Journal of Agricultural and Resource Economics 28(3):558–579

Optimal On-Farm Grain Storage by Risk-Averse Farmers

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

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 on-farm 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 findings show, contrary to the "sell all or nothing" risk-neutral rule, risk-averse farmers will spread sales out over the storage season. As farmers become more risk averse, the optimal strategy is to sell more grain at harvest and spread sales over the storage season, even though this practice reduces expected return. This result is more consistent with observed farmer behavior than the "sell all or nothing" risk-neutral rule.

Key words: grain storage, risk aversion, stochastic dynamic programming

Introduction

On-farm grain storage can have important effects on farm profitability. Thus, it is not surprising many research and extension programs have investigated the optimal timing of sales from storage (e.g., Fackler and Livingston; Ferris; Lence, Kimle, and Hayenga; Tronstad and Taylor; and Zulauf and Irwin). By storing grain at harvest and waiting for prices to rise over the storage season, farmers can sometimes obtain higher total returns, even after accounting for storage costs. Of course, this strategy is risky because prices sometimes fall over the storage season or do not rise enough to cover storage costs. Furthermore, stocks must be sold at some point in time, and determining the best time to sell can be difficult.

In a recent article, Fackler and Livingston solve the optimal on-farm storage problem using stochastic dynamic programming, and apply their model to on-farm storage of soybeans in Illinois. One of the innovations in their analysis is that they model irreversibility by assuming, once sold, on-farm stocks cannot be replenished until the next harvest. Speculative repurchase of grain is ruled out on grounds that on-farm storage is a business marketing activity, and transaction costs from speculative repurchase will discourage such behavior. This appears to be a significant innovation because irreversibility is shown to have an important impact on optimal grain storage decisions, and certainly, very little speculative repurchase of grain is observed among operational grain farms (Sartwelle et al.). Fackler and Livingston's optimal storage strategy under irreversibility is an "all or nothing" rule—sell everything now if the current price is high

---

Jing-Yi Lai is assistant professor, Graduate Institute of Economics, Nan-Hua University, Taiwan. While this research was undertaken, she was a graduate student at Michigan State University. Robert J. Myers and Steven D. Hanson are professors, Department of Agricultural Economics, Michigan State University. The authors thank two anonymous reviewers and Russ Tronstad for helpful comments, and the Michigan Agricultural Experiment Station for its financial support.

Review coordinated by Russell Tronstad.

enough, otherwise wait and sell nothing. The decision to sell is determined by a "cutoff price" that may change over time in response to changing information about the state of the market (and hence expectations about future price levels).

One concern about Fackler and Livingston's "all or nothing" rule is that it seems at odds with the way operating grain farmers actually behave. For example, on-farm corn stocks in Michigan were estimated at 120 million bushels (mb) on January 1, 2002, after the 2001 corn harvest. But stocks had fallen to 80 mb by March 1, 54 mb by June 1, and 16 mb by September 1, 2002 [U.S. Department of Agriculture (USDA)]. Clearly, sales out of on-farm storage are spread out over the storage season, at least at the aggregate state level. Of course, these aggregate data could be reflecting a large number of different individual farmers each making "all or nothing" sell decisions at different points in time. Nevertheless, both the aggregate data and anecdotal evidence suggest many, if not most, farmers with on-farm storage facilities spread sales out over the storage season, rather than selling everything at one time.

The "all or nothing" nature of Fackler and Livingston's optimal decision rule stems from the linearity of their return function and transition equation. Clearly, any model extension which makes the return function strictly concave will result in optimal sales which are spread out over the storage season, at least to some extent. How might this be accomplished? As shown by Tronstad and Taylor, allowing for a nonlinear tax schedule can introduce such a concavity. Other possibilities are allowing for liquidity constraints and/or differential borrowing and lending costs which encourage farmers to spread sales out in order to meet continuous consumption needs. However, an obvious source of return function concavity which has received very little attention in the on-farm storage literature is farmer risk aversion. If farmers are risk averse, they will have an incentive to spread sales out over the storage season to diversify the risk of selling everything at one time. Hence, incorporating farmer risk aversion into the optimal on-farm storage model appears to have the potential to explain partial sales at different times during the storage season, as is often observed among operational grain farms.

Based on a search of the literature, the only published study of optimal on-farm storage accounting explicitly for farmer risk aversion is an analysis by Berg, who examined on-farm wheat storage in the European Union (EU). Berg did indeed find that spreading sales out over the storage season can be optimal for risk-averse EU wheat farmers. But while this study provides useful insights, the model is quite restrictive. In particular, Berg assumed triangular price distributions and that the probability distribution for price in any one period is independent of price outcomes in previous periods. These appear to be very restrictive assumptions given what we know currently about the probability structure of most grain price movements (e.g., Yang and Brorsen; Baillie and Myers; Wang et al.).

The objective of this study is to show how the optimal on-farm grain storage problem can be solved assuming risk-averse farmers and incorporating more realistic price probability distributions. Stochastic dynamic programming in a discrete-time framework is used to derive optimal storage rules under risk aversion. We continue to use Fackler and Livingston's irreversibility assumption, but extend their risk-neutral analysis by allowing for farmer risk aversion. In our model, an optimal partial selling rule is derived, as opposed to the "all or nothing" rule obtained by Fackler and Livingston. 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 on-farm storage of corn in Michigan. Results provide empirical support for spreading sales out over the storage season as a diversification strategy. Sensitivity analysis is conducted to investigate how various factors influence the optimal timing of sales out of storage. Finally, the performance of the optimal storage rules is evaluated by comparing the mean and standard deviation of storage returns under a range of alternative storage rules.

The Theoretical Model

Consider a risk-averse farmer with on-farm storage facilities deciding when to sell grain over the storage season. The storage season begins at the current harvest, ends before next year's harvest, and is divided into T equal-spaced decision nodes. At the beginning of each decision node, the farmer has a current grain stock (s_t), observes the current market price (p_t), and chooses an amount of the commodity (q_t) to sell in the spot market. Following Fackler and Livingston, we impose the restriction $0 \le q_t \le s_t$, thereby ensuring sales must be nonnegative (grain cannot be repurchased for speculative purposes) and less than or equal to the current storage level (no short selling allowed). The behaviors ruled out by these restrictions would be speculative and involve additional transaction costs which discourage most farmers from engaging in them.[1]

The farmer pays a one-time per unit cost, $c > 0$, to move grain into or out of storage, and for each period grain is carried over in storage there is an additional cost of $k > 0$ per unit. Because farmers are risk averse, they may try to manage storage risk by hedging stocks on the futures market. Let f_t be the current futures price for grain deliverable at the next decision node, and b_t be sales of futures contracts at t which are held until they mature at the next decision node, $t+1$.[2] Consistent with previous assumptions about nonspeculation with cash positions, we rule out speculation on futures by imposing the restriction $b_t \ge 0$. Because futures contracts are marked to market at the next decision node, any profits or losses from futures trading are realized at that time. There is a per unit transaction cost from trading futures (round-trip cost paid when the futures position is taken out) of δ. This also includes costs of attaining and processing the information required to trade futures.

With these assumptions, compounded storage returns (π_t) evolve over the storage season according to the following transition equations:

(1a) $$\pi_1 = (1+r)\left[p_0 q_0 - (c+k)(s_0 - q_0) - \delta b_0\right] + (f_0 - f_1)b_0$$

and

(1b) $$\pi_{t+1} = (1+r)\left[\pi_t + (p_t - c)q_t - k(s_t - q_t) - \delta b_t\right] + (f_t - f_{t+1})b_t,$$

$$\text{for } t = 1, \ldots, T-1,$$

[1] These transaction costs include not only transportation, brokerage, etc., but also the costs of collecting sufficient information and marketing expertise to be able to speculate with grain repurchase or short selling.

