysis of the Impacts of Well Yield and Groundwater Depth on Irrigated Agriculture." *Journal of Hydrology* 523: 86–96.

Golden, B., and B. Guerrero. 2017. "The Economics of Local Enhanced Management Areas in Southwest Kansas." *Journal of Contemporary Water Research & Education* 162 (1): 100–111.

Golden, B., and J. Johnson. 2013. "Potential Economic Impacts of Water-Use Changes in Southwest Kansas." *Journal of Natural Resources Policy Research* 5 (2-3): 129–145.

Guerrero, B., et al. 2016. "Economic Feasibility of Converting Center Pivot Irrigation to Subsurface Drip Irrigation." *Journal of the American Society of Farm Managers and Rural Appraisers*: 77–88.

Houk, E., G. Taylor, and M. Fraiser. 2000. “Valuing the Characteristics of Irrigation Water in the Platte River Basin.” Paper presented at the 2000 Western Agricultural Economics Association Annual Meetings. https://ageconsearch.umn.edu/record/36454?ln=en.

Hrozencik, R.A., et al. 2017. “The Heterogenous Impacts of Groundwater Management Policies in the Republican River Basin of Colorado.” *Water Resources Research* 53 (10): 757–10, 778.

Klein, R.N., et al. 2018. “2017 Nebraska Crop Budgets.” University of Nebraska Lincoln, Institute of Agriculture and Natural Resources, Lincoln, NE. https://cropwatch.unl.edu/Economics-Real-Estate/2018-Nebraska-Crop-Budgets.pdf.

Lansford, R.R., et al. 1983. “The On-Farm Impacts of the Declining Ogallala Aquifer for the New Mexico High Plains.” *Journal of the American Society of Farm Managers and Rural Appraisers* 47 (1): 31–36.

Lansford, R.R., et al. 1984. “Irrigation Scheduling Models: An Economic Analysis.” *Journal of the American Society of Farm Managers and Rural Appraisers* 48 (2): 58–63.

Maupin, M., and N. Barber. 2000. “Estimated Withdrawals from Principal Aquifers in the United States.” United States Geological Survey.

Peck, D.E., and R.M. Adams. 2010. “Farm-Level Impacts of Prolonged Drought: Is a Multi-Year Event More Than the Sum of Its Parts?” *The Australian Journal of Agricultural and Resource Economics* 54: 43–60.

Steduto P., et al. 2012. “Crop Yield Response to Water.” In *Irrigation and Drainage Paper 66*. Food and Agriculture Organization of the United Nations. https://www.fao.org/3/i2800e/i2800e.pdf.

Steduto, P., et al. 2009. “AquaCrop: The FAO Crop Model to Simulate Yield Response to Water: 1. Concepts and Underlying Principles.” *Agronomy Journal* 101: 426–437.

Sharma, V. 2018. “Methods and Techniques for Soil Moisture Monitoring.” University of Wyoming, College of Agriculture and Natural Resource Extension.

USDA. 2012. “County Summary Highlights.” *Census of Agriculture* 1 (2): 1–55.

U.S. Geological Service. 2013. “High Plains Water-Level Monitoring Study.” https://ne.water.usgs.gov/ogw/hpwlms/.

Willis, K.P. 2019. “Farm-Level Economic Assessment of Alternative Groundwater Management Strategies Over the Ogallala Aquifer in Southeastern Wyoming.” Master’s Thesis. University of Wyoming, Department of Agricultural and Applied Economics.

Young, K.B., et al. 2004. “Value of Water Conservation Improvements on Arkansas Rice Farms.” *Journal of the American Society of Farm Managers and Rural Appraisers* :119–126.

