NCCC-134
APPLIED COMMODITY PRICE ANALYSIS, FORECASTING AND MARKET RISK MANAGEMENT

Optimal On-Farm Storage

by

Paul L. Fackler and Michael J. Livingston

Suggested citation format:

Fackler, P. L., and M. J. Livingston. 1996. "Optimal On-Farm Storage." Proceedings of the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management. Chicago, IL. [http://www.farmdoc.uiuc.edu/nccc134].

Optimal On-Farm Storage

Paul L. Fackler and Michael J. Livingston*

When transactions costs prohibit an agricultural producer from replenishing grain stocks during the post-harvest marketing season, sales out of storage may be viewed as irreversible investments. The irreversibility of sales decisions transforms the dynamic marketing problem into one that is analogous to the optimal exercise of a financial option. A procedure is developed to solve the producer's marketing/storage problem and is applied to the cases of North Carolina and Illinois soybeans. Decision rules derived from this procedure are shown to be practically significant relative to simple marketing strategies that ignore the irreversible nature of the sales decision.

Introduction

Agricultural producers with on-site storage facilities may have considerable flexibility in post-harvest marketing strategies since they can choose the timing and quantities of sales out of storage. A simple example is the myopic strategy that transforms the producer's dynamic problem into a static problem. If the producer is modeled as being able to make marketing decisions at discrete periods during the post-harvest marketing season, then the myopic strategy calls for sales decisions to be made on the basis of the prospects for sales in the current and in the next period only. In fact, this is the optimal strategy if stocks can always be replenished next period at no cost. Bid/ask spreads, loading, unloading, and transport costs, however, are likely to be large for the average producer.

For many producers a more realistic assumption is that these costs are high enough to preclude any replenishing of stocks once they have been sold. Under this assumption sales out of storage can be viewed as irreversible investments: stocks can not be recovered once they are sold. Recently a considerable literature has developed concerning the effects of irreversibility on optimal investment decisions. McDonald and Siegel(1986), Brennan and Schwartz(1985), and Dixit and Pindyck(1994) recognized that making an irreversible investment is equivalent to exercising a financial option and drew on the finance literature to solve for the value of this investment option. The essential result of their analyses is that the standard net present value rule is suboptimal because it ignores the cost of exercising an irreversible investment option. The optimal rule when an investment is irreversible is to invest when the expected present value of future payoffs exceeds the sunk investment cost plus the value of exercising the investment option. Viewing current sales out of storage as irreversible suggests that the myopic rule, which is akin to the standard net present value rule, is not optimal and calls for sales to be made too soon.

* The authors are, respectively, an Associate Professor and a Graduate Research Assistant at the Department of Agricultural and Resource Economics, North Carolina State University, Raleigh NC.

The practical significance of this result depends on the magnitude of the value of exercising an irreversible investment option. This will depend on the expected future payoffs, the discount rate, the time remaining on the option, and the degree of volatility of expected payoffs. It is well known that the value of a financial option increases with volatility and the time remaining on the option. This implies that an increase in uncertainty makes one less willing to commit to an irreversible decision. Furthermore, Pindyck(1991) notes that increases in interest rates lower the value of exercising irreversible investment options.

To investigate the effects of irreversibility on optimal marketing decisions, a stochastic dynamic programming model is developed for a risk-neutral producer. The dynamic marketing problem is shown to be a particular example of an optimal stopping problem wherein the optimal sales rule reduces to a simple condition on the current price: sell everything when the current price is high; otherwise store everything into the next period.

The optimal marketing rules are derived in the next section as the solution to a stochastic dynamic programming problem. The following section extends the analysis to the continuous-time case. The next section provides a brief discussion of how to calculate optimal and suboptimal decision rules for the discrete-time case. An empirical example illustrates the results of these calculations in the next section using estimated models for North Carolina and Illinois soybean prices. The optimal decision rules are shown to be practically significant relative to suboptimal rules, and comparative statics exercises illustrate the effects of changing interest rates, unit storage costs, and conditional volatilities on the optimal rule. The paper is closed with a brief summary and some concluding remarks.

Optimal Marketing Rules

A producer with on-site storage facilities can decide how much of production to store and how much to sell following the harvest. Periodically through the marketing season, the producer can reevaluate the sales decision in light of current information. We model the producer's optimization problem as an N-period, discrete, stochastic, dynamic programming problem. The following assumptions formalize the structure of the problem:

A1: The producer is a risk-neutral price-taking profit-maximizer intending to maximize the expected present value of N-period wealth, given an initial stock of harvested production, s_0.

A2: The bid/ask spread and loading, unloading, and transport costs prohibit the producer from replenishing storage stocks once they are sold, thus sales can never be negative.

A3: All stocks must be sold by the end of the marketing period, at t=N, and thus sales can not exceed harvested stocks in any period of the marketing season.