[2] Farmers will generally have a choice regarding which futures contract (maturity date) they use for hedging. Here they are assumed to use a nearby future (maturing at the next decision node) because it will generally have higher liquidity than distant maturities. Specifically, when the next decision node is reached, if some grain is still going to be left in storage, then continued hedging requires the futures positions to be rolled over into the next maturing contract. Therefore, futures positions are only held for one decision period and then liquidated, but if storage continues, new positions can be taken out in the next maturing contract. Generalizing the model to allow the farmer to sell futures contracts over a range of future maturities does not alter the main results which follow, but would make the analysis more complicated.

where s_0 is total grain available for initial storage or sale (i.e., total production harvested at $t = 0$), and r is the (assumed constant) interest rate. Storage is subject to the transition equation:

$$s_{t+1} = s_t - q_t. \tag{2}$$

At each period $t = 0, 1, ..., T-1$, the farmer chooses an amount to sell (q_t) and an amount to hedge (b_t) to maximize the expected utility of final compounded return at the end of the storage season:

$$\max_{\{q_t\}_{t=0}^{T-1}, \{b_t\}_{t=0}^{T-1}} E_0 U(\pi_T), \tag{3}$$

subject to the transition equations (1) and (2), the constraints $0 \le q_t \le s_t$ and $b_t \ge 0$, and a Markov probability process for prices $(\tilde{p}_{t+1}, \tilde{f}_{t+1}) \sim g_t(p_{t+1}, f_{t+1} | p_t, f_t)$, where g_t is the (possibly time-varying) probability density function for cash and futures prices at $t+1$ conditional on current cash and futures prices at t. The $U(\cdot)$ function is an increasing and concave von Neumann-Morgenstern utility function representing farmer risk preferences.[3]

The storage problem can be solved using discrete-time stochastic dynamic programming (Bertsekas; Miranda and Fackler). Defining the state vector $\mathbf{x}_t = (\pi_t, s_t, p_t, f_t)$ and value function $v_t(\mathbf{x}_t)$, then Bellman's equation for the problem is given by:

$$v_T(\mathbf{x}_T) = U(\pi_T) \tag{4a}$$

and

$$v_t(\mathbf{x}_t) = \max_{q_t, b_t} E_t\{v_{t+1}(\mathbf{x}_{t+1})\}, \quad \text{for } t = 0, 1, ..., T-1, \tag{4b}$$

subject to the transition equations (1) and (2) and the constraints $0 \le q_t \le s_t$ and $b_t \ge 0$. Second-order conditions for a maximum are satisfied by the concavity of U.

This problem has no closed-form solution, even for simple assumptions on the form of the utility function and probability distributions. However, some interesting insights can still be gained about the form of the solution by examining first-order necessary conditions for an optimum. Although the storage and hedging decisions are clearly interrelated, the necessary conditions for futures hedging are examined first:

$$E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}\left[f_t - f_{t+1} - \delta(1+r)\right]\right\} \le 0 \tag{5a}$$

and

$$b_t E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}\left[f_t - f_{t+1} - \delta(1+r)\right]\right\} = 0. \tag{5b}$$

[3] Defining utility to depend on grain storage returns only imposes an implicit assumption that grain storage decisions are separable from other production and marketing decisions undertaken by the farmer. While this assumption is questionable, there is a long history of adopting it in applied risk analysis (see Newbery and Stiglitz). The separability assumption can be made less restrictive by examining storage behavior under a range of alternative risk aversion levels, as is done in the numerical analysis which follows.

If the futures market is unbiased, $f_t = E_t(f_{t+1})$, and futures transaction costs δ are high enough, then no positive value of b_t will satisfy (5a) with strict equality. In this case, it is optimal not to use futures ($b_t = 0$). Furthermore, even when futures transaction costs are low enough to encourage futures selling, it will never be possible to fully hedge all future return risk from the storage operation using futures contracts alone, as long as there is basis risk (see Benninga, Eldor, and Zilcha; Lence; and Myers and Hanson). Hence, even when an optimal hedging strategy is followed, there will still be residual return risk facing the farmer if storage is undertaken. Consequently, even when farmers are hedging using futures, residual basis risk will still influence the optimal storage decision.

Next, the necessary conditions are examined for optimal choice of sales (storage) at each t. As shown in the appendix, these necessary conditions can be written as follows:[4]

(6a)
$$\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t}(p_t - c) - \frac{\partial v_t(\mathbf{x}_t)}{\partial s_t} - \lambda_t \le 0,$$

(6b)
$$q_t\left[\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t}(p_t - c) - \frac{\partial v_t(\mathbf{x}_t)}{\partial s_t} - \lambda_t\right] = 0,$$

and

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

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

First, suppose 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 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 (6b) and (6c), we have:

(7)
$$\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t}(p_t - c) \le \frac{\partial v_t(\mathbf{x}_t)}{\partial s_t}.$$

In this case, the marginal value (in terms of expected utility of compounded storage returns) of selling the first bushel of grain out of storage must be less than or equal to the marginal value (again in terms of expected utility of compounded storage returns) of keeping that bushel of grain in storage. If (7) is satisfied with strict inequality, the farmer would actually like to buy more grain to store but is precluded from doing so by the constraint $q_t \ge 0$.

[4] Note, at $t = 0$, the marginal return from current period sales would be p_t instead of $(p_t - c)$ because grain sold right away does not go into storage. Thus, there is no charge for taking it back out of storage.

Third, suppose 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 (6b) and (6c) imply:

$$\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t}(p_t - c) \geq \frac{\partial v_t(\mathbf{x}_t)}{\partial s_t}. \tag{8}$$

In this case, the marginal value (again in terms of expected utility of compounded storage returns) 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 (8) is satisfied with strict inequality, the farmer would like to sell more, 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 (6b) and (6c) imply:

$$\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t}(p_t - c) = \frac{\partial v_t(\mathbf{x}_t)}{\partial s_t}. \tag{9}$$

In this final case, the farmer is indifferent between selling or keeping the marginal bushel of grain. The concavity of $U(\cdot)$, together with the previous result that futures hedging will not eliminate all storage risk in the presence of transaction costs and basis risk, ensures $v_t(\mathbf{x}_t)$ is concave in π_t and s_t. This 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 interior solution (case four) which distinguishes this model from the one in Fackler and Livingston. The linearity of the (risk-neutral) Fackler and Livingston 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 rise high enough, everything will be sold, and if current prices drop 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, even if some of the risk can be hedged on futures markets. While this is an interesting theoretical possibility, it is not yet clear how empirically relevant the possibility of partial sales really is. Can we find a reasonable problem and realistic set of parameter values for which significant partial sales are part of the optimal strategy? In the next section, this question is addressed using an empirical application to on-farm corn storage in Michigan.

An Application

While there is no closed-form solution to the theoretical model developed in the previous section, it can be solved numerically if assumptions are made about risk preferences and the model parameters are estimated. In this section, we solve the model numerically for the optimal on-farm storage problem faced by a farmer using the Saginaw market in Michigan to market corn. The 1997 corn crop year is used to calibrate the analysis. We attempt to ensure the structure of the problem is as realistic as possible. However, we emphasize that the main goal of the application is not to provide specific extension advice to farmers. Rather, the objectives are to illustrate how a numerical analysis of the model can be undertaken, and to demonstrate that partial sales can indeed be optimal for a realistic on-farm storage problem with risk-averse farmers and a reasonable set of parameter values.

Michigan corn farmers make their first marketing decision at harvest (defined here as the first week of November) when they decide how much production to sell immediately and how much to store to 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 or her corn. Any storage left over at the first week of September is automatically sold to make way for the coming new crop.