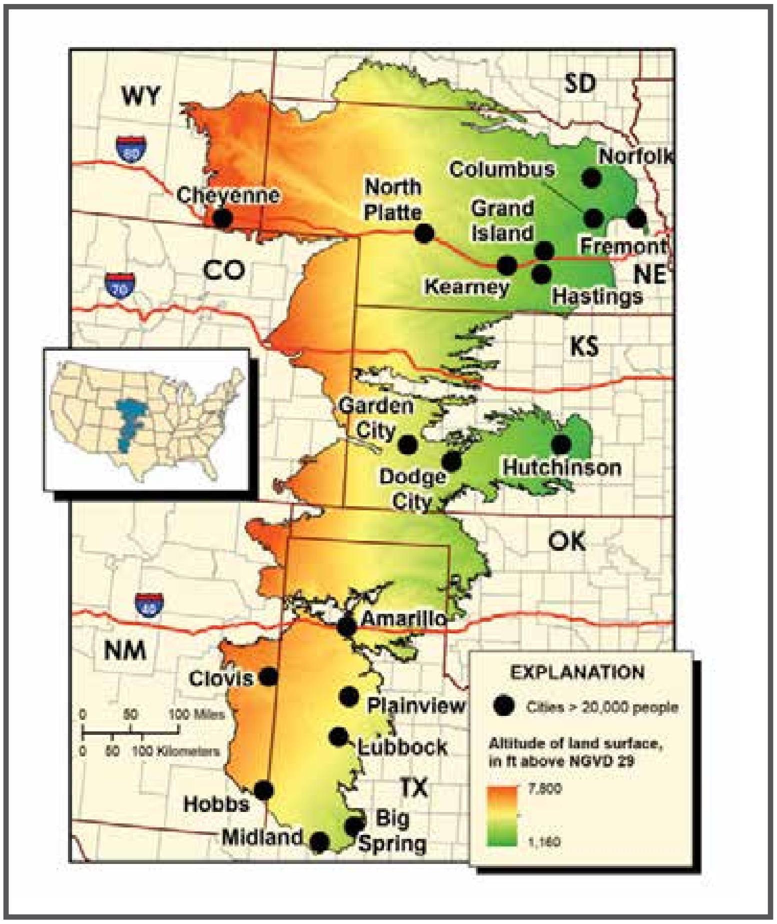

Figure 1. High Plains Aquifer region (Source: USGS, 2013)

Within-Season Timing	X1	Above average Average → W2 (X1, S1) Below average	Above average Average → W3 (X1, S1, Below average	Above average Average → W4 (X1, S1, W2, S2, W3, S3) Below average	E(π)
	Planting Time	**Early Season**	**Mid Season**	**Late Season**	**Harvest**
States of Nature by Stage		Stage 1 (**S1**): Early-season precipitation	Stage 2 (**S2**): Mid-season precipitation	Stage 3 (**S3**): Late-season precipitation	
Choice Variables	**X1**: Choose crops	**W2**: Choose early-season irrigation	**W3**: Choose mid-season irrigation	**W4**: Choose late-season irrigation	

Notes: Crop choices at X1 are corn, alfalfa, winter wheat (all irrigated or dryland), dry edible beans (irrigated), dryland crop rotation. Irrigation level choices at W2, W3, and W4 are full, deficit, or zero. Irrigation level choices are made conditional on crops planted, earlier precipitation, and earlier irrigation decisions.

Figure 2. Visual representation of decision points throughout the irrigation season