A4: Prices follow a Markov process, and the conditional cumulative distribution function for p_{t+1} given p_t shifts to the right when p_t increases. That is, there is positive autocorrelation in prices. Furthermore, it is assumed that the correlation coefficient is strictly less than 1+r, where r denotes a constant interest rate.[1]

---

1. This assumption is taken to ensure that the programming problem can be solved. In the applications to North Carolina and Illinois soybeans, the autocorrelation coefficients on lagged

A5: The producer discounts the future at rate $\beta=1/(1+r)$.

The structure of the problem is completed by the following specification for period-t profits, two state transition equations, and the following variable definitions:

$$\pi_t = q_t p_t - (s_t - q_t)k$$
$$s_{t+1} = s_t - q_t$$
$$p_{t+1} \sim f_t(p_{t+1} \mid p_t)$$

where

π_t: the net gain per period, sales revenues less storage costs
s_t: the carry-in stocks level at period t, a deterministic state variable
p_t: the price at period t, a random state variable
q_t: the amount sold at period t for price p_t, the control variable
k: a constant per unit storage rate
f_t: a probability density function for p_{t+1} conditional on p_t.

The value function for any decision period, t, denotes the optimized value of stocks for that period and, for our problem, satisfies the following Bellman equation:

$$V_t(p_t,s_t) = \max_{q_t} q_t p_t - (s_t - q_t)k + \beta E_t\left[V_{t+1}(p_{t+1},s_{t+1})\right] \quad \text{s.t. } 0 \le q_t \le s_t.$$

The irreversibility of sales out of storage is embodied in this simple inequality constraint.

It can be shown that the optimal value function for this problem, V_t, has two important features that greatly simplify the optimal decision rule. First, the value function is linear in s_t, therefore, the problem can be expressed in terms of the per-unit value function, v_t. Second, at any given period it is optimal to either sell all stocks or to sell none. The value function can, therefore, be written as the maximum of the value of immediate sales and the value of waiting:[2]

$$(1) \qquad v_t(p_t) = \max\left(p_t, \beta E_t\left[v_{t+1}(p_{t+1}) \mid p_t\right] - k\right).$$

The dynamic programming problem is solved backwards in time by initializing the N-period value function implied by (A3). Since all stocks must be sold by period N, the N-period per-unit value function is

$$(2) \qquad v_N(p_N) = p_N \, .$$

The optimal sales strategy can be summarized as follows:

i) $q_t = s_t$ if $v_t(p_t) = p_t$, sell everything; or
ii) $q_t = 0$ if $v_t(p_t) \ge p_t$, store everything.

---

price are less than one and so satisfy this assumption.

2. The problem is a special example of what is commonly referred to as an optimal stopping problem; see Dixit and Pindyck (1994) for a discussion of optimal stopping problems in the context of irreversible investment decisions.

If we define c_t to be the root of $v_t(p_t)-p_t$, so that $v_t(c_t)-c_t=0$, then the optimal rule can also be expressed in the following manner:

i) sell everything if $p_t \geq c_t$; or

ii) store everything if $p_t < c_t$.

This is a particularly simple decision rule that requires only knowledge of the current period's price and cutoff price, c_t.

The decision to sell is analogous to the decision to exercise an American option. In the option pricing literature it is common to decompose the value of an option into its intrinsic value (what it is worth if exercised) and its time value (its value if held). It is well known that a financial option should only be exercised when the intrinsic value is positive and the time value is zero. In the marketing problem, the intrinsic value of stocks is the current price, which is always positive. It is, therefore, optimal to refrain from selling as long as stocks in storage have time value and to exercise when their value is equal to p_t.

Optimal Rules in Continuous Time

In the discrete-time dynamic programming approach to optimal sales timing, the decision maker reassesses previous decisions at fixed intervals over the marketing season. As formulated, there is no cost to reassessment and it can, therefore, be done continuously.

Assumptions (A1)-(A5) hold and remain unchanged. In continuous time, the marketing problem is also greatly simplified by defining the value function as the optimized unit value of stocks. This is done, without loss of generality, by setting the initial harvested stocks, s_0, to unity. The structure of the continuous-time optimal stopping problem is given by a specification for the stochastic differential equation for price, and the following variable definitions:

$$dp = a(p,t)dt + b(t)dz;$$

where

dp: the instantaneous change in current price at time t;

dz: the increment to a standard Wiener process;

a(p,t): a deterministic function of price and time, the drift coefficient;

b(t): a deterministic function of time, the instantaneous standard deviation;

r: a constant interest rate per unit time; and

k: a constant unit storage rate per unit time.

The value function is given by

$$(3) \qquad v(p_t,t) = \max\left(p_t,\ -kdt + (1+rdt)^{-1}E_t\left[v(p_t+dp,t+dt)\right]\right);$$

subject to the stochastic differential equation for price. Notice that the inequality that embodies the irreversible nature of the sales decision need not be explicitly modeled. It is implicit in the formulation of the producer's marketing problem as an optimal stopping problem. If the first term in the max(.,.) function is the largest, stopping is optimal, and everything is sold for the current price. Otherwise, continuation is optimal, and everything is stored into the next instant of time.