To keep the problem simple, transaction costs for trading futures are assumed to be high enough that farmers choose not to hedge on futures markets. This assumption is adopted for two reasons. First, it is consistent with the fact that few farmers make extensive use of futures hedging (Musser, Patrick, and Eckman; Sartwelle et al.; Asplund, Forster, and Stout). Second, by excluding futures hedging from the numerical example, we focus on storage decisions rather than hedging decisions, and provide the maximum incentive for farmers to spread sales out over the storage season. This approach seems reasonable, given the main objective of the numerical example is to show that partial sales will occur for at least some farmers under some conditions.

It is also emphasized that while allowing for futures hedging may reduce the incentive to spread sales out over the storage season, it will not eliminate this incentive entirely (because of residual basis risk, as explained earlier). Indeed, one way to observe how the introduction of futures hedging (a reduction in risk) might influence the storage choice would be to investigate what happens when the farmer becomes less risk averse (a reduction in risk aversion). In both cases, the storage decision rule should move in the same direction (toward the "all or nothing" risk-neutral rule and away from the partial sales risk-averse rule). Hence, by examining how optimal storage rules change when the farmer's degree of risk aversion is reduced (but always assuming no futures trading), it is possible to assess the direction of the effects on the optimal storage rule of allowing futures hedging (lowering of futures transaction costs). The advantage of this approach is that it allows us to meet the main objectives of the numerical application without overly complicating the state space for the analysis.

Although the control and state spaces for the problem might best be viewed as continuous, a discrete approximation is used 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 one of 11 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 11 possible proportions. The state space for selling prices p_t was specified as 15 possible price states ranging from \$1.60 to \$4.40 per bushel in 20¢ 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 \$1.70 or below.

The state space for compounded storage returns could be approximated similarly to the storage and price state spaces. However, this would require value function interpolation because for a given current storage level, price level, and compounded storage return, and a given current sales choice, there would be no guarantee next period's compounded storage return would lie in any discrete state space that might be chosen for storage returns. To avoid the additional approximation error resulting from value function interpolation, the state space for the compounded storage return was defined to be time dependent based on all possible feasible values the storage returns could take over the time horizon of the problem, given the state and control spaces for price, storage, and sales as defined previously.[5] This procedure increases computation time, because of the large number of compounded storage return states which must be evaluated, but does eliminate approximation error that would otherwise be caused by value function interpolation.

Three alternative values were used for the annual interest rate and monthly storage costs, $r = 5\%$, 10%, and 15%, and $k = 0¢$, $1¢$, and $2¢$ per bushel. Interest rates will vary for individual farmers with differing opportunity costs of capital. Hence, it seems sensible to investigate the sensitivity of results over a range of interest rates. Similarly, storage costs may vary for different farmers because drying and handling costs will depend on the farm's capital structure.[6] However, it is important to note that variable on-farm storage costs for corn in Michigan are very small, consisting only of some electricity to run fans, and perhaps some minor fuel, labor, and maintenance costs. Hence, variable costs of on-farm storage are estimated to run between 0¢ and 1¢ per bushel per month (Ferris). The 2¢ cost rate is included to show the sensitivity of results to an extreme upper bound.

The remaining cost parameter c, which represents the cost of moving corn into and out of on-farm storage bins, was set to zero. While this cost may be important in some applications, it is not a major additional cost of on-farm corn storage in Michigan—i.e., most Michigan farmers who own storage bins move the harvest from their fields into the bins, even if they are planning on selling immediately. This is done to collect the corn for transport to the elevator by semi-trailer (Hilker).

Thus, there are no significant additional charges from moving corn into and out of on-farm storage in Michigan, other than those expected to be accrued during the normal harvesting process. Furthermore, many elevators have an in-charge at harvest time when the demand for their services is high, and then drop the charge later in the storage season as throughput falls. Clearly, this additional in-charge for selling immediately at harvest can offset any additional cost from moving grain into storage at harvest and out again later in the storage season.

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. A constant relative risk aversion (CRRA) utility function was assumed: $U(\pi_T) = \pi_T^{1-R}/(1-R)$, $R > 0$.[7] The parameter R denotes the coefficient of relative risk aversion and is set to one

---

[5] The state space for compounded storage returns is computed starting at the harvest period when there is a single state of zero (returns from storage must be zero before any sales take place). To compute the state space at subsequent decision nodes, we chose an interest rate r, and cost of storage parameters c and k, and took every feasible time path in the state space for storage decisions and applied every feasible price path sequence in the state space for price, to obtain every possible path for the compounded storage return π_t at every period over the time horizon of the problem. All of these possible values were then used at each decision node as the state space for that node. The resulting state space for compounded storage returns is time dependent, and gets very large toward the end of the storage season.

[6] Note, variable costs only are counted because we are looking at the optimal storage decision for given (i.e., sunk) storage facilities, not analyzing new investment in storage facilities.

[7] Because sales and current stock levels are expressed as a proportion of the total harvest, compounded storage returns must be scaled by the total harvest as well. Thus a storage return level of 1.351 would indicate total compounded returns at that date would really be \$1.351 times the amount of the initial harvest in bushels. It can be shown, in the case of a CRRA utility function, this transformation has no effect on the optimal decision rule because the CRRA is homogeneous of degree $1 - R$ (Varian). Hence, the use of this transformation is innocuous and is only used to make the results easier to interpret.

Table 1. Corn Price Model Estimates

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

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

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

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:

Likelihood Ratio (LR) = 74.9968

Ljung-Box Q-Statistics (p-values in parentheses):

$Q(5)$ = 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: $\ln(p_i)$ is the logarithm of cash price at Saginaw, Michigan, in week i. The constant was excluded from the conditional mean equation because its estimated value was very close to zero and not statistically different from zero. Likelihood-ratio tests were used to determine the order of the seasonal functions in mean and variance; n is the observation number in the season which corresponds to the current i (n = 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 lth-order serial correlation in the residuals; and $Q^2(l)$ is a test for lth-order serial correlation in the squared standardized residuals.

of seven possible values, R = {0.0001, 0.5, 1.0, 2.0, 3.0, 5.0, 10.0}, in order to compare results across a range of farmer risk aversion.[8] Computation of the price transition probabilities required a 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 transition probability matrices which map the stochastic price states across the decision nodes. 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 underlying price process is specified as an autoregressive seasonal model for the conditional mean of changes in log prices, 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).

[8] Because the CRRA utility function is not defined under exact risk neutrality, we set R = 0.0001 (near risk neutrality) to approximate the optimal risk-neutral strategy. Henceforth, this case will be simply described as "risk neutral." Similarly, the CRRA utility function converges to logarithmic utility if R = 1.0, so logarithmic utility was used in this case.

After investigating several alternative models for goodness of fit, a preferred model was chosen and estimated (see table 1).[9] The logarithm of weekly Saginaw corn prices was found to be nonstationary, and the conditional mean and variance of the change in log prices appear to vary over time as a result of both stochastic and seasonal factors. The Ljung-Box Q-statistics show the chosen model is well specified in terms of removing autocorrelation from both the errors and squared standardized errors of the change in log prices.

The weekly econometric model can be used to generate transition probabilities corresponding to the 15 possible price states in the price state space. The transition probabilities are allowed to be different at each decision node, t, and so five different sets of transition probability matrices are computed. The transition probabilities are generated from the estimated weekly price model (table 1) using simulation.

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 November price state ($1.60) is selected and its natural logarithm is taken to give the first possible value of the log price state in November. Next, one realization of the change in log price from November to January is simulated by taking the weekly econometric model estimated in table 1 and making a sequence of eight or nine (whichever is appropriate) random draws on ε_i values, making certain the dynamics of both the log price and the variance of the ε_i are tracked properly over the two-month period.[10] Finally, we take the simulated change in log prices from November to January, add this to the initial log price in November, and exponentiate to give *one* possible realization of the price level in the first week of January, conditional on the price in the first week of November being $1.60.

This process is repeated 10,000 times using a random number generator for the ε_i, and a check is made of 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 possible alternative price states that could occur in the first week of January.

