Staff Paper

OPTIMAL POST-HARVEST GRAINS STORAGE BY RISK AVERSE FARMERS

Jing-Yi Lai, Robert J. Myers, and Steven D. Hanson

Staff Paper No. 01-50 **November, 2001**

Department of Agricultural Economics
MICHIGAN STATE UNIVERSITY
East Lansing, Michigan 48824

OPTIMAL POST-HARVEST GRAIN STORAGE BY RISK AVERSE FARMERS

Jing-Yi Lai, Robert J. Myers, and Steven D. Hanson
myersr@msu.edu

Abstract

Most previous research on post-harvest grain storage by farmers has assumed risk-neutral behavior and/or made restrictive assumptions about underlying price probability distributions. In this study we solve the optimal post-harvest storage problem for a risk averse farmer under more general assumptions about underlying price distributions. The resulting model is applied to Michigan corn farmers and results show that, contrary to the sell all-or-nothing risk-neutral rule, risk averse farmers will spread sales out over the storage season. The optimal pattern for sales by Michigan corn farmers is to sell approximately 50% of corn at harvest in November (a risk-reduction strategy) and approximately 40% in May (a return-enhancing strategy).

(33 Total Pages)

OPTIMAL POST-HARVEST GRAIN STORAGE BY RISK AVERSE FARMERS

Many farmers store all or part of their grain at harvest in an attempt to profit from expected price increases over the storage season (Sartwelle et al.). On-farm storage provides flexibility in choosing the timing of sales and can provide increased returns to farmers selling at the right time (at a relatively high price) versus the wrong time (at a relatively low price). But while on-farm storage can increase expected returns it also increases risk. A farmer who sells at harvest has no remaining price risk but one who stores grain for future sale faces the risk that prices will move lower over the storage season. Some of this risk could be managed by using futures or options but this entails transaction costs and might also eliminate some of the potential reward if prices move higher.

Most studies of on-farm grain storage do not account for farmer risk aversion. Early research evaluated the performance of essentially arbitrary *a priori*, fixed marketing strategies to see how they might influence the mean and variability of returns (see the review by Zulauf and Irwin). This approach has the advantage of not having to assume *anything* about farmer behavior or risk preferences. The obvious disadvantage, however, is that while an arbitrary storage strategy may generate positive expected returns, there is no way to determine if such a strategy is optimal given a well-specified farmer objective function. On the other hand, studies of *optimal* on-farm storage have found it difficult to incorporate farmer risk aversion.

In a recent study, Fackler and Livingston (FL) develop a simple dynamic programming solution to the optimal post-harvest grain storage problem in a useful framework that is relatively easy to implement. In particular, one advantage of FL is that they impose a non-negativity constraint (farmers cannot buy back grain into storage once it has been sold) and a short selling constraint (farmers cannot sell more grain than they have in storage). These seem like reasonable restrictions to impose for on-farm storage because transaction costs generally preclude such behaviors and, indeed, such behaviors are rarely observed in practice. However, FL also assume farmers are risk neutral which causes the solution to take the form of an optimal stopping rule—at each point in time it is optimal to either wait and sell nothing or to liquidate all stocks and stop the process. These features are at odds with observed on-farm storage behavior because farmers appear to be risk averse and to spread grain sales out over the storage season (USDA). Other studies that have examined various aspects of optimal grain storage decisions assuming risk neutral storage behavior include Tronstad and Taylor, and Lence, Kimle, and Hayenga.

A few researchers have incorporated risk aversion into optimal on-farm storage models. For example, Berg allows for farmer risk aversion in his dynamic programming study of on-farm wheat storage in the EU. Berg finds that partial sales (i.e. spreading sales out over the storage season) can be optimal for risk averse farmers. However, Berg assumed triangular probability distributions for prices, and that price probability distributions in any one period are completely independent of price outcomes in past periods. These seem like very restrictive assumptions given what we now know about the probability structure of most grain price movements (e.g. Yang and Brorsen; Baillie and Myers; Wang, et al.).

The objective in this study is to develop optimal strategies for post-harvest grain storage assuming risk averse farmers and more realistic assumptions about price probability distributions. Stochastic dynamic programming in a discrete-time framework is used to derive optimal post-harvest marketing decisions under these conditions. The results extend FL's risk neutral model by incorporating farmer risk aversion. In our model an optimal partial selling rule is derived, as opposed to the sell-all-or-nothing rule developed by FL. This means it is optimal to spread sales out over the storage season, as is more commonly observed in practice. The resulting optimal storage rules are applied to the post-harvest storage problem faced by Michigan corn producers. Results provide empirical support for spreading sales out over the storage season as a risk management strategy. Sensitivity analysis is conducted to investigate how various factors influence the optimal timing of sales. Finally, the performance of the optimal storage rules are evaluated and their economic value is measured using a willingness-to-pay concept.

Theoretical Model

Consider a risk-averse farmer with on-farm storage facilities who intends to store harvested grain to profit from an expected increase in cash price during the storage season. The storage season begins at the current harvest, ends before next year's harvest, and is divided into T equal-sized decision nodes. At the beginning of each decision node, the farmer has a current stock level s_t, observes the current market price p_t, and chooses an amount of the commodity q_t to sell in the spot market. We follow FL and impose the restriction $0 \leq q_t \leq s_t$ to ensure that sales must be non-negative (no re-purchase of grain allowed) and less than or equal to the current storage level (no short selling allowed). These assumptions are because transaction costs would generally preclude such behaviors. We also assume no storage hedging on futures markets. This simplifies the model and is consistent with the fact that few farmers make extensive direct use of futures and options, with nearly two-thirds of all grain marketed directly through cash sales (Musser, Patrick and Eckman; Sartwelle et al.).

Each period the producer pays a storage cost of $\alpha > 0$ per unit of stock to carry storage into the next period. By definition, storage is subject to the transition equation:

$$s_{t+1} = s_t - q_t. \tag{1}$$

We assume farmers are interested in the compounded cash flows from their storage operations which will be called "wealth", w_t. At any period wealth satisfies the transition equation:

$$w_{t+1} = (1 + r)[w_t + p_t q_t - \alpha(s_t - q_t)] \tag{2}$$

where r is the (assumed constant) interest rate and initial wealth, w_0, is given.

The farmer is assumed to choose a sequence of sales $\{q_t\}_{t=0}^{T-1}$ to maximize the expected utility of terminal wealth:

(3) $$\max_{\{q_t\}_{t=0}^{T-1}} E_0 U(w_T)$$

subject to the transition equations (1) and (2), the constraint $0 \le q_t \le s_t$, and a Markov probability process for prices $\tilde{p}_{t+1} \sim f_t(p_{t+1} \mid p_t, z_t)$ where f_t is the probability density function for the cash price at $t+1$ conditional on the current cash price p_t and a vector of current information variables z_t. The $U(\cdot)$ function is an increasing and concave von Neumann-Morgenstern utility function representing farmer risk preferences.