Expanding $E_t[v(p_t+dp,t+dt)]$ by Ito's Lemma gives

$$E_t\left[v(p+dp,t+dt)\right] = v(p,t) + \left[v_t(p,t) + a(p,t)v_p(p,t) + \frac{1}{2}b(t)^2v_{pp}(p,t)\right].$$

The value function is given by the second term in the max(.,.) function whenever it is larger than the current price.[3] Multiplying both sides of equation (3) by (1+rdt), and letting dt go to zero, gives the following partial differential equation that must hold whenever the current price is in the continuation region:

$$(4) \qquad \frac{1}{2}b(t)^2 v_{pp} + a(p,t)v_p + v_t - rv - k = 0 .$$

This partial differential equation holds only in the continuation region, which is itself unknown.[4] In order to solve simultaneously for the value function and the free boundary that separates the stopping region from the continuation region, two additional restrictions are usually placed on the value function. The value function is set equal to its value in the stopping region at the free-boundary price, c(t), so that

$$(5) \qquad v(c(t),t) = c(t) .$$

Finally, a constraint on the partial derivative of the value function evaluated at the free boundary is imposed[5]

$$(6) \qquad v_p(c(t),t) = 1 .$$

Equation (4) and the free boundary, c(t), are found subject to (5) and (6), the terminal condition that all stocks must be sold by t=N, the parameters of the diffusion process for price, and values for storage and interest rates. The optimal decision rule, for any moment in time over the post-harvest marketing season, is to sell everything as soon as the current price, p_t, equals or surpasses the free-boundary price at that time, c(t). The continuous-time solution is elegantly and compactly represented as the solution to a free-boundary problem. All that is

---

3. Subscripts denote partial derivatives.

4. Continuous-time optimal stopping problems of this nature are sometimes referred to as free-boundary problems. The value function and the continuation region over which the partial differential equation is satisfied are both unknown. Thus the free boundary separating the stopping region from the continuation region must be solved for, along with the value function itself.

5. Free-boundary prices denote points of indifference between selling everything and storing everything, since the first and second terms in (3) are equal for these prices. Equations (5) and (6) are referred to as the value-matching and smooth-pasting conditions, respectively. The value-matching condition is imposed, because the value function is continuous in price and, at the free boundary, must equal the current price. The smooth-pasting condition constrains the partial derivative of the value function with respect to price equal to the partial derivative of the return to stopping with respect to price. For more on smooth-pasting conditions see Dixit and Pindyck (1994): 130-32.

needed to make optimal marketing decisions is the free boundary and knowledge of the current price. In order to solve the free-boundary problem, parameter estimates for the stochastic differential price equation are needed. It can be shown that estimates of these parameters can be obtained using the estimation model in our empirical example. The analysis presented below, however, does not solve for the continuous-time cutoff prices. Instead, we use the discrete-time results derived in the first section, wherein the cutoff price for any period, t, is the root of $v_t(p_t)-p_t$, where $v_t(p_t)$ is given by equation (1). The next section discusses briefly how these roots are calculated and introduces two suboptimal marketing strategies which will be used for purposes of comparison.

Optimal and Suboptimal Cutoff Prices

Although the optimal decision rule is simple, the calculation of the cutoff prices is somewhat involved, because closed-form expressions do not exist. Instead, the cutoff prices are calculated numerically using the backwards recursion defined implicitly by equations (1) and (2). This is most easily accomplished by approximating the per-unit value functions as piecewise linear functions over some grid of price points and using trapezoid integration to compute expectations. The cutoff prices are then found as the roots of the approximated per-unit value functions.

Since the calculation of the optimal cutoff prices is somewhat involved, it is useful to compare the optimal cutoff prices to cutoff prices that are easier to calculate. A myopic marketing strategy considers only the prospects for sales in the current and the next period, implicitly treating the next period as the last period. The myopic value function can be expressed as

$$v_t^{\,m}(p_t) = \max\left(p_t, \beta E[p_{t+1} \mid p_t] - k\right).$$

A more interesting comparison is with the so-called open-loop-with-feedback strategy. The open-loop strategy refrains from selling all stocks in the current period if the discounted conditional expected unit value of sales in a future period exceeds the current price. The open-loop control, however, does not take into account that the timing of sales can change in future periods. The open-loop value function is defined to be

$$v_t^{\,o}(p_t) = \max_{0 \le i \le N-t} \beta^i E[p_{t+i} \mid p_t] - k\sum_{j=1}^{i-1} \beta^j .$$

In both cases, cutoff prices can be computed that solve $v_t(c_t)=c_t$. It can be shown that the myopic cutoff price can be no greater than the open-loop-with-feedback cutoff price, which can be no greater than the optimal cutoff price:

$$c_t^{\,m} \le c_t^{\,o} \le c_t .$$

In the empirical example, optimal, myopic, and open-loop cutoff prices are compared to get a feel for the practical significance of the optimal marketing strategy relative to these suboptimal strategies that are much easier to implement but which ignore the irreversibility of sales.