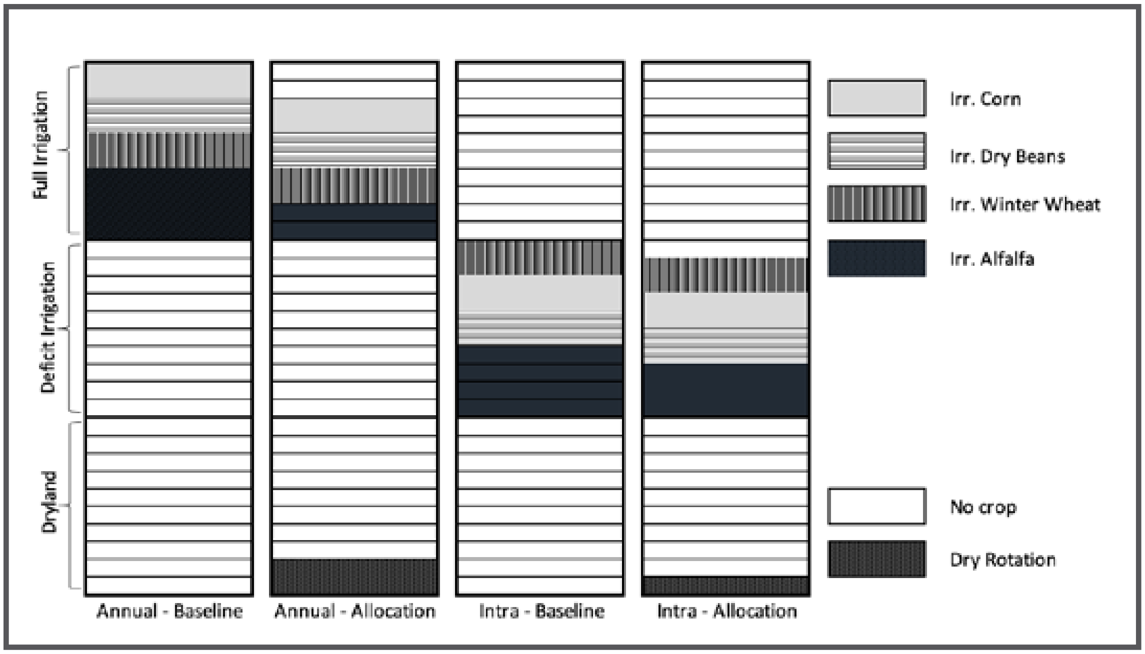

Figure 3. Crop mix and irrigation levels by model version and scenario (each bar represents one half-pivot)

Table 1. Comparison of Average Returns Over Variable Costs and Water Use by Model Version and Scenario

Scenario	Average ROVC	Water Use	
	$	Average ac-in/ac	Total ac-in
Annual Model - Baseline	$242,090	12.90	8,384
Annual Model - Allocation	$191,292	9.44	6,135
Intra-Season Model - Baseline	$245,160	12.28	7,981
Intra-Season Model - Allocation	$221,165	9.89	6,430

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

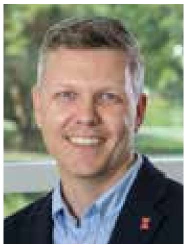

By Joseph P. Janzen

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

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

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

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

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 . 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

$$Gross\ Return = \frac{P_{t+1}^{D} - P_{t}^{N}}{P_{t}^{N}} \tag{3}$$

This calculation limits the number of observations available for two reasons. First, many farms in my sample have no near-to-harvest sales $(Sales_t^N = 0)$, 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.

Kuethe, 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.