The problem can be solved using discrete time stochastic dynamic programming (Bertsekas; Miranda and Fackler). Defining $x_t = (w_t, s_t, p_t, z_t)$ as the state vector and $v_t(x_t)$ as the value function, Bellman's equation for the problem is given by:

(4.a) $$v_T(x_T) = U(w_T)$$

(4.b) $$v_t(x_t) = \max_{q_t} E_t\{v_{t+1}(x_{t+1})\} \text{ for } t = 0, 1, \ldots, T-1$$

subject to the transition equations (1) and (2) and the constraint $0 \le q_t \le s_t$. It is shown in the appendix that necessary conditions for a solution are:

(5.a) $$q_{T-1} = s_{T-1}$$

(5.b) $$\frac{\partial v_t(x_t)}{\partial w_t}\, p_t - \frac{\partial v_t(x_t)}{\partial s_t} - \lambda_t \le 0 \quad \text{for } t = 0, 1, \ldots, T-2$$

(5.c) $$q_t\left[\frac{\partial v_t(x_t)}{\partial w_t}\, p_t - \frac{\partial v_t(x_t)}{\partial s_t} - \lambda_t\right] = 0 \quad \text{for } t = 0, 1, \ldots, T-2$$

(5.d) $$\lambda_t\,(q_t - s_t) = 0 \qquad\qquad \text{for } t = 0, 1, \ldots, T-2$$

and $q_t \ge 0$, $\lambda_t \ge 0$, and $s_t - q_t \ge 0$. Here, λ_t is the shadow value of relaxing the short selling constraint that requires $q_t \le s_t$.

At the last decision node $T-1$ the optimal decision is clearly to set $q_{T-1} = s_{T-1}$ (sell everything left in storage, if any) whenever $p_{T-1} > 0$. This is because stocks have no value in the final period T but wealth does. This ensures that bins are emptied before the next harvest. At all periods prior to $T-1$ the farmer faces a trade-off—either sell all or part of total stocks now (if any is left) and receive p_t, or sell nothing and wait to see if prices rise. There are four cases to consider.

First, suppose that s_t is zero (there is no storage left). In this case the (trivial) optimal strategy is to set $q_t = 0$ because this is the only choice in the opportunity set.

Second, suppose that current stocks are positive $(s_t > 0)$ and the optimal choice is still to sell nothing and wait $(q_t = 0)$. Then $q_t < s_t$ and $\lambda_t = 0$. Furthermore, from (5b) and (5c) we have

$$\frac{\partial v_t(x_t)}{\partial w_t} p_t \leq \frac{\partial v_t(x_t)}{\partial s_t}. \tag{6}$$

In this case the marginal value (in terms of expected utility of terminal wealth) of selling the first bushel of grain out of storage must be less than or equal to the marginal value (in terms of expected utility of terminal wealth) of keeping that bushel of grain in storage. So the farmer chooses not to sell anything and if (6) is satisfied with strict inequality would actually buy grain to store except for the constraint $q_t \geq 0$.

Third, suppose that current stocks are positive $(s_t > 0)$ and the optimal choice is to sell everything $(q_t = s_t)$. Then $\lambda_t \geq 0$ and (5b) and (5c) imply

$$\frac{\partial v_t(x_t)}{\partial w_t} p_t \geq \frac{\partial v_t(x_t)}{\partial s_t}. \tag{7}$$

In this case the marginal value (again in terms of expected utility of terminal wealth) of selling the last bushel of grain is greater than or equal to the marginal value of keeping that bushel of grain in storage. If (7) is satisfied with strict inequality the farmer would like to sell more if he/she had it but is constrained by $q_t \leq s_t$.

Fourth, suppose the optimal strategy satisfies $0 < q_t < s_t$. Then $\lambda_t = 0$ and (5b) and (5c) imply that

$$\frac{\partial v_t(x_t)}{\partial w_t} p_t = \frac{\partial v_t(x_t)}{\partial s_t}. \tag{8}$$

The farmer is indifferent between selling or keeping the marginal bushel of grain. The concavity of $U(\cdot)$ ensures that $v_t(x_t)$ is concave in q_t which implies there may be a wide range of conditions (price, storage costs, risk preferences, etc.) under which there are interior solutions satisfying $0 < q_t < s_t$.

It is the existence of an interior solution (case four) that distinguishes this model from the one in FL. The linearity of the (risk neutral) FL model leads to a "sell all or nothing" rule whereas the (risk averse) model here provides an incentive to spread sales out over the storage season. Of course, if current prices get high enough everything will be sold and if current prices get low enough nothing will be sold. At some intermediate range of current prices, however, the risk averse farmer will sell some grain and keep some in storage.

An Application

In order to implement the theoretical model of the previous section and investigate the partial sales phenomenon empirically we need to make a number of assumptions and solve the model numerically. Here we apply the model to the optimal post-harvest storage problem faced by corn farmers in Michigan.

Assume Michigan corn farmers make their first marketing decision at harvest (which we define as the first week of November) when they decide how much to sell right away and how much to store and sell later. We then break the storage season up into five additional decision nodes (the first weeks of January, March, May, July and September). At each of these six decision nodes (harvest plus the five through the storage season) the farmer can sell some or all of his/her corn. Any storage left over at the first week of September is automatically sold to make way for the coming new crop.

The control and state spaces for this problem might best be viewed as continuous but we use a discrete approximation to facilitate a discrete state and control space solution technique. The control space is specified as a proportion of the total harvest, s_0, sold in any period. The proportion of the harvest sold at any decision node is assumed to take on one of eleven possible values $q_t \in \{$ 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0} . The state space for current stocks is also specified as a proportion of the total harvest and so obviously consists of these same eleven possible proportions. The state space for selling prices was specified as 15 possible price states ranging from $1.60 to $4.40 per bushel in 20 cent increments. Each price in the state space is viewed as the mid-point of the underlying continuous price interval. The upper bound price state of $4.40 per bushel represents the price interval of $4.30 or above while the lower bound price state of $1.60 per bushel represents the price interval of $1.50 or below.

Wealth is also expressed as a proportion of the total harvest. We assumed an initial wealth state of zero and computed the state space for wealth as all possible feasible combinations of wealth per unit of total harvest that could come from any combination of feasible price and marketing (sales) strategies over the entire time horizon of the storage season. An initial wealth state of zero implicitly assumes the storage operation is viewed as separable from other farm production and consumption decisions, and so the storage decision only depends on wealth generated by that decision. To compute the wealth states we need to make assumptions about interest rates (the rate of return on accumulated wealth) and storage costs. To examine the solution's sensitivity to interest rate and storage cost assumptions we used annualized interest rates of r = 5%, 10%, and 15% and monthly storage costs of α = 0, 1, and 2 cents per bushel. Benefits of specifying the wealth states to conform with all feasible price outcomes and sales decisions are that specification of the value function does not require interpolation (a discrete value function value can be computed for every current state in the state space), and it is efficient because infeasible wealth outcomes do not need to be evaluated.