To generate the entire matrix of transition probabilities, this process was repeated 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, this entire procedure was repeated for transitioning between all decision nodes (November–January, January–March, March–May, May–July, July–September). This process

[9] Initial starting values for the estimation were found by applying ordinary least squares (OLS) to the conditional mean part of the model and then examining the autocorrelation function and partial autocorrelation function of the squared OLS residuals to obtain reasonable starting values for the GARCH parameters. We also investigated a range of alternative starting values and spot checked the likelihood value over a wide grid of alternative parameter values to ensure the maximum-likelihood estimates were a global maximum.

[10] To track these dynamics properly, it is necessary to specify starting values for $\Delta\log(p_{i-2})$, ε^2_{i-1}, and σ^2_{i-1} to initialize the simulation, and then compute the dynamic path of the change in log prices and its conditional variance recursively (using the estimated model from table 1 and beginning from these starting values). Also, care was taken to update the conditional variance at each step before simulating the next ε_i value. In the absence of any other conditioning information, we set the initial value of $\Delta\log(p_{i-2})$ to zero, and the initial values of ε^2_{i-1} and σ^2_{i-1} equal to their unconditional expectation (the estimated unconditional variance of ε_i). Any sensitivity of the results to these starting values will diminish as simulation extends further into the future.

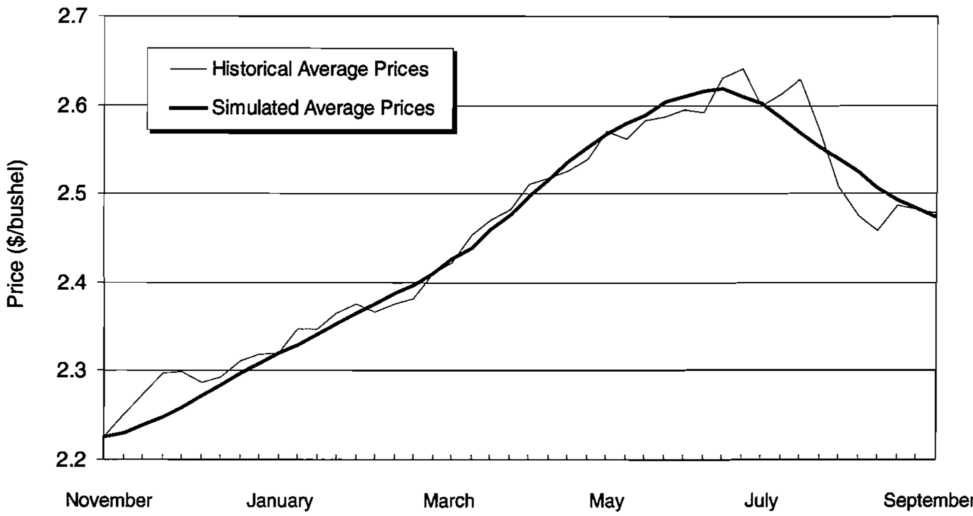

Figure 1. Comparison of simulated average and historical average price movements

gives five different transition probability matrices which can be input into the dynamic programming algorithm.[11]

We conducted a number of experiments to validate the econometric model and estimated transition probabilities. In one of these experiments the initial price in November was set equal to its historical data mean of $2.24 and this price was simulated forward over the entire storage season using the econometric model and 10,000 replications. We then computed the average simulated value (across the 10,000 replications) at each week over the storage season and compared this to the actual historical average price at each week over the season. Results are graphed in figure 1. Average historical prices clearly rise through the storage season until about July, when prices begin to fall in anticipation of the coming harvest. This is the expected pattern, 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 result suggests the econometric model performs well in capturing the seasonal movement in Saginaw corn prices.

The pattern of volatility movements predicted by the econometric model was examined by forecasting the seasonal component of the conditional variance of changes in log prices. This seasonal component is graphed in figure 2 and shows corn price volatility is at a minimum during the winter months of December through February, then rises during spring planting and the growing season. Volatility reaches a maximum at the end of summer in July and August, and finally falls again after harvest and into the

[11] As noted by an anonymous reviewer, the estimated econometric model is not entirely consistent with the discrete Markov probability structure assumed for the transitions between price states in the dynamic programming (DP) model. The econometric model is based on weekly price observations, while the decision nodes in the DP model are each two months apart. Furthermore, because the econometric model is in first differences and has a second-order lag, the econometric model is not first-order Markov (current price is not a sufficient statistic for predicting next week's price). We do not view these inconsistencies as a serious problem. The aim in our analysis is to take the preferred econometric model of weekly prices (which happens not to be first-order Markov) and use this to obtain a preferred estimate of a first-order Markov approximation for bi-monthly price movements which is required to implement the DP algorithm. This is exactly what the simulation procedure accomplishes.

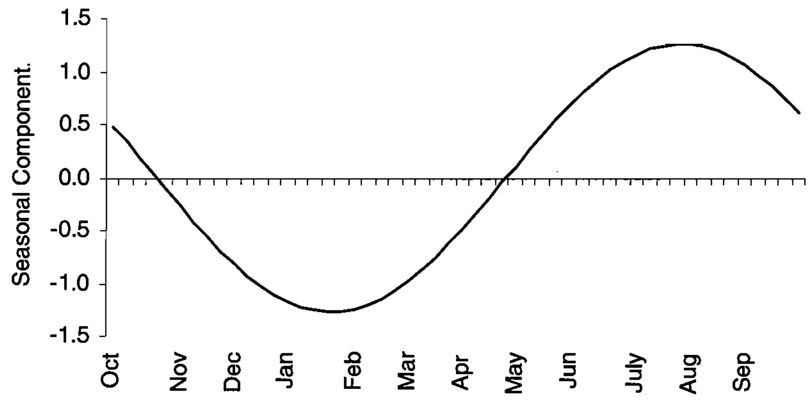

Figure 2. Seasonal component of conditional volatility

winter months. This is exactly the pattern we would expect because price movements (in either direction) should tend to be greater during the growing season when even mild changes in weather conditions can have a big impact on expectations regarding the coming harvest (and hence on current corn prices).

In a final validation check, we evaluated several price probability distributions implied by our estimated transition probability matrices. There are many of these implied distributions depending on the initial starting price. However, the information in the transition probability matrices can be summarized by computing the unconditional distribution of price outcomes at each of the decision nodes. These unconditional distributions were computed by first obtaining the relative frequencies with which historical November prices fell into each of the discrete price intervals over the sample period from 1975 to 1996. The resulting vector of relative frequencies, $\mathbf{P}_0$, was used as the estimate of the (discrete) unconditional marginal distribution of November prices.[12] The unconditional distribution of prices at each future decision node was then calculated by computing the product $\{\mathbf{P}_0 \mathbf{P}_1 \mathbf{P}_2, ..., \mathbf{P}_j\}$ for decision nodes $\{j = 1, 2, 3, 4, 5\}$ using transition probability matrices $\{\mathbf{P}_1, \mathbf{P}_2, ..., \mathbf{P}_j\}$, which give the probabilities of transitioning from prices in decision node $j - 1$ to prices in decision node j.

The resulting unconditional discrete probability distributions are shown in figure 3 for each decision node. There are several interesting features of these unconditional probability distribution estimates. First, each distribution is truncated from below with significant probability massed at the $1.60 minimum price level. This observation is consistent with the fact that the loan rate supported market prices at a minimum price level over most of the sample period used for the econometric model. Second, although it is difficult to see from the graphs, the distributions show a rise in mean prices over the storage season, until September when mean price falls slightly. Furthermore, the distributions become more spread out as the storage season progresses (variance increases). Both of these features are consistent with the historical data.