An Empirical Example

The analysis is applied to an estimated model of weekly spot price movements for North Carolina and Illinois soybeans. The data are weekly (Thursday) soybean prices over the post-harvest marketing period, assumed to begin in November and run through the end of June, for the years 1976-1992 (Illinois) and 1976-1993 (North Carolina). The North Carolina prices are cash bid prices at a soybean crushing facility in Fayetteville reported in North Carolina Department of Agriculture Market Grain Reports. The Illinois prices are central Illinois elevator average bid prices reported in the Wall Street Journal (obtained from the TechTools data base).

The estimated model allows for seasonal variation in both the conditional mean and conditional variance of prices. Specifically,

$$(7) \qquad p_{t+\Delta t} = m(t+\Delta t) - \alpha\left(p_t - m(t)\right) + \varepsilon_{t+\Delta t} \; ;$$

where

$$(8) \qquad \varepsilon_{t+\Delta t} \sim N\left(0, v(t+\Delta t) - \alpha^2 v(t)\right).$$

Both $m(t)$ and $v(t)$ are seasonally periodic functions, i.e., $m(t+i)=m(t)$, for integer i (when time is measured in years). This stochastic model can be shown to be the discrete time processes corresponding to an Ito process of the form

$$dp = \rho\left(a(t) - p\right)dt + b(t)dz,$$

where $a(t)$ and $b(t)$ are seasonally periodic.

With data observed at even time increments, this model can be estimated using the iterative GLS procedures described in McNew and Fackler(1994) in the form

$$p_t = \mu(t) - \alpha p_{t-1} + \varepsilon_t \; , \text{ with } \varepsilon_t \sim N\left(0, \sigma^2(t)\right).$$

The seasonally periodic functions $\mu(t)$ and $\sigma^2(t)$ are estimated using truncated Fourier series of the form

$$\mu(t) = \eta_0 + \sum_{i=1}^{n} \left(\eta_i \cos(2\pi it) + \theta_i \sin(2\pi it)\right)$$

and

$$\sigma^2(t) = \psi_0 + \sum_{i=1}^{n} \left(\psi_i \cos(2\pi it) + \omega_i \sin(2\pi it)\right).$$

The specifications for $\mu(t)$ and $\sigma^2(t)$ imply similar Fourier series specifications for m(t) and v(t), respectively; and the coefficients of the latter are obtained using simple transformations on the coefficients of the former. Once the coefficients of m(t) and v(t) are obtained, it is possible to estimate conditional mean and variance functions for decision horizons of arbitrary length.[6]

The conditional heteroskedasticity is confirmed using Breusch-Pagan/Godfrey tests. For both sets of price data, sequential F-tests are used to arrive at the order specifications for the conditional mean and variance functions.[7] Maximum likelihood parameter estimates, and asymptotic standard errors and p-values for Illinois are reported in Table 1.

Table 1

Variable	Estimate	Standard Error	P-Value
η_0	0.117	0.059	0.047
η_1	0.017	0.026	0.510
θ_1	0.034	0.039	0.388
η_2	-0.024	0.018	0.181
θ_2	-0.014	0.024	0.573
α	0.980	0.008	0.000
ψ_0	0.066	0.022	0.003
ψ_1	-0.045	0.021	0.030
ω_1	-0.034	0.032	0.281
ψ_2	0.009	0.013	0.525
ω_2	0.031	0.022	0.152
ψ_3	0.007	0.010	0.498
ω_3	-0.021	0.006	0.000

Using these parameter estimates, optimal, open-loop, and myopic cutoff prices are calculated for North Carolina and Illinois. Weekly cutoff prices for Illinois are shown in Figure 1. The analysis uses a base case with an annual interest rate of 5% and a monthly storage cost of $0.05 per bushel. Figure 2 shows cutoff prices for the monthly decision-horizon for Illinois. Notice that the optimal cutoffs are always larger than the open-loop cutoffs which, in

---

6. Note that α is an estimate of $e^{-\rho\Delta t}$. Conditional means and variances for arbitrarily long decision horizons are obtained from equations (7) and (8) by raising the estimate of α to the power appropriate to the Δt of interest.