The final assumptions required to operationalize the model are a utility function and a set of transition probabilities for transitioning from one price state to another. We assumed a constant relative risk aversion (CRRA) utility function $U(w_T) = w_T^{1-b}/(1 - b), \quad b > 0$ where w_T is the terminal wealth per bushel of initial harvest .[1] The parameter b denotes the coefficient of relative risk aversion and is set to one of two possible values in order to compare results for near risk neutrality (b = 0.0001) with those from risk averse behavior (b = 5).[2] Computation of the price transition probabilities required some detailed empirical analysis.

Computation of Price Transition Probabilities

The probability density functions for corn prices faced by the farmer at each decision node, t, are represented by a set of discrete-time transition probability matrices which map the stochastic price states across the marketing periods. The underlying stochastic structure is estimated using weekly cash corn closing prices each Wednesday at Saginaw, Michigan starting the first week in October of 1975 and ending the last week in September of 1996. The estimated price process is represented by an autoregressive seasonal model for the conditional mean of the logarithmic price changes, and a generalized autoregressive conditional heteroskedastic t-distribution model (GARCH-t) with seasonality for the conditional variance of the innovations (Fackler, Bollerslev). These specifications have been found to do a good job of representing the probability structure of weekly grain price movements (Yang and Brorsen). After investigating several alternative models for goodness of fit a preferred model was chosen and estimated (table 1). Weekly Saginaw corn prices were found to be nonstationary and the conditional mean and variance of price changes appear to vary over time as a result of both stochastic and seasonal factors. The Ljung-Box Q-statistics show the model is well specified in terms of removing autocorrelation from both the errors and squared standardized errors of the change in log prices.

The estimated price model can be used to generate transition probabilities that correspond to the fifteen possible price states in the price state space. These transition probabilities are then used to solve the model and determine the optimal marketing rules. We allow the transition probabilities to be different at each decision node, t, and so we actually compute five different sets of transition probabilities. The transition probabilities are generated from the estimated weekly price model (table 1) using simulation techniques.

The simulation process begins by going to the first decision node (the harvest period specified as the first week of November). Then the first possible initial price state ($1.60) is selected. Using this price state as a starting point, the weekly econometric model is simulated eight to nine weeks ahead to get one random draw for the price state in the first week of January (the next decision node). The simulation of this one random draw is obtained by making random draws on a sequence of eight to nine ϵ_i values and tracing through the price and variance outcomes using the econometric model (see table 1). In the absence of any information about the initial value of the conditional variance of ϵ_i, an estimate of the unconditional variance of ϵ_i is used to initialize the simulation of σ_i^2.

After this process is complete we now have one random draw on the price for the first week of January given that the price during the first week of November is $1.60. Now repeat this process 10,000 times using a random number generator for the ϵ_i and check the relative frequency with which the price outcomes fall into each of the price intervals in the price space. These relative frequencies are used as the transition probabilities for transitioning from the initial price state $1.60 in the first week of November to each of the alternative price states in the state space in the first week of January.

To generate the entire matrix of transition probabilities we repeated this process for every possible initial price state at the harvest period in the first week of November, {$1.60, $1.80, $2.00, $2.20, $2.40, $2.60, $2.80, $3.00, $3.20, $3.40, $3.60, $3.80, $4.00, $4.20, $4.40}. This provides the entire matrix of transition probabilities for transitioning from the first week of November to the first week of January. To generate a transition probability matrix for every decision node we repeated this whole procedure for transitioning between all decision nodes (November-January, January-March, March-May, May-July, July-September). This gives us five different transition probability matrices that can be input into the dynamic programming algorithm.

To validate these transition probabilities, and make sure they are a reasonable approximation to the underlying probability distribution of prices, we conducted a number of experiments. In one of these experiments we set the initial price in November equal to its historical data mean of $2.24 and simulated this price forward over the entire storage season using 10,000 replications. The resulting relative frequencies for the price outcomes in each month are shown in figure 1. As expected, the initial price has a big impact on the mean of future price distributions. That is, there is a high probability that prices will remain around $2.24, especially in the early part of the storage season. Notice, however, that the probability distributions tend to shift to the right as the storage season progresses (higher probability of prices increasing than decreasing), as well as becoming more spread out (prices at the end of the storage season more uncertain than prices at the beginning). Towards the end of the storage season in September, however, there appears to be a higher probability that prices will be lower compared to immediately preceding decision nodes. This seasonal pattern of price movements can be seen more clearly by comparing actual historical average prices for each week over the storage season with the associated average simulated price (again using 10,000 replications and starting all of the simulations from the historical average harvest price of $2.24 per bushel). A graph of these historical and simulated weekly averages is shown in figure 2. Prices clearly have a tendency to rise throughout the storage season until about June when prices begin to fall in anticipation of the coming harvest. This is the kind of pattern we would expect and the simulated average prices do a good job of tracking the actual historical average weekly prices, though the path of simulated average prices is smoother (as expected). This should provide some confidence that the calculated transition probabilities are generating outcomes that are consistent with the underlying probability structure of actual corn price movements over the storage season.

Dynamic Programming Algorithm

The algorithm used to solve the storage model is a discrete time, discrete state and control dynamic programming algorithm based on the approach of Miranda and Fackler and programmed and solved in

GAUSS. The algorithm proceeds by evaluating the value function at every (discrete) point in the state and control space and so can take a long time to solve (in this case several hours on a personal computer).

Results

We first solved the model for the case of a risk neutral farmer (b = 0.0001). This leads to the same "sell all or nothing" marketing strategy derived in FL (as expected). Hence, in this case the optimal marketing strategy is defined by a cutoff price for each decision node, such that if the current price is above the cutoff price sell everything and if it is below then sell nothing.

The cutoff prices at each decision node for the risk neutral case (b = 0.0001) are presented in table 2. Using the base case (the first column) as an example, for the month of November the cutoff price is $4.10 per bushel. Hence, if the cash price in the first week of November is $4.10 per bushel or higher it is optimal to sell everything in storage; otherwise, the optimal marketing strategy is to retain the entire stock. The results show the optimal cutoff price starts with a relatively high value at the beginning of the marketing season and decreases as the end of the marketing season approaches. This is because at the earlier stages of the marketing season there are more future time periods in which prices might rise. The cutoff prices for different assumptions about storage costs and interest rates are also presented in table 2. The results show the cutoff price declines as storage cost increases. This is intuitive because storage cost is a negative income incurred when holding the grain. An increase in storage costs suggests a decrease in the value of waiting to sell stocks in the future, which makes storage less desirable, and lowers the current cash price needed to compensate for the foregone benefits of future sales. Similarly, an increase in interest rates triggers a lower cutoff price because the potential increase of interest income represents an increased opportunity cost of holding the grain, which makes storage less attractive.