Finally, as the storage season progresses, more probability gets massed at the upper-bound price level of $4.40 (figure 3). This occurs because the discrete approximation masses extra transition probability at the upper bound of the discrete price space (rather

[12] The raw historical frequencies were smoothed using a four-span Hanning linear smoother to smooth out "bumps" in the historical frequencies.

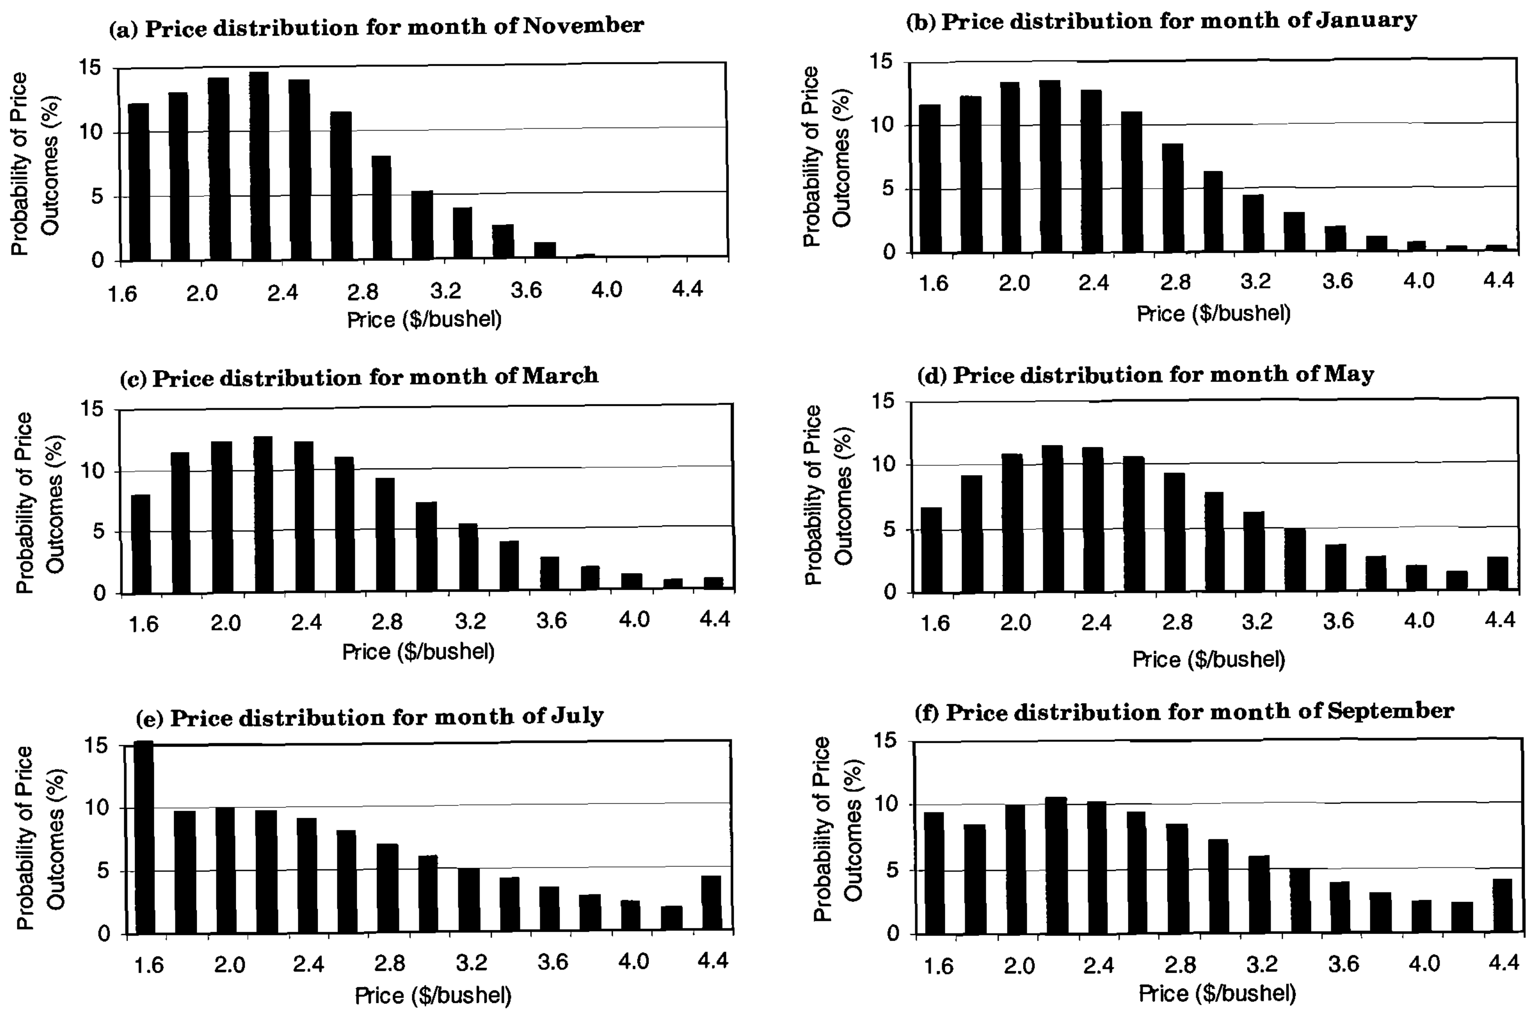

Figure 3. Estimated discrete unconditional price distributions for each decision node

than spreading it over higher prices above the upper bound). Hence, when the current price state is relatively high, the conditional mean of prices next period will be understated by the estimated transition probabilities. Any approximation error introduced by this truncation effect will have little impact on optimal storage decisions at low or intermediate current price states (because the conditional mean of future prices will be estimated well in these cases), though it may have more of an impact on optimal storage decisions at high current price states (because the conditional mean of future prices will be understated in this case). This problem is intrinsic to all discrete approximations of continuous probability spaces, and any attempt to reduce the truncation effect must come at the expense of increasing the dimension of the state space, thus making the problem more difficult to solve.[13]

The Dynamic Programming Algorithm

A discrete time, discrete state, and control space, dynamic programming algorithm based on the approach of Miranda and Fackler is used to solve the model. Programmed and solved in GAUSS, the algorithm proceeds by evaluating the value function at every (discrete) point in the state and control space and takes several hours to solve on a personal computer.[14]

Results

We first solved the model for the case of a risk-neutral farmer ($R = 0.0001$). This leads to the same "sell all or nothing" marketing strategy obtained by Fackler and Livingston (as expected). Hence, in this case, the optimal marketing strategy is defined by a cutoff price for each decision node. If the current price is at or 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 are presented in table 2. Using the base case (the first numeric column) as an example, for the month of November the cutoff price is \$4.10 per bushel. Thus, 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 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 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 cutoff price declines as storage cost

[13] Another way of thinking about the discrete approximation problem is that the estimated discrete transition probability matrices impose an implicit reflecting barrier at the lower and upper price bounds of the discrete state space (prices bounce along or off the barrier instead of breaking through; see Dixit and Pindyck, p. 83). This leads to approximation error if the true underlying price distribution does not contain reflecting barriers (as is implicitly assumed in the econometric model). However, if the true underlying price distribution does have reflecting barriers (say, at a lower bound because of government support programs and at an upper bound because of freedom of entry into a profitable market), then using estimated transition matrices which impose the reflecting barriers actually may provide a better estimate of the underlying price probability distribution than an estimate that is barrier free.

[14] The reason for the long computation time is the large dimensionality of the state space, particularly the state space for compounded storage returns at decision nodes well into the storage season. As discussed previously, the state space for compounded storage returns was defined by evaluating every possible return state that could be achieved under every possible price and storage outcome. This turns out to be a very large number. The tradeoff is that we obtain a more accurate solution, because value function interpolation is not required.