7. The sequential F-test begins by specifying very high-order seasonal functions for the mean and variance. A joint F-test of the hypothesis that the last two terms in a seasonal function are zero is carried out until a specification is reached wherein the last two terms are statistically significant. This analysis specifies first- and second-order seasonal functions for the conditional means for North Carolina and Illinois, respectively, and third-order seasonal functions for the conditional variance for both North Carolina and Illinois.

turn, are always larger than the myopic cutoffs. This holds until period t=N-1, at which time, all cutoff prices are the same, since the next period, t=N, is the last period. The maximum difference between the weekly optimal and open-loop cutoffs is $0.21, occurring at the beginning of the marketing season and declining as the end of the season approaches. This implies that the ability to recognize that future optimal controls can change makes a practical difference in the calculation of optimal cutoff prices when sales out of storage are irreversible, especially over the first two-thirds of the marketing season. The maximum difference between the weekly optimal and myopic cutoff prices is $1.87, occurring around the second week of January, at which time the differential decreases monotonically as the end of the marketing season approaches. The interpretation of this empirical result is that the time value of the irreversible sales option is largest towards the beginning of the marketing season and steadily declines as the end of the season approaches. In fact, this is a well known theoretical result concerning the time value of financial options. At the beginning of the marketing season there are simply more future time periods in which to consider selling. Furthermore, when sales out of storage are irreversible, the myopic cutoffs greatly understate the unit value of storage stocks and call for sales to be made too early. These results are obtained for the North Carolina case as well.

Table 2 contains ex post results of applying the different marketing strategies to the historical weekly Illinois

Figure 1

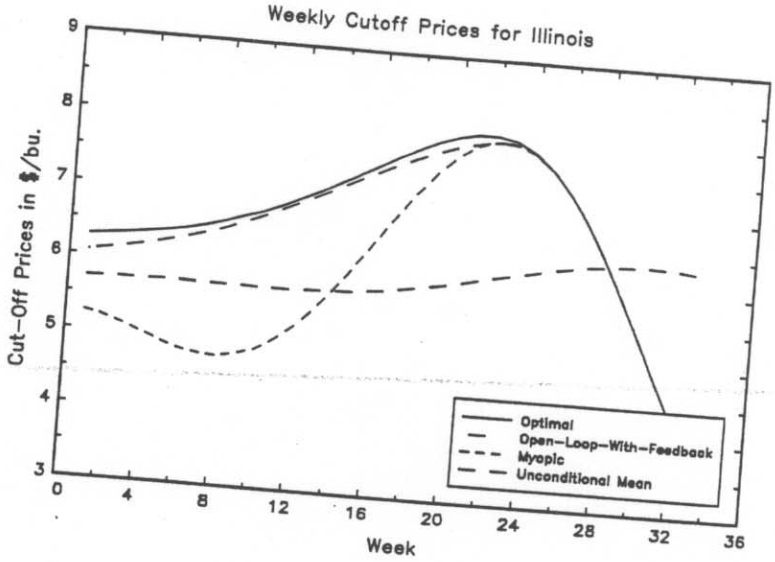

Figure 2

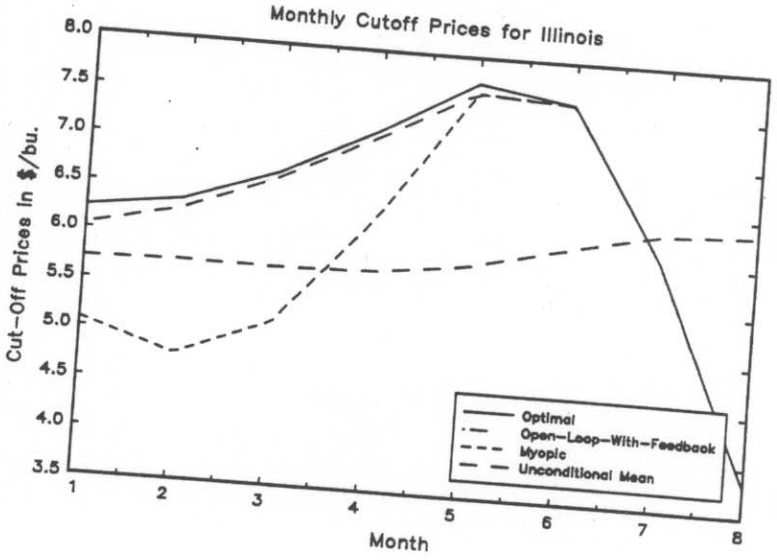

price data. The first column contains the prices for the first week

Table 2

		k=$0.00/bushel				k=$0.05/bushel			
Year	p_1	π_{optim}	π_{op}	π_{myopic}	$\pi_{optim-fh}$	π_{optim}	π_{op}	π_{myopic}	$\pi_{optim-fh}$
1976	6.55	8.63	8.25	6.55	6.54	6.55	6.55	6.55	6.55
1977	5.58	6.87	6.87	5.98	5.58	6.53	6.53	5.58	5.33
1978	6.89	6.95	6.89	6.89	7.11	6.89	6.89	6.89	6.89
1979	6.08	5.94	5.94	6.08	6.70	6.51	6.08	6.08	5.95
1980	8.53	8.53	8.53	8.53	8.53	8.53	8.53	8.53	8.53
1981	6.39	6.03	6.03	6.39	6.78	6.39	6.39	6.39	6.39
1982	5.41	5.84	5.84	5.50	5.80	5.63	5.63	5.41	5.44
1983	8.45	8.45	8.45	8.45	8.45	8.45	8.45	8.45	8.45
1984	6.27	5.51	5.51	6.27	6.53	5.25	6.27	6.27	6.19
1985	5.12	5.12	5.12	5.12	5.20	4.78	4.78	4.95	4.85
1986	4.89	5.31	5.31	5.31	4.78	4.97	4.97	4.63	4.44
1987	5.25	7.18	7.18	5.74	5.26	6.45	6.45	5.25	5.02
1988	7.75	7.75	7.75	7.75	7.75	7.75	7.75	7.75	7.75
1989	5.56	5.74	5.74	5.77	5.91	5.56	5.56	5.56	5.51
1990	5.84	5.58	5.58	5.84	6.15	5.27	5.27	5.84	5.81
1991	5.51	5.93	5.93	5.59	5.50	5.45	5.45	5.51	5.18
average	6.25	6.58	6.56	6.36	6.41	6.31	6.35	6.23	6.14
average weeks of storage		29	29	11	29	26	28	6	26
first cutoff price		6.97	6.76	5.80	6.97	6.28	6.06	5.25	6.28