Moving to the case of risk aversion (b = 5) the optimal marketing strategy for the harvest period (first week of November) is represented graphically in figure 3 for the base case of 10% interest rate and storage cost of $0.01 per bushel per month. Because this is the first decision node the unique storage state is 100% of the total harvest and the unique initial wealth state is zero. Therefore, the optimal marketing strategy depends only on the current price. The results show that if prices are at or above $4.10 at harvest the optimal strategy is to sell everything and if it is $1.70 or below the optimal strategy is to sell nothing and wait. At intermediate prices, however, there are partial sales. If the current price is between $1.70 and $1.90 the optimal strategy is to sell 50% of the crop at harvest; if it is between $1.90 and $3.70 the optimal strategy is to sell 60% of the crop at harvest, and so on. Risk aversion creates a clear incentive to sell at least some of the harvest right away and not store it.

The optimal marketing strategy for this same risk averse (b = 5) farmer in January is more difficult to represent because the state space has many more dimensions (the possible storage and wealth states are no longer unique). To demonstrate the rule using two-dimensional graphs we construct three separate graphs for three different wealth states—high wealth (in which previous sales were made at relatively high prices), middle wealth (in which previous sales were made at intermediate prices) and low wealth (in which previous sales were made at relatively low prices). For each of these wealth levels there can still be any one of 11 possible storage states, $s \in \{ 0, 0.1, ..., 1.0\}$, depending on how much corn was sold in November. The optimal strategy for $s = 0$ is simply to sell nothing (because all of the corn has already been sold). The other 10 possible storage states are each represented by a different line in each panel of figure 4, with each of the panels representing a different wealth (previous price) level.

Optimal strategies for January are more complex and less intuitive than for November. Results for January are consistent in the sense that very low prices lead to no sales while very high prices lead the farmer to sell everything. At intermediate prices there may be partial sales. Notice, however, that in some cases, particularly when there is a lot of storage left (minimal sales in November) the optimal strategy is to sell less at intermediate prices than at low prices (see the negative slope of some of the decision rules in figure 4 over some regions of the price space). The reason for this is that the current price state not only represents the benefits of immediate sale but is also a signal determining the probability of prices going higher in the future. Therefore, a higher current price encourages more sales because it allows risk-free collection of a relatively high return, but it may also signal a higher probability of price increases in the future which would discourage current sales. If the latter effect dominates the former then a higher current price may lead to less current sales.

The results in figure 4 also show that, for a given storage level and current price, the optimal marketing strategy at high wealth states is to sell less in the current period than in low wealth states. This is consistent with the CRRA utility function because absolute risk aversion decreases with an increase in wealth and so the farmer is more willing to take the gamble of waiting to see if prices rise the higher is his/her wealth level (the more money he/she has already received from previous sales).

Optimal storage rules for the other decision nodes in March, May, and July were also computed but results are not shown here to conserve space. These results show that partial sales may be optimal at each decision node (except the last) when farmers are risk averse, and that the price required to encourage sales falls as we move through the storage season (as expected). One way to summarize these results is to use the estimated transition probabilities and optimal storage rules to compute the likelihood of sales occurring in various months under both risk neutrality and risk aversion, using the base case of r = 10% and storage costs of $0.01 per bushel per month. These results are summarized in figure 5. Under risk neutrality optimal sales in any given year will always occur in one month only (sell all-or-nothing rule), but the figure shows the expected distribution of those sales over repeated samples (years). The resulting distribution is unimodal with sales occurring in May 83% of the time and sales occasionally occurring in March, July, and September as well. Under risk aversion the expected distribution of optimal sales over the storage season is bi-modal with 52% of sales occurring right at harvest, 39% of sales in May, and other sales occasionally occurring in other months. Clearly, the optimal strategy under risk aversion is to spread sales out over the storage season but most of the sales will occur either at harvest (a risk reduction strategy) or in May (an expected profit generating strategy).

Performance of the Optimal Storage Rule

One way to evaluate the performance of the optimal storage rules is to compare outcomes in these cases to what would have happened if all corn was sold at harvest (no storage). Using the transition probability matrices and simulation results the mean and unconditional variance of final wealth was computed under three scenarios—sell everything at harvest, apply the optimal risk neutral storage rule, and apply the optimal risk averse storage rule. Results are provided in the first two lines of table 3 using an interest rate of 10% and a storage cost of $0.01 per bushel per month. Selling everything at harvest generates the lowest mean wealth and the highest unconditional wealth variance. On average, prices rise enough over the storage season so that selling at harvest generates the lowest mean return. Nevertheless, if all corn is sold at harvest there is no risk of future price declines so, from this perspective, selling everything at harvest is a risk minimizing strategy. The unconditional variance measure indicates the year-to-year variability in final wealth and should not be misconstrued as a measure of short-term risk. This unconditional variance measure shows that, even though selling at harvest is a risk minimizing strategy because it eliminates the risk of experiencing price declines during the storage season, from a long-run perspective this strategy still generates considerable year-to-year variation in returns.

The optimal storage strategy for a risk neutral farmer generates the highest mean wealth but also the highest risk because sales rarely occur at harvest under this strategy, and so farmers are usually fully exposed to the risk of price declines over the storage season. Nevertheless, the unconditional variance estimate shows that long-run, year-to-year variability in returns under this storage rule is slightly lower than for the case of selling everything at harvest.

The optimal storage strategy for a risk averse farmer generates a mean return that is higher than selling at harvest but lower than applying the optimal risk-neutral strategy. Risk is reduced because some sales usually occur at harvest, but not completely eliminated because some corn is usually put into storage for later resale as well. Long-run, year-to-year variability in returns is reduced considerably by using this "diversification" strategy.

Another way of looking at performance is to compute the farmer's willingness to pay (WTP) for the optimal marketing strategy and use this as a quantitative estimate of the economic value of the marketing decision rule. We define WTP as the maximum certain monetary amount the farmer is willing to pay for the right to apply the optimal marketing strategy rather than sell all production immediately at harvest. WTP values are calculated by solving the equation

$$U[\hat{w}_T + \mathrm{WTP}(1 + r)^T] = E_0 U(w_T^*) \tag{9}$$

using the bisection method. In equation (9), w_T^* is the terminal wealth under the optimal storage rule, $\hat{w}_T$ is the terminal wealth under a sell everything at harvest rule, and the probabilities for computing the expectation are calculated using the relative frequencies of the historical harvest period prices together with the transition probability matrices computed earlier.

Results for WTP are reported in the last line of table 3, again using an interest rate of 10% and a storage cost of $0.01 per bushel. The economic value of the optimal marketing strategy for a risk neutral farmer ($0.138 per bushel) is greater than the value for a risk averse farmer ($0.042 per bushel) because the risk averse farmer sells some corn at harvest even when he/she is following the optimal (risk averse) storage strategy, while the risk neutral farmer would generally not sell anything at harvest when he/she is following the optimal (risk neutral) strategy. Therefore the risk averse farmer does not get as much value out of being able to apply his/her optimal strategy as the risk neutral farmer does (compared to the case of selling everything at harvest). Of course, the WTP figure for the risk neutral case only estimates what a risk neutral farmer would be willing to pay to use this storage rule. If the farmer was actually risk averse he/she would be willing to pay less to use the optimal risk neutral rule than he/she would be willing to pay to be able to use the optimal risk averse rule.