Table 2. Cutoff Prices Under Risk Neutrality for Different Storage Costs and Interest Rates

	Base Case	Changes in Storage Costs		Changes in Interest Rates	
Month	r = 10%, k = \$0.01	k = \$0.02	k = \$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

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 in interest income represents an increased opportunity cost of holding the grain, which makes storage less attractive.

Moving to the case of risk aversion, we initially set relative risk aversion at $R = 5$ to examine how the optimal storage strategy changes under risk aversion. For $R = 5$, the optimal marketing strategy for the harvest period (first week of November) is represented graphically in figure 4 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 compounded return from sales state is zero (because nothing has been sold yet). Therefore, the optimal marketing strategy depends only on the current price. As shown from the results, 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 farmer in January is more difficult to represent because the state space has many more dimensions (the possible storage and compounded returns from storage states are no longer unique). In particular, instead of one storage state of $s = 1.0$ (as in November), there are now 11 possible January storage states, $s \in \{0, 0.1, ..., 1.0\}$, depending on how much corn was sold in November. And instead of one possible compounded storage return state, there are many, depending on the amount of November sales and the price at which previous sales in November took place. To demonstrate the optimal storage decision using two-dimensional diagrams, we constructed three separate graphs, one for each of three possible current compounded storage return states—a high compounded return state (corn sold at harvest in November was sold at a relatively high price), an intermediate current compounded return state (corn sold at harvest in November was sold at an intermediate price), and a low compounded return state (corn sold at harvest in November was sold

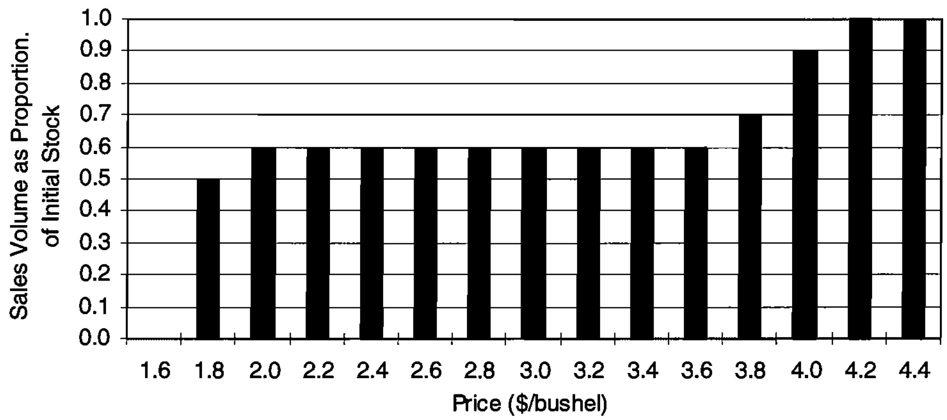

Figure 4. Optimal November marketing policy for risk-averse farmers

at a relatively low price). Then for each of these possible compounded storage return states, we computed how the optimal storage decision changes with changes in current (January) price and storage levels (see figure 5). Each line in each panel of figure 5 reports optimal January sales as a function of the observed January price, given a particular level of current storage and a particular level of the incoming compounded storage return.

Optimal strategies for January are more complex and less intuitive than for November. The simplest result (not included in the graphs) is that if $s = 0$, the (trivial) optimal strategy is to sell nothing because all of the stock has already been sold. Other results for January are consistent in the sense that very low prices lead to no additional sales, while very high prices lead the farmer to sell everything currently available. At intermediate prices, there may be partial sales. Notice, however, 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 5 over some regions of the price space). This is because the current price state not only represents the benefits of immediate sale, but is also a signal indicating the probability of prices going higher in the future. Therefore, a higher current price encourages more sales now 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.

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. When farmers are relatively risk averse, these results suggest partial sales may be optimal at each decision node (except the last), and the price required to encourage current sales falls with movement through the storage season (as expected).

One way to summarize the optimal storage results is to use the estimated price probabilities and optimal storage rules to compute the expected frequency of optimal sales occurring in each month under risk neutrality and under various degrees of risk aversion, assuming the optimal storage rule for that degree of risk aversion is being followed. These results are summarized in table 3 assuming $r = 10\%$, storage costs of \$0.01 per bushel per month, and R = 0.0001, 0.5, 1.0, 2.0, 3.0, 5.0, and 10.0.