of November for the corresponding year. The next three columns are historical, discounted, per-bushel returns for the optimal, open-loop, and myopic marketing strategies, respectively.[8] The fourth column is the second column plus the discounted returns from a fully-hedged position, wherein the hedge is marked-to-market at each rollover date and at the date the stocks are liquidated. The next four columns are the same, but the monthly storage cost, k, is changed from $0.00/bushel to $0.05/bushel. The annual interest rate is 5% in both columns. For the case of zero storage costs, the average optimal per-bushel return is the highest, followed closely by the average open-loop return. Historical annual returns are the same under both strategies, except for the years 1976 and 1978. The average return for the optimal marketing strategy assuming a fully-hedged futures position is next, followed closely by the myopic strategy, and the average return from always selling in the first week of November.[9] Also notice that when storage occurs, the optimal and open-loop strategies store for an average of 29 weeks over the marketing season,followed by an average of only 11 weeks for the myopic strategy. For the case of positive storage costs, the average return for the open-loop strategy is slightly higher than for the optimal strategy. Again, the historical annual

8. Per-bushel returns are discounted back to the first week of November, so that all listed returns are directly comparable.

9. Under the implicit assumption that knowledge of the current futures price does not affect conditional mean and variance forecasts, optimal cutoff prices do not need to be modified.

Figure 3

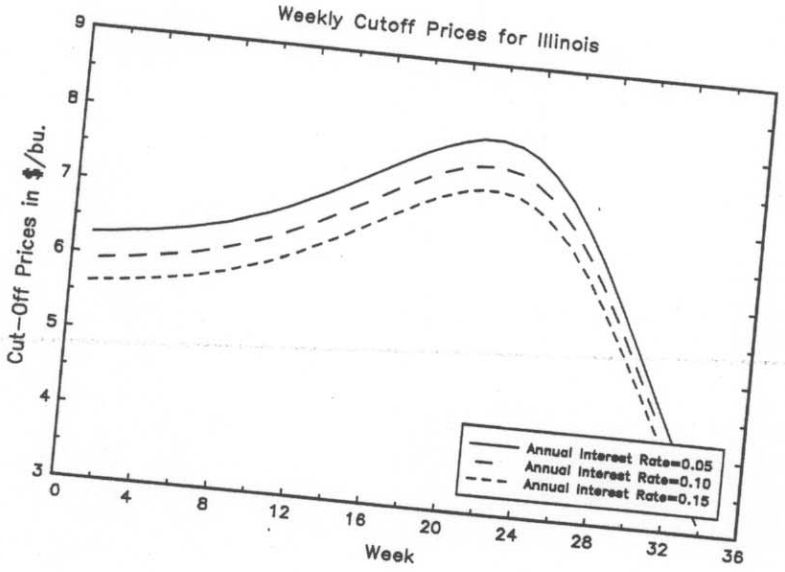

Figure 4

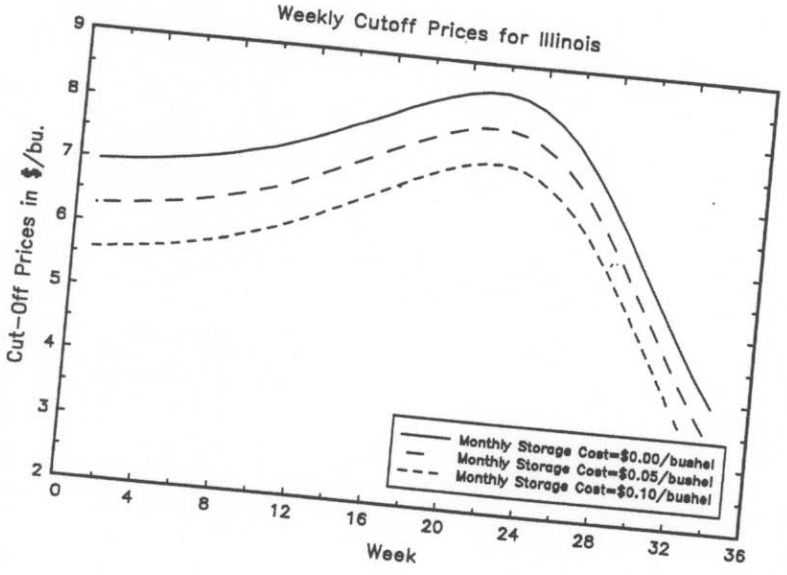

returns under both strategies differ for only two years. The optimal strategy, however, still outperforms the other strategies. Furthermore, the increase in the unit storage cost decreases the average time stocks spend in storage for all of the marketing strategies. Since such a small sample is used in this historical experiment, however, the results are not general but, nonetheless imply that the optimal and open-loop strategies produce similar results with respect to average unit returns.