These dollar per bushel WTP amounts for Michigan corn farmers translate to between 5% and 15% of net farm income, which might be considered economically significant, thus highlighting the importance of a good marketing strategy.

Conclusions

This study extends the risk neutral farm storage model of FL to the case of a risk averse farmer in a discrete time, discrete state and control space framework. Results provide both theoretical and empirical support for the optimality of partial sales over the storage season, as opposed to the simple sell everything or sell nothing strategy derived in FL. This partial sales behavior is more consistent with what we actually observe farmers doing when they make their storage decisions. The optimal distribution of sales over the storage season depends on the degree of farmer risk aversion, as well as storage costs, interest rates, and the underlying probability distribution of prices.

An application of the model to farm storage of corn in Michigan shows that risk averse farmers will sell a proportion of their corn crop early (generally right at harvest) unless harvest prices are extremely low. They do this even though it reduces their expected terminal wealth (since they cannot then gain from the expected rise in prices over the storage season) because it also reduces risk. This result shows that risk aversion is capable of explaining the observed behavior of partial sales over the storage season without assuming that farmers are somehow myopic or failing to optimize.

The economic value of optimal storage rules for corn in Michigan was estimated by computing the farmer's WTP for the optimal marketing rule, as opposed to just selling all corn at harvest. WTP estimates correspond to about 5%-15% of average net farm income to farmers in Michigan, depending on their degree of risk aversion. This might be considered an economically significant amount, thus highlighting the importance of farmers implementing good post-harvest grain marketing strategies.

Footnotes

1. Because sales and current stock levels are expressed as a proportion of the total harvest, wealth levels must be scaled by the total harvest as well. Thus a "wealth" level of 1.351 would indicate total wealth would really be $1.351 times the amount of the initial harvest in bushels. It can be shown that, in the case of a CRRA utility function this transformation has no affect on the optimal decision rule because the CRRA is homogeneous of degree 1-*b* (Varian). Hence, the use of this transformation is innocuous and is only used to make the results easier to interpret.
2. The CRRA utility function is not defined under exact risk neutrality so we set $b = 0.0001$ (near risk neutrality) to approximate the optimal risk-neutral strategy. Henceforth, this case will be simply described as "risk neutral."

References

Baillie, R.T., and R.J. Myers. “Bivariable GARCH Estimation of the Optimal Commodity Futures Hedge.” *J. of App. Econometrics* 6(1991):109-124.

Berg, E. “A Sequential Decision Model to Determine Optimal Farm-Level Grain Marketing Policies.” *Europ. Rev. of Agr. Econ.* 14(1987):91-116.

Bertsekas, D.P. *Dynamic Programming and Stochastic Control*. New York: Academic Press, 1976.

Bollerslev, T. “A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return.” *Rev. of Econ. and Stat.* 69(August 1987):542-47.

Fackler, P.L. “Futures Price Volatility: Modeling Non-Constant Variance.” Paper presented at AAEA Meetings, Reno, Nevada, July 1986.

Fackler, P. L., and M. J. Livingston. “Optimal On-Farm Storage.” *NCR 134, Commodity Price Analysis, Forecasting and Risk Management* (April 1996):176-90.

Lence, S., K. Kimle, and M.L. Hayenga. “A Dynamic Minimum Variance Hedge.” *Amer. J. of Agr. Econ.* 75(November 1993):1103-1171.

Miranda, M.J. and P.L. Fackler. *Applied Computational Economics and Finance*. Forthcoming, MIT Press.

Musser, W.N., G.F. Patrick, and D.T. Eckman. “Risk and Grain Marketing Behavior of Large Scale Farmers.” *Rev. of Agr. Econ.* 18(January 1996):65-77.

Sartwelle J., D. Obrien, W. Tierney Jr., and T. Eggers. “The Effect of Personal and Farm Characteristics Upon Grain Marketing Practices.” *J. Agr. and Apl. Econ.* 32(April 2000): 95-111.

Tronstad, R. and C.R. Taylor. "Dynamically Optimal After-Tax Grain Storage, Cash Grain Sale, and Hedging Strategies." *Amer. J. of Agr. Econ.* 73(February 1991):75-88.

USDA, NASS. *Agricultural Statistics 2001*. US Governmental Printing Office, 2001.

Varian, H.R. *Microeconomic Analysis*. New York: W.W. Norton, 1984.

Wang, H., S.D. Hanson, R.J. Myers, and J.R. Black. "The Effects of Crop Yield Insurance Designs on Farmer Participation and Welfare." *Amer. J. of Agr. Econ.* 80(November 1998):806-820.

Yang, S-R., and B.W. Brorsen. "Nonlinear Dynamics of Daily Cash Prices." *Amer. J. of Agr. Econ.* 74(August 1992):706-715.

Zulauf, C.R. and S.H. Irwin. "Market Efficiency and Marketing to Enhance Income of Crop Producers." *Rev. of Agr. Econ.* 20(Fall/Winter 1998): 308-331.