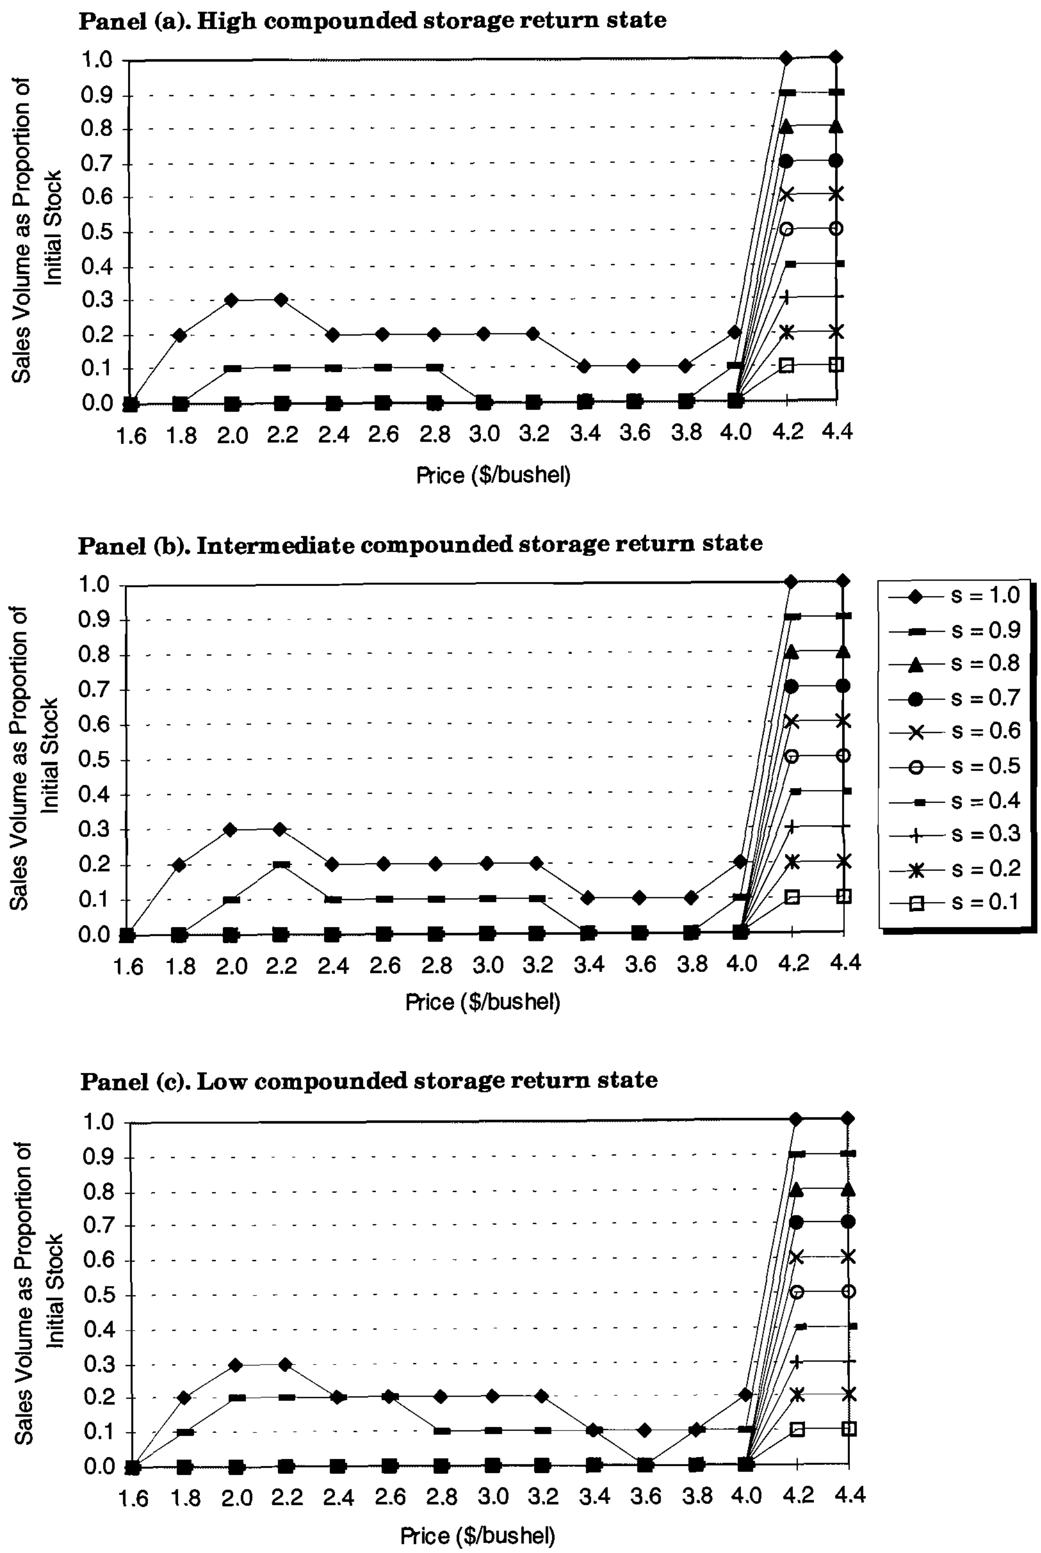

Figure 5. Optimal January marketing policy for risk-averse farmers

Table 3. Distribution of Expected Optimal Marketing Volumes Under Different Degrees of Risk Aversion

Relative Risk Aversion	November	January	March	May	July	September
$R = 0.0001$	0%	10%	0%	56%	22%	12%
$R = 0.5$	0%	10%	0%	63%	15%	11%
$R = 1.0$	7%	4%	0%	66%	13%	10%
$R = 2.0$	9%	2%	0%	73%	8%	8%
$R = 3.0$	17%	2%	0%	67%	7%	7%
$R = 5.0$	40%	2%	3%	45%	4%	6%
$R = 10.0$	65%	2%	4%	23%	3%	3%
1996–97 Aggregate Michigan Farm Corn Marketings	27%	27%	15%	8%	11%	12%

Under risk neutrality, optimal sales in any given year will always occur in one month only (sell all or nothing rule), but table 3 shows the expected frequency of those sales occurring at particular months over repeated samples (years). Under risk neutrality, it is most often optimal to sell in May (56% of the time), but sometimes sales occur in July and September (22% and 12%, respectively), and occasionally it is even optimal to sell in January (10% of the time). Notice, however, it is almost never optimal for a risk-neutral farmer to sell at harvest in November. As the farmer becomes more and more risk averse, the frequency of November sales grows dramatically, reaching 65% at $R = 10$. Furthermore, as the degree of risk aversion rises and the frequency of November sales increases, the frequency of May sales declines until it reaches just 23% at $R = 10$ (down from 56% under risk neutrality). Clearly, the optimal strategy under risk aversion is to spread sales out over the storage season, but most of the time sales will occur either at harvest (a risk-reduction strategy) or in May (an expected profit-generating strategy).

For purposes of simple comparison, aggregate Michigan farm marketings of corn are reported in the last row of table 3. The Michigan data are aggregated across individuals at a point in time, and so are not directly comparable to the optimal storage results in table 3 (which refer to expected frequencies over time of marketings from a single farmer). Nevertheless, there is clearly much more corn being sold at harvest in Michigan, and less being sold later in the storage season around May, than would be predicted from the optimal storage results under risk neutrality.

Performance of the Optimal Storage Rules

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), or a naive strategy of selling an equal amount of corn at every decision period (even sales), or a strategy of allocating sales according to the proportion of aggregate monthly corn marketings by Michigan farmers during the crop year (see the last row of table 3). The latter comparison is designed to examine how the optimal rules compare to simply following the historical temporal pattern of aggregate corn marketing by farmers.

To conduct the performance evaluation, we used the transition probabilities to compute the conditional mean and conditional standard deviation of final compounded storage returns (at the end of the storage season), conditional on each possible initial harvest period price in the state space, under six different storage rules: (*a*) sell everything at harvest, (*b*) sell an equal amount at each decision period, (*c*) allocate sales according to the proportion of aggregate monthly corn marketings by Michigan farmers during the crop year, (*d*) apply the optimal risk-neutral storage rule, (*e*) apply the optimal storage rule under $R = 5$, and (*f*) apply the optimal risk-averse storage rule under $R = 10$. Results are provided in table 4 using an interest rate of 10% and a storage cost of \$0.01 per bushel per month.

Selling everything at harvest is the least risky strategy because final compounded storage returns are then known with certainty (zero standard deviation). However, this strategy only returns the current initial harvest price (compounded through to the end of the storage season using the fixed 10% interest rate). This return can be high when the initial harvest price is high, but will be quite low when the initial harvest price is low (see table 4). When the initial harvest price is low, farmers could increase their expected return by storing rather than selling immediately, but this would also expose them to the risk of further price declines during the storage season.

The optimal storage rule under risk neutrality always generates the highest expected return (as anticipated), but also has the highest risk (standard deviation) because storing and waiting for prices to rise exposes the farmer to the risk that price will not rise or actually fall. An exception to this outcome occurs at very high current prices when the optimal risk-neutral strategy is to sell everything immediately, which is of course also a risk-free strategy. The naive "even sales" and "historical marketings" strategies still continue to store some corn, even in the face of these relatively high initial prices, which is why their returns are still exposed to some risk (see the bottom part of table 4).

As the degree of farmer risk aversion increases, the mean return from an optimal storage strategy generally declines relative to the risk-neutral strategy because the farmer is diversifying by selling at least some corn at harvest and otherwise spreading sales out over the storage season. However, the farmer also faces less risk than would have occurred under the optimal risk-neutral rule (smaller standard deviation than the optimal risk-neutral strategy). Of course, at high initial harvest prices (those at or above \$4.10), the optimal strategy is to sell everything immediately, irrespective of the degree of risk aversion. At these prices, therefore, the expected return from an optimal strategy equals the current price compounded through to the end of the storage season and there is no remaining risk, regardless of the degree of risk aversion (see the bottom part of table 4).

From the results in table 4, it is quite clear the optimal storage rules generally provide higher expected returns than the naive "no storage," "even sales," and "historical marketings" strategies, and in many cases also have lower risk than the simple "even sales" or "historical marketings" rules. Hence, the optimal strategies generally dominate the "even sales" and "historical marketings" strategies irrespective of the degree of risk aversion. The "no storage" rule is universally the least risky, but imposes a very high cost in terms of foregone expected returns. Which of the optimal strategies a farmer might prefer will depend on his or her willingness to trade off risk and expected return (i.e., the degree of farmer risk aversion).