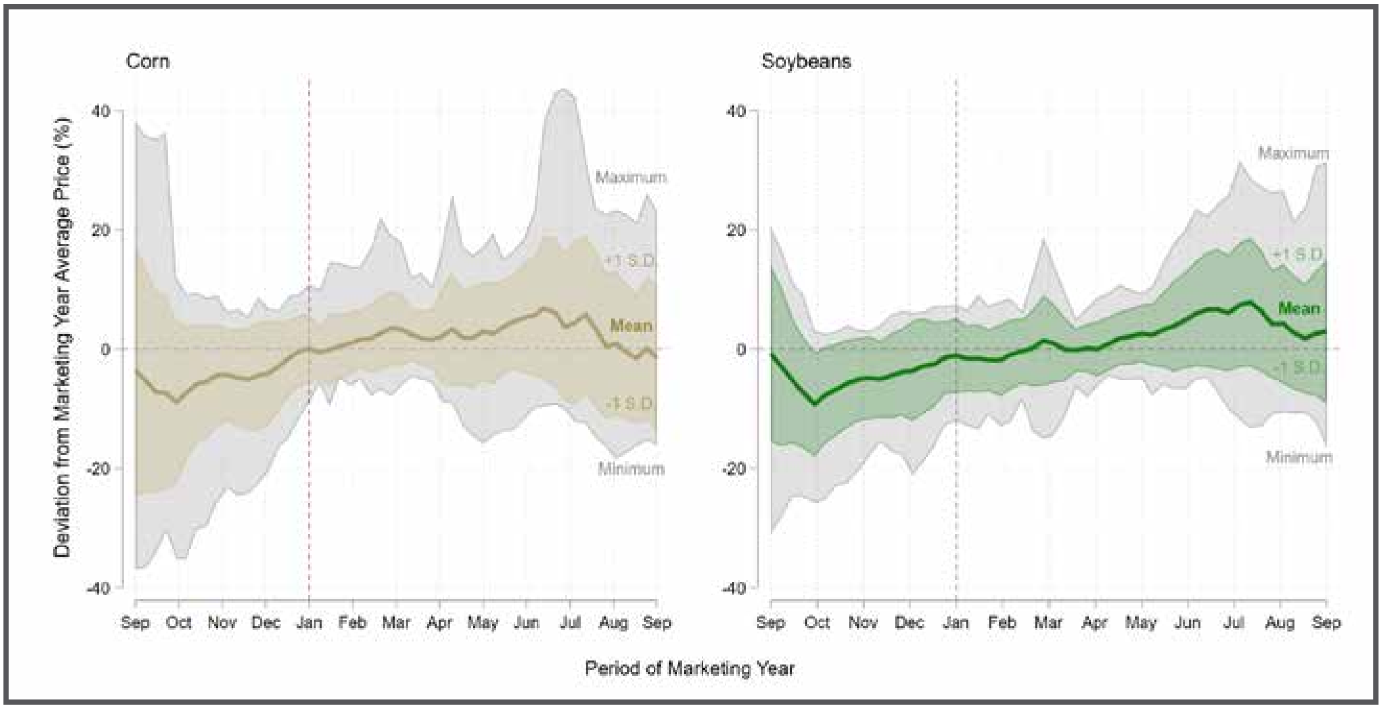

Figure 1. Typical and extreme seasonal variation in corn and soybean prices in Central Illinois, 2004–2020 (Mean (thick) lines indicate the average across all years of the weekly deviation from the marketing year average price, shaded areas indicate the range given by one standard deviation above and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

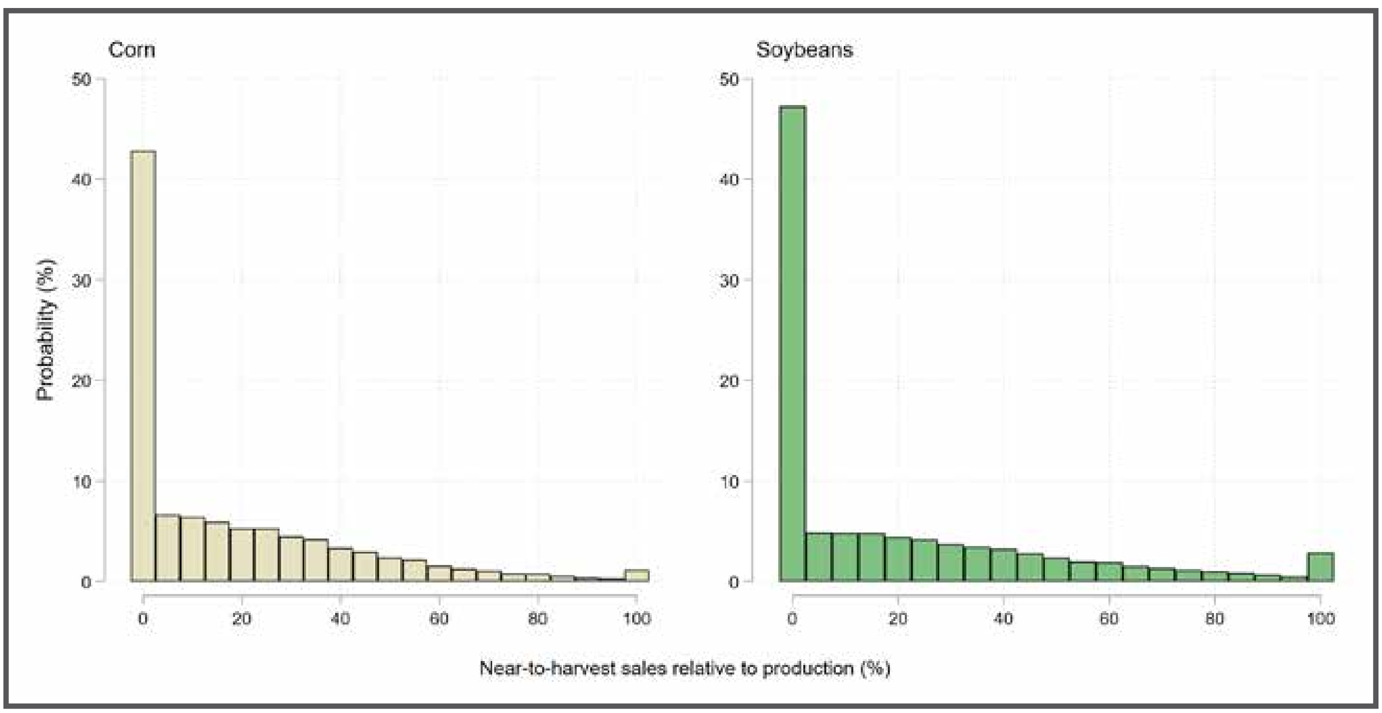

Figure 2. Distribution across farms and years of the proportion of corn and soybean sales made near to harvest