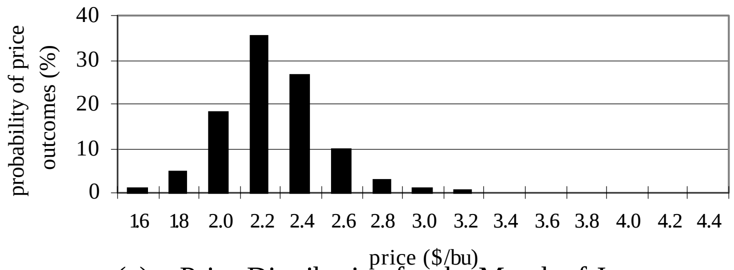

(a) Price Distribution for the Month of January

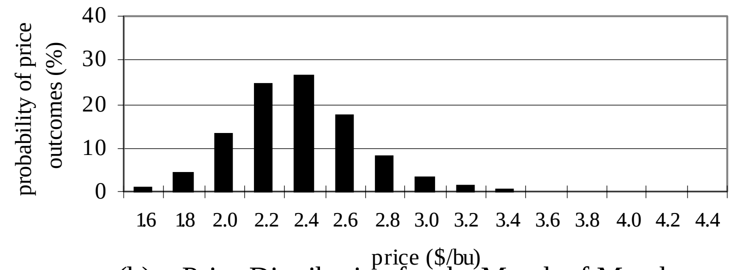

(b) Price Distribution for the Month of March

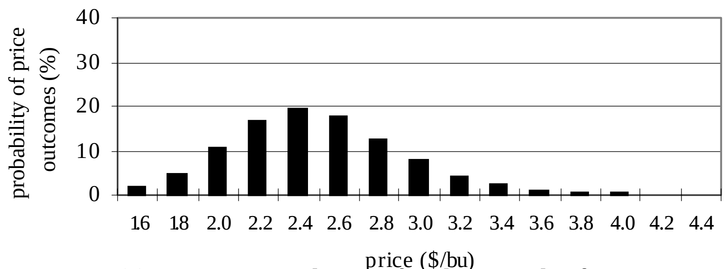

(c) Price Distribution for the Month of May

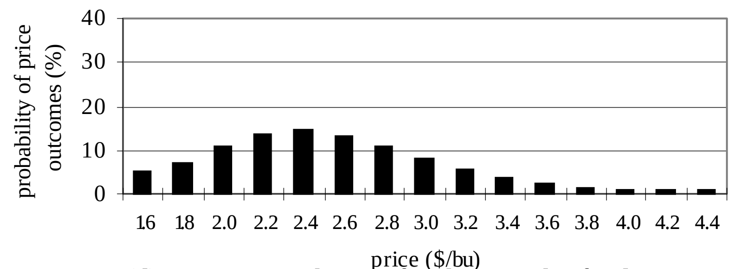

(d) Price Distribution for the Month of July

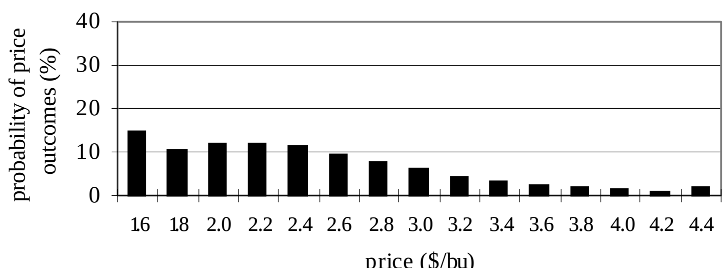

(e) Price Distribution for the Month of September

Figure 1. Future Price Distributions Conditional on a November Price of $2.24

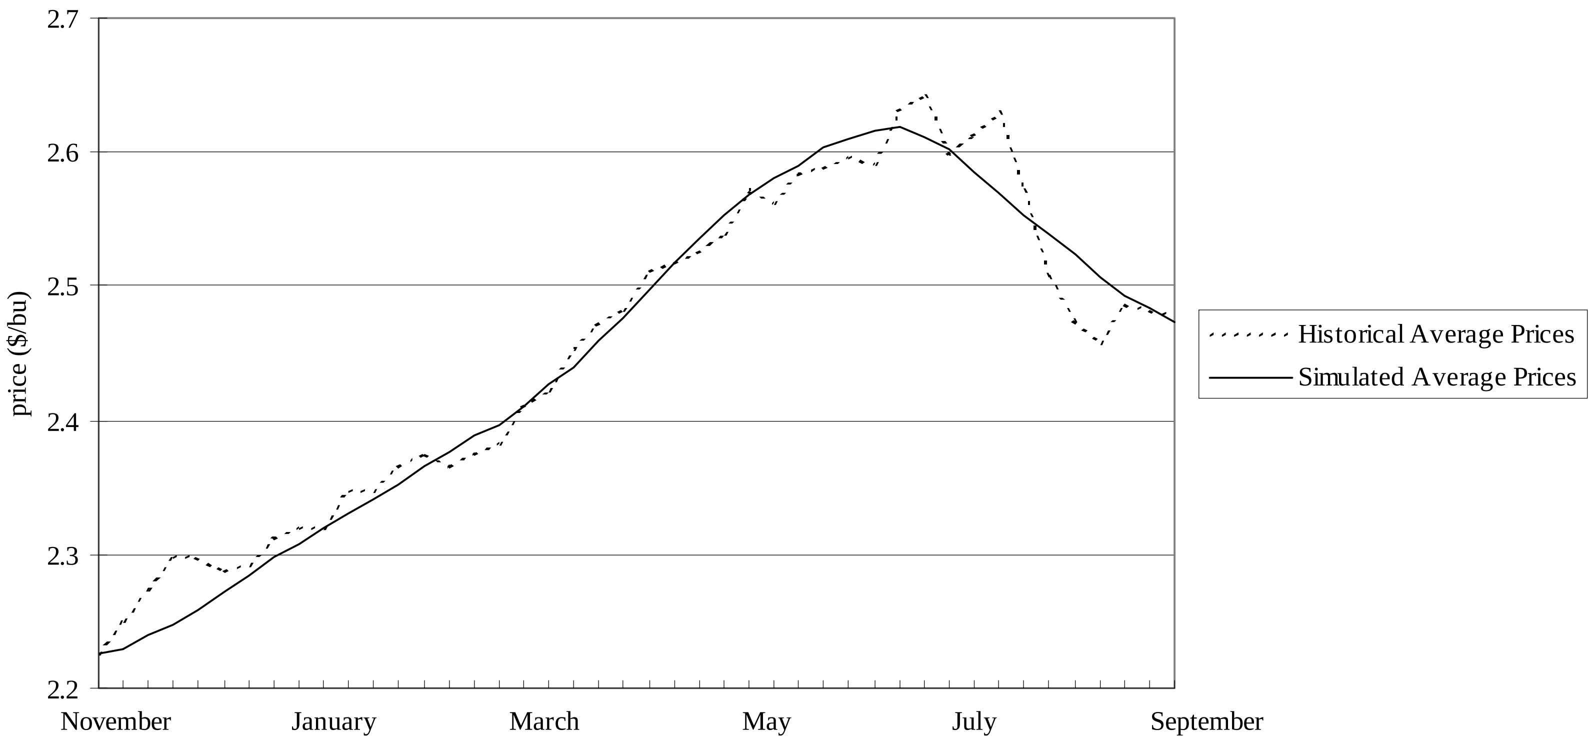

Figure 2. Comparison of Simulated Average and Historical Average Price Movements

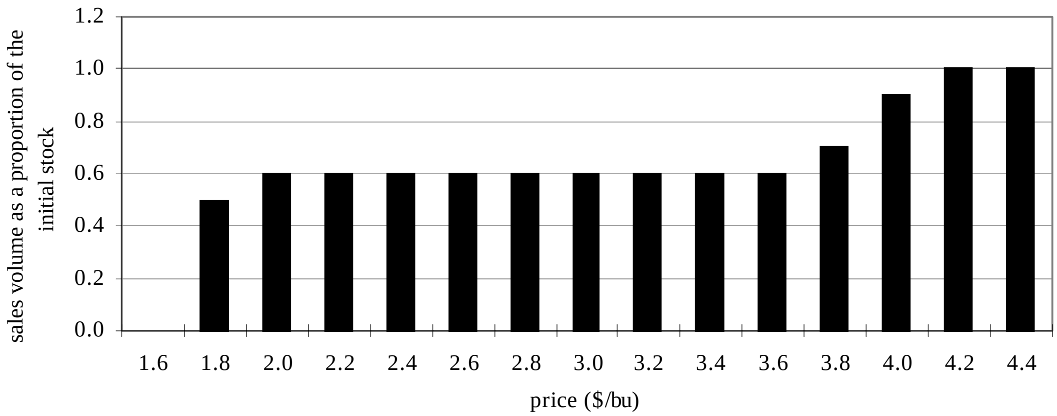

Figure 3. Optimal November Marketing Policy for Risk Averse Farmers

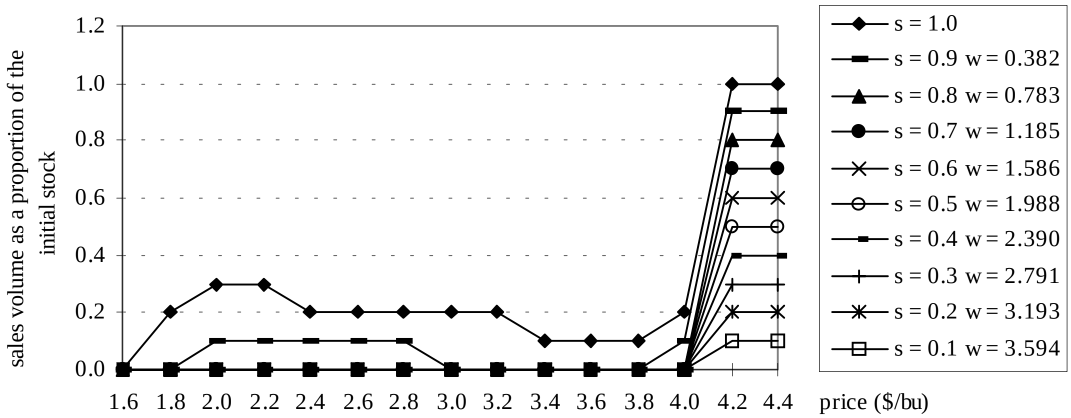