Table 4. Mean and Standard Deviation of Final Compounded Storage Returns Under Alternative Storage Strategies Conditional on Different Initial Harvest Period Prices (interest rate = 10%, storage cost = $0.01/bushel/month)

		STORAGE STRATEGY					
Initial Price ($/bu.)	Performance Measure	No Storage	Even Sales	Historical Marketings	Optimal Storage (R = 0.0001)	Optimal Storage (R = 5)	Optimal Storage (R = 10)
1.6	Mean	1.76	1.86	1.83	1.95	1.92	1.86
	Std. Dev.	(0)	(0.334)	(0.251)	(0.370)	(0.255)	(0.156)
1.8	Mean	1.98	2.03	2.01	2.11	2.04	2.02
	Std. Dev.	(0)	(0.345	(0.263)	(0.410)	(0.186)	(0.105)
2.0	Mean	2.20	2.24	2.22	2.33	2.25	2.23
	Std. Dev.	(0)	(0.354)	(0.270)	(0.459)	(0.180)	(0.090)
2.2	Mean	2.42	2.46	2.44	2.56	2.48	2.45
	Std. Dev.	(0)	(0.363)	(0.277)	(0.505)	(0.200)	(0.100)
2.4	Mean	2.64	2.68	2.66	2.79	2.70	2.67
	Std. Dev.	(0)	(0.370)	(0.283)	(0.547)	(0.218)	(0.109)
2.6	Mean	2.86	2.92	2.89	3.03	2.92	2.89
	Std. Dev.	(0)	(0.377)	(0.290)	(0.585)	(0.231)	(0.117)
2.8	Mean	3.08	3.14	3.11	3.26	3.15	3.12
	Std. Dev.	(0)	(0.384)	(0.298)	(0.615)	(0.245)	(0.123)
3.0	Mean	3.30	3.37	3.34	3.49	3.34	3.34
	Std. Dev.	(0)	(0.388)	(0.304)	(0.630)	(0.251)	(0.125)
3.2	Mean	3.52	3.59	3.56	3.71	3.60	3.56
	Std. Dev.	(0)	(0.387)	(0.306)	(0.628)	(0.249)	(0.124)
3.4	Mean	3.74	3.79	3.77	3.92	3.81	3.47
	Std. Dev.	(0)	(0.385)	(0.309)	(0.611)	(0.239)	(0.120)
3.6	Mean	3.96	3.98	3.97	4.11	4.02	3.99
	Std. Dev.	(0)	(0.377)	(0.307)	(0.570)	(0.220)	(0.110)
3.8	Mean	4.18	4.13	4.14	4.28	4.21	4.20
	Std. Dev.	(0)	(0.366)	(0.300)	(0.515)	(0.149)	(0.098)
4.0	Mean	4.40	4.25	4.29	4.43	4.40	4.40
	Std. Dev.	(0)	(0.351)	(0.290)	(0.446)	(0.043)	(0.043)
4.2	Mean	4.62	4.35	4.42	4.62	4.62	4.62
	Std. Dev.	(0)	(0.331)	(0.271)	(0.000)	(0.000)	(0.000)
4.4	Mean	4.84	4.42	4.51	4.84	4.84	4.84
	Std. Dev.	(0)	(0.309)	(0.247)	(0.000)	(0.000)	(0.000)

Conclusions

This study extends the risk-neutral on-farm storage model of Fackler and Livingston 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 by Fackler and Livingston. 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 risk-averse farmers will sell a proportion of their corn crop early (generally right at harvest) unless harvest prices are extremely low. Farmers use this strategy even though it reduces their expected compounded storage return because it also reduces risk. This result confirms that risk aversion is capable of explaining the observed behavior of partial sales over the storage season without assuming farmers are somehow myopic or failing to optimize.

Performance comparisons indicate that optimal storage rules can generate considerably higher expected returns than selling everything at harvest, though at the cost of increased risk. However, part of this risk can be mitigated (at the cost of part of the gain in expected return) by optimally diversifying sales over the storage season based on current observed price levels, expectations about future price movements, and the degree of farmer risk aversion.

[Received November 2001; final revision received August 2003.]

References

Asplund, N. M., D. L. Forster, and T. T. Stout. "Farmers' Use of Forward Contracting and Hedging." *Rev. Futures Mkts.* 8(1989):24–87.

Baillie, R. T., and R. J. Myers. "Bivariable GARCH Estimation of the Optimal Commodity Futures Hedge." *J. Appl. Econometrics* 6(1991):109–24.

Benninga, S., R. Eldor, and I. Zilcha. "The Optimal Hedge Ratio in Unbiased Futures Markets." *J. Futures Mkts.* 4(Summer 1984):155–59.

Berg, E. "A Sequential Decision Model to Determine Optimal Farm-Level Grain Marketing Policies." *Eur. Rev. 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. Econ. and Statis.* 69(August 1987):542–47.

Dixit, A. K., and R. S. Pindyck. *Investment Under Uncertainty.* Princeton NJ: Princeton University Press, 1994.

Fackler, P. L. "Futures Price Volatility: Modeling Non-Constant Variance." Paper presented at annual meetings of the AAEA, Reno NV, July 1986.

Fackler, P. L., and M. J. Livingston. "Optimal On-Farm Storage." *Amer. J. Agr. Econ.* 84(August 2002):1103–71.

Ferris, J. "Using Seasonal Cash Price Patterns for Selling Decisions on Corn, Soybeans, and Wheat." NCR Extension Pub. No. 217, Michigan State University, 1996.

Hilker, J. Professor, Department of Agricultural Economics, Michigan State University. Personal communication, December 2002.

Lence, S. H. "On the Optimal Hedge Under Unbiased Futures Prices." *Econ. Letters* 47(1995):385–88.

Lence, S., K. Kimle, and M. L. Hayenga. "A Dynamic Minimum Variance Hedge." *Amer. J. Agr. Econ.* 75(November 1993):1103–71.

Miranda, M. J., and P. L. Fackler. *Applied Computational Economics and Finance.* Cambridge MA: MIT Press, 2002.

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

Myers, R. J., and S. D. Hanson. "Optimal Dynamic Hedging in Unbiased Futures Markets." *Amer. J. Agr. Econ.* 78(1996):13–20.

Newbery, D. M. G., and J. E. Stiglitz. *The Theory of Commodity Price Stabilization.* Oxford: Clarendon Press, 1981.

Sartwelle, J., D. O'Brien, W. Tierney, Jr., and T. Eggers. "The Effect of Personal and Farm Characteristics upon Grain Marketing Practices." *J. Agr. and Appl. 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. Agr. Econ.* 73(February 1991):75–88.
U.S. Department of Agriculture/National Agricultural Statistics Service. *Agricultural Statistics 2001.* USDA/NASS/Statistical Reporting Service. Washington DC: Government Printing Office, 2002.
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. Agr. Econ.* 80(November 1998):806–20.
Yang, S.-R., and B. W. Brorsen. "Nonlinear Dynamics of Daily Cash Prices." *Amer. J. Agr. Econ.* 74(August 1992):706–15.
Zulauf, C. R., and S. H. Irwin. "Market Efficiency and Marketing to Enhance Income of Crop Producers." *Rev. Agr. Econ.* 20(Fall/Winter 1998):308–31.

Appendix: Derivation of the Optimality Conditions for Storage

Differentiating Bellman's equation [text equation (4b)] with respect to q_t after using a Lagrangian to account for the no short-selling constraint ($q_t \le s_t$) gives:

(A1a) $$E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}(1+r)(p_t - c + k) - \frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial s_{t+1}}\right\} - \lambda_t \le 0,$$

(A1b) $$q_t\left[E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}(1+r)(p_t - c + k) - \frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial s_{t+1}}\right\} - \lambda_t\right] = 0,$$

and

(A1c) $$\lambda_t(q_t - s_t) = 0.$$

Here, λ_t is the shadow value of relaxing the short-selling constraint. From the envelope theorem we have:

(A2a) $$\frac{\partial v_t(\mathbf{x}_t)}{\partial s_t} = -k(1+r)E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}\right\} + E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial s_{t+1}}\right\}$$

and

(A2b) $$\frac{\partial v_t(\mathbf{x}_t)}{\partial \pi_t} = (1+r)E_t\left\{\frac{\partial v_{t+1}(\mathbf{x}_{t+1})}{\partial \pi_{t+1}}\right\}.$$

Now substituting (A2a) and (A2b) into (A1) results in equations (6) in the text.