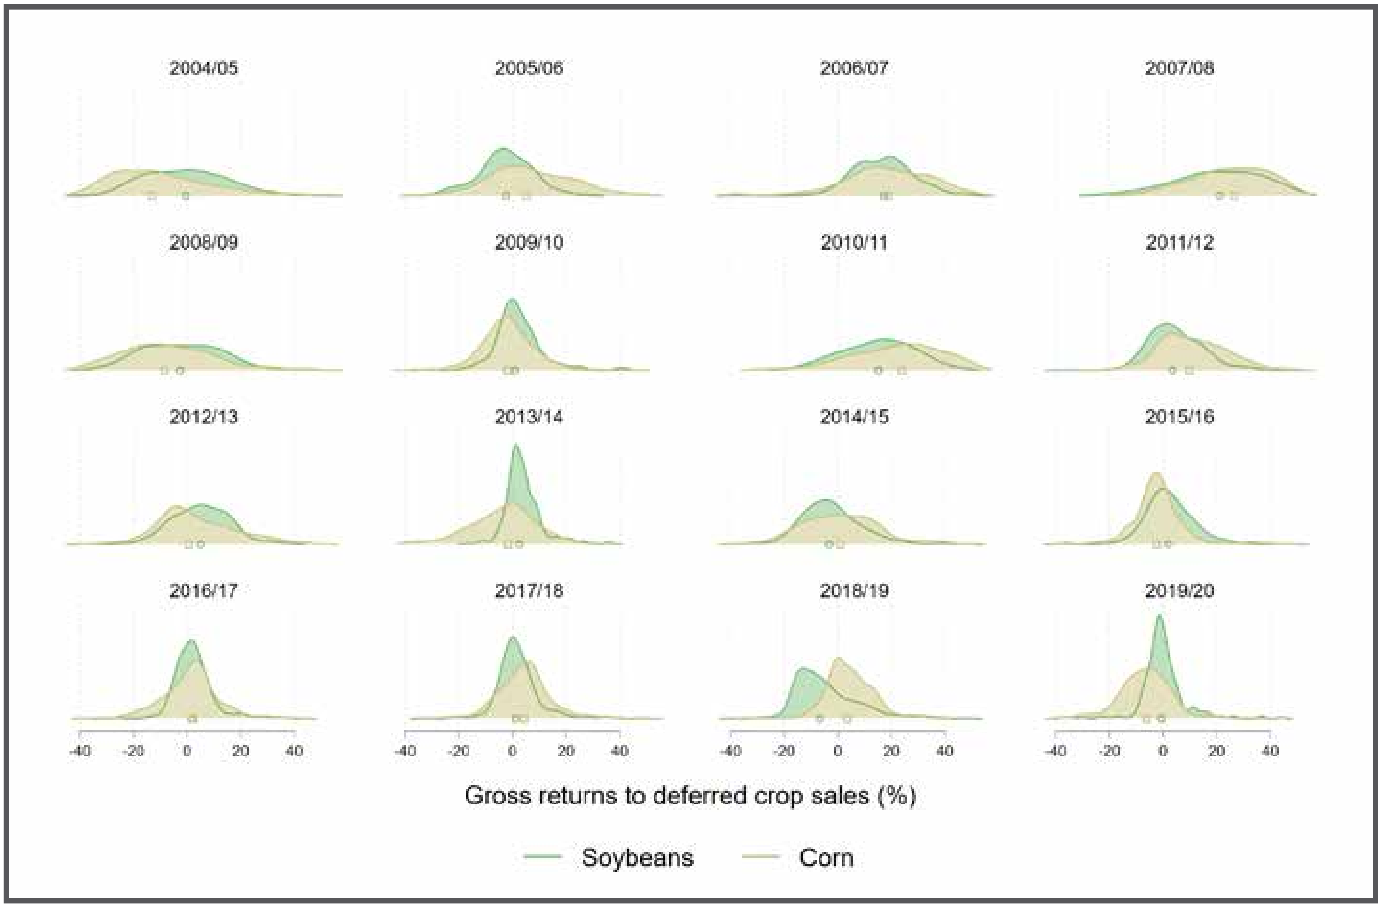

Figure 3. Distribution of gross returns from deferred (post-January 1) crop sales by commodity and marketing year, 2004/05 to 2019/20 (Markers below each distribution function indicate the median gross return in that marketing year)

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