Panel a High Wealth Level

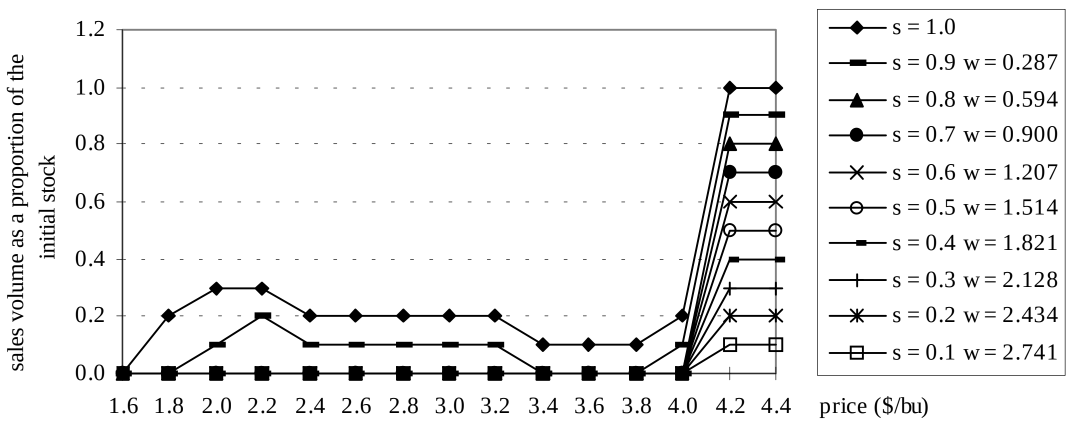

Panel b Middle Wealth Level

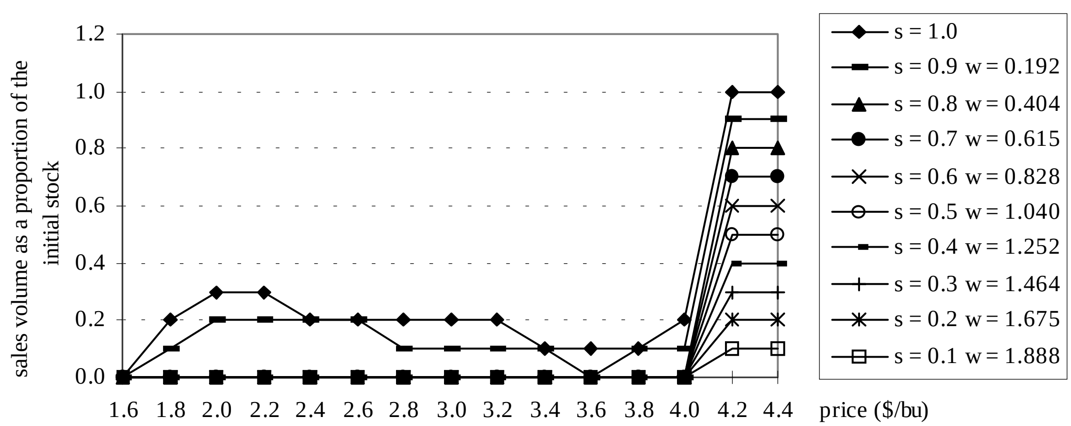

Panel c Low Wealth Level

Figure 4. Optimal January Marketing Policy for Risk Averse Farmers

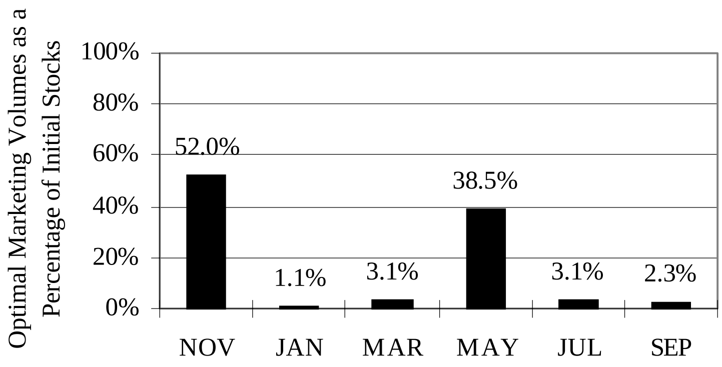

(a) Risk Aversion

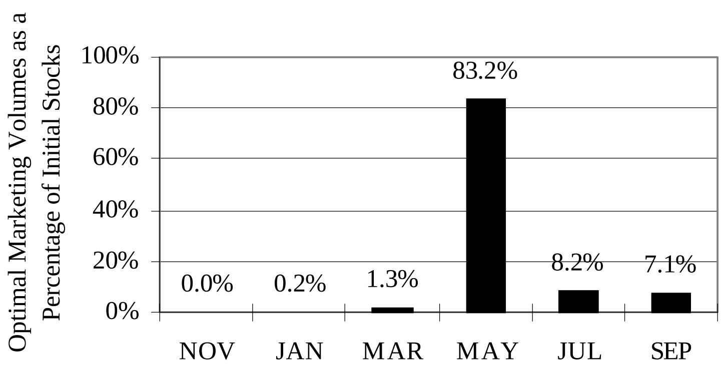

(b) Risk Neutrality

Figure 5. Expected Distribution of Sales Over the Storage Season

Table 1. Corn Price Model Estimates

$$\Delta p_i = \gamma \Delta p_{i-2} + \sum_{j=1}^{2} [d_{1j} \cos(2\pi \frac{k}{52/j}) + d_{2j} \sin(2\pi \frac{k}{52/j})] + \epsilon_i$$

$$\epsilon_i | \Omega_{i-1} \sim t(0, \sigma_i^2, v)$$

$$\sigma_i^2 = \omega + \alpha \epsilon_{i-1}^2 + \beta \sigma_{i-1}^2 + \psi_1 \cos(2\pi k/52) + \psi_2 (\sin(2\pi k/52)$$