Figures 3-6 illustrate some comparative statics results of varying the underlying model parameters. Each presents the base case and two alternative, weekly, optimal cutoff price paths. Figures 3 and 4 demonstrate the intuitively reasonable result that cutoff prices decline as either the interest rate or the per-unit storage cost are increased. Both kinds of increases make storage less desirable. Interestingly, interest rate and cost changes cause approximate parallel shifts in the cutoff price paths.

Figure 5 shows the effect of changing the level of uncertainty of the conditional price forecast. The figure shows optimal cutoff prices for conditional volatilities of one half and

Figure 5

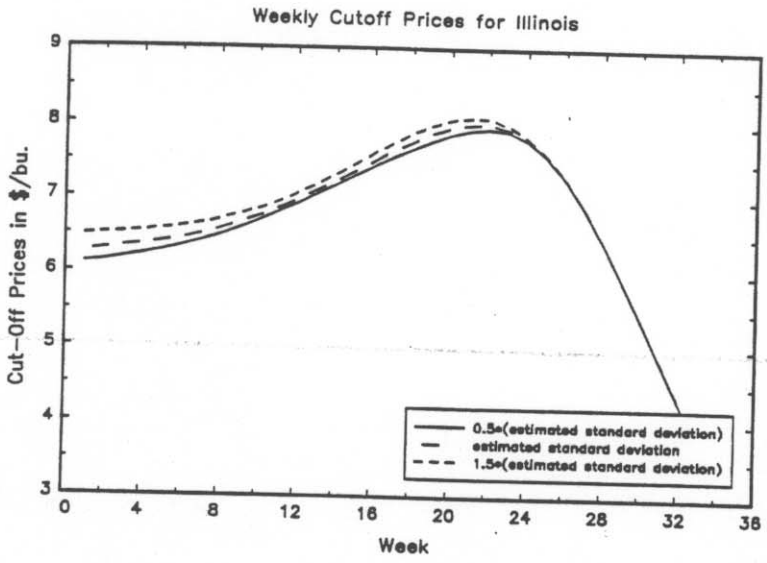

Figure 6

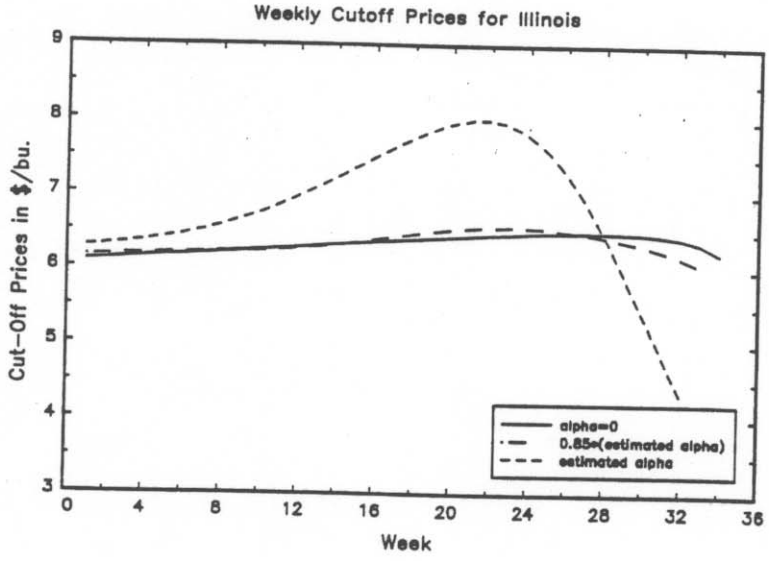

one and a half times the estimated conditional volatilities of the base case. As expected, the time value of the irreversible sales option increases with increases in conditional volatility.

Finally, Figure 6 shows optimal cutoff price paths for different levels of the mean-reversion parameter, α. As this figure shows, the cutoff price paths are very sensitive to this parameter. Great care, therefore, should be taken to ensure accurate estimates of this parameter. Small changes can affect the optimal cutoff prices dramatically.

Summary and Conclusions

When transactions costs prohibit agricultural producers from replenishing storage stocks once they are sold, sales out of storage during the post-harvest marketing period can be viewed as irreversible investments. This irreversibility confers an additional return to holding stocks in any one decision period; namely, the option value of being able to store stocks beyond the next immediate period. The optimal cutoff prices simply take this additional option value into account and so correctly value the return to storage in any given period.