Parameter	Estimate	p-Value	Parameter	Estimate	p-Value
γ	0.0506	0.0488	ω	1.6255	0.0078
d_{11}	-0.5776	0.0000	α	0.0934	0.0013
d_{21}	-0.3893	0.0007	β	0.7658	0.0000
d_{12}	-0.0021	0.4929	ψ_1	0.6222	0.0614
d_{22}	0.3211	0.0029	ψ_2	-1.1059	0.0000
			v^{-1}	0.1913	0.0000

Statistics

LR = 74.9968

Ljung-Box Q-statistics

$Q(5)$ = 3.6719	(0.2991)	$Q^2(5)$ = 1.1303	(0.7698)
$Q(15)$ = 13.1055	(0.4397)	$Q^2(15)$ = 12.8787	(0.4572)
$Q(20)$ = 17.9130	(0.4614)	$Q^2(20)$ = 16.3879	(0.5655)

Notes: p_i is the logarithm of cash price at Saginaw, Michigan in week i. Likelihood ratio tests were used to determine the order of the seasonal functions in mean and variance, k is the observation number in the season which corresponds to the current i (k = 1, 2, ... 52), and LR is the likelihood ratio test statistic for GARCH(1,1) with conditional normal errors ($1/v = 0$) against conditional t-distributed errors ($1/v > 0$). $Q(l)$ is a test for l-th order serial correlation in the residuals; and $Q^2(l)$ is a test for l-th order serial correlation in the squared standardized residuals. The values in parentheses following each Q test statistic are the corresponding p-values for that test.

Table 2. Cutoff Prices for Different Storage Costs and Interest Rates for Nearly Risk Neutral Producers

Month	Base Case	Changes in Storage Costs		Changes in Interest Rates	
	r = 10% α = \$0.01	α = \$0.02	α = \$0	r = 15%	r = 5%
November	\$4.1	\$3.9	\$4.1	\$3.9	\$4.1
January	\$4.3	\$4.1	\$4.3	\$4.1	\$4.3
March	\$4.1	\$4.1	\$4.1	\$4.1	\$4.1
May	\$1.9	\$1.7	\$2.9	\$1.9	\$3.1
July	\$1.7	\$1.7	\$1.7	\$1.7	\$1.7
September	0	0	0	0	0

Table 3. Performance of Alternative Storage Rules

Performance Measure	Sell Everything at Harvest	Optimal Storage Under Near Risk Neutrality	Optimal Storage Under Risk Aversion
Expected Wealth/Bushel	\$2.237	\$2.375	\$2.301
Unconditional Variance of Wealth/Bushel	0.494	0.454	0.101
Willingness to Pay/Bushel	\$0	\$0.138	\$0.042

Appendix

Derivation of the Optimality Conditions

The state vector for our problem is defined as $x_t = (w_t,\ s_t,\ p_t,\ z_t)$ and Bellman's equation can be written:

(A.1a) $$v_T(x_T) = U(w_T)$$

(A.1b) $$v_t(x_t) = \max_{q_t} E_t\{v_{t+1}(x_{t+1})\} \quad ; \quad t = 0,\ 1,\ \ldots,\ T-1$$

where the maximization is subject to the transition equations:

(A.2a) $$w_{t+1} = (1 + r)\ [w_t + p_t q_t - \alpha(s_t - q_t)] \quad ; \quad t = 1,\ 2, \ldots,\ T$$

(A.2b) $$s_{t+1} = s_t - q_t \quad ; \quad t = 0,\ 1,\ \ldots,\ T-1$$

and the feasibility constraints:

(A.2c) $$q_t \geq 0 \quad ; \quad t = 0,\ 1,\ \ldots,\ T-1$$

(A.2d) $$q_t \leq s_t \quad ; \quad t = 0,\ 1,\ \ldots,\ T-1$$

We solve the problem by iterating on Bellman's equation.

At the final decision period $t = T-1$ the constraint (A.2d) is clearly binding because, from (A.1a) and (A.2a), stocks have no value in the terminal period T but wealth does. Hence, the solution at $t = T-1$ is

$q_{T-1} = s_{T-1}$. In all prior periods the necessary conditions for solving Bellman's equation subject to the transition equations and feasibility constraints are given by:

(A.3a)
$$E_t \left\{ \frac{\partial v_{t+1}(x_{t+1})}{\partial w_{t+1}} (1 + r)(p_t + \alpha) - \frac{\partial v_{t+1}(x_{t+1})}{\partial s_{t+1}} \right\} - \lambda_t \leq 0$$

(A.3b)
$$q_t \left[E_t \left\{ \frac{\partial v_{t+1}(x_{t+1})}{\partial w_{t+1}} (1 + r)(p_t + \alpha) - \frac{\partial v_{t+1}(x_{t+1})}{\partial s_{t+1}} \right\} - \lambda_t \right] = 0$$

(A.3c)
$$\lambda_t (q_t - s_t) = 0$$

and $q_t \geq 0$, $\lambda_t \geq 0$, and $s_t - q_t \geq 0$. Here, λ_t is the shadow value of relaxing the short selling constraint. It will also be useful to note that from the envelope theorem:

(A.4)
$$\frac{\partial v_t(x_t)}{\partial s_t} = -\alpha (1 + r) \, E_t \left\{ \frac{\partial v_{t+1}(x_{t+1})}{\partial w_{t+1}} \right\} + E_t \left\{ \frac{\partial v_{t+1}(x_{t+1})}{\partial s_{t+1}} \right\}, \text{ and}$$

(A.5)
$$\frac{\partial v_t(x_t)}{\partial w_t} = (1 + r) \, E_t \left\{ \frac{\partial v_{t+1}(x_{t+1})}{\partial w_{t+1}} \right\}$$

Now using (A.4) and (A.5), (A.3) can be expressed

(A.6a)
$$\frac{\partial v_t(x_t)}{\partial w_t} p_t - \frac{\partial v_t(x_t)}{\partial s_t} - \lambda_t \leq 0$$

(A.6b) $$q_t\left[\frac{\partial v_t(x_t)}{\partial w_t} - \frac{\partial v_t(x_t)}{\partial s_t} - \lambda_t\right] = 0$$

(A.6c) $$\lambda_t(q_t - s_t) = 0.$$

These are the conditions given in the text.