Furthermore, the optimal cutoffs are empirically significant relative to other suboptimal marketing strategies which ignore the irreversible nature of the sales decision. In particular, if the current price is above the myopic cutoff price but below the optimal cutoff price, it is optimal to hold stocks even though the return from holding and selling in the next period is less than the current price.

It is shown, however, that optimal cutoff prices are highly sensitive to the model used to predict future prices and to underlying parameters like storage costs and interest rates. Before particular results of this analysis are applied in practice, therefore, care must be taken to ensure that forecast models are accurate and that the appropriate interest and storage costs are used. The forecast model used in the present analysis assumes that disturbances to the price process are additively normal. The reason for taking what may turn out to be a restrictive assumption on the stochastic price process is that it allows us to explicitly derive conditional transition densities for arbitrarily long decision horizons which are appropriate to an initial specification of instantaneous price changes as an Ito diffusion. This allows estimation of the parameters of the Ito diffusion process using discretely sampled data, which will be important in future research for calculating the free boundary of the continuous-time optimal stopping problem. It may be the case, however, that the costs of assuming normality outweigh the benefits. Research is presently underway to assess the restrictiveness of the normality assumption.[10]

It is a simple exercise to extend the methodology of this paper to the case in which knowledge of the current futures price affects conditional mean and variance forecasts of the cash price.[11] In this case, optimal cutoff prices are functions of the current futures price. Research is presently underway to calculate cutoff prices for this bivariate case under the assumption that weekly storage stocks are fully hedged in the futures market. It would also be useful to analyze the effects of optimal hedging strategies and the tax consequences of storage on optimal marketing decisions.

---

10. It may, however, be possible to model log price changes as an Ito diffusion similar to the one specified in the present analysis, preserving the tractable form of the discrete-time transition densities. Prices could then be modeled as being lognormally distributed, an assumption that may turn out to be less restrictive.

11. In the interest of brevity this extension is not included but is available from the authors upon request.

Appendix

This appendix demonstrates that the optimal storage problem can be expressed as a per unit optimal stopping problem. Recall that the Bellman equation for the problem is

$$V_t(p_t,s_t) = \max_{q_t} q_t p_t - (s_t-q_t)k + \beta E_t\left[V_{t+1}(p_{t+1},s_{t+1})\right] \quad \text{s.t. } 0\le q_t\le s_t.$$

The maximization problem can be formulated in terms of the Lagrangean function

$$L(q_t,\lambda_t) = p_t q_t - k(s_t-q_t) + \beta E_t\left[V_{t+1}(p_{t+1},s_{t+1})\right] + \lambda_t(s_t-q_t),$$

with the associated Kuhn-Tucker conditions

$$\text{(A1)} \qquad p_t - \beta E_t\left[\frac{\partial V_{t+1}(p_{t+1},s_{t+1})}{\partial s_{t+1}}\right] - \lambda_t \le 0,\ q_t \ge 0,\ \text{C.S.}$$

and

$$\text{(A2)} \qquad s_t - q_t \ge 0,\ \lambda_t \ge 0,\ \text{C.S.},$$

(C.S. denotes complementary slackness). The envelope theorem can be used to derive an expression for the partial derivative of the value function with respect to stocks:

$$\frac{\partial V_t(p_t,s_t)}{\partial s_t} = \beta E_t\left[\frac{\partial V_{t+1}(p_{t+1},s_{t+1})}{\partial s_{t+1}}\right] - k.$$

Noting that $\partial V_N/\partial s_N = p_N$ is not a function of s_N, it can be seen by induction that $\partial V_t/\partial s_t$ is not a function of s_t, so that V_t is linear in s_t. This allows us to write the first of the Kuhn Tucker conditions is terms of the per unit value function:

$$p_t - \beta E_t\left[v_{t+1}(p_{t+1})\right] - \lambda_t \le 0,\ q_t \ge 0,\ \text{C.S.}$$

Since this inequality does not depend on q_t, the Kuhn-Tucker conditions imply that q_t is either s_t or 0, depending on a condition on the random state variable, p_t. The value function is therefore the larger of the returns from selling now and the expected returns from holding:

$$v_t(p_t) = \max\left(p_t, \beta E_t\left[v_{t+1}(p_{t+1})\right] - k\right).$$

References

Brennan, M.J. and E.S. Schwartz. "Evaluating Natural Resource Investments." *Journal of Business*. 58(1985):135-157.

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

McDonald, R. and D.R. Siegel. "The Value of Waiting to Invest." *Quarterly Journal of Economics*. 101(1986):707-727.

McNew, K.P. and P.L. Fackler. "Nonconstant Optimal Hedge Ratio Estimation and Nested Hypotheses Tests." *The Journal of Futures Markets*. 14(1994):619-635.

Pindyck, R.S. "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature*. 29(1991):1110-1148.