Australian Journal of Agricultural Economics, Vol. 22, No. 3 (December 1978) pp. 158-174.

AUSTRALIAN WHEAT STORAGE: A DYNAMIC PROGRAMMING APPROACH*

CHRIS M. ALAOUZE†, N. H. STURGESS and A. S. WATSON

University of Melbourne

Dynamic programming is used to examine whether Australia should store wheat for subsequent sale at higher prices. The dynamic programming model is developed assuming that the demand for Australian wheat is perfectly elastic at the world price. An important consequence of this assumption is that the algorithm used in computing the optimal policy takes a very simple form. It is concluded that interest rate and wheat price in the following season are the major determinants of storage policy. When prices are stable, rule-of-thumb policies, such as storing when prices are below average, are sub-optimal.

Introduction

In this paper the problem of determining the optimal level of Australian wheat storage is analysed using control theory. Underlying the analysis is the assumption that the Australian Wheat Board (AWB) attempts to maximise the sum of current and future discounted expected net revenue from wheat sales. The inventory model presented in this paper is developed using the simplifying assumption that the demand curve for Australian wheat is infinitely elastic at the world price. The model was analysed using sets of simulated prices and a dynamic programming algorithm.

The model has the following features:

(a) It is assumed that the area sown to wheat is fixed (at the 1968/69 level) and that the yield distribution is stationary and normal.

(b) The cost of storage function is assumed constant over time. Only variable costs were considered and these were estimated for the 1973/74 season.

(c) The data from which price series were simulated span the period 1949/50 to 1973/74.

Thus the model is a simplification which represents no particular period in time. We believe it is a useful abstraction which permits inferences about some aspects of past and future wheat storage policy. Our decision to analyse the model for simulated prices and a fixed time horizon, rather than treat price as a state variable and find the steady state solution for an infinite time horizon, is based on the following considerations. Firstly, the steady state solution would result in a more 'operational' storage policy in the sense that if the model were realistic the AWB could use the solution as a basis for its inventory policy. Unfortunately, due to recent policy changes by major wheat importing and exporting countries, our understanding of future patterns of supply

* This work was supported by a grant from the Rural Credits Development Fund, Reserve Bank of Australia. We acknowledge the helpful comments of the anonymous referees.

† Present address is: Department of Economics, University of Wisconsin.

and demand in the world wheat market is speculative[1], thus limiting the operational usefulness of a steady state solution based on historical prices.

Secondly, a great deal is known about the supply and demand conditions that generated the historical price series for the period under consideration. The annual price series consists of a long period of low, stable prices (1953/54 to 1971/72) between two periods of high prices. The sequence of stable prices was the direct result of the large inventories held by the United States and Canada and, after 1968/69, by Australia. Over much of this period Australia was in a price-taking position (the exceptions are identified in the following section). From the viewpoint of the stable price period, the factors which caused the periods of high prices can be regarded as episodic.

This suggests two interesting questions concerning historical policy:

(a) If the AWB knew the mean of the stable price series with certainty, how would a policy of storing wheat in seasons of below average prices compare with the optimal policy?

(b) What is the optimal storage policy in seasons immediately preceding the occurrence of episodic price increase?

Approximate answers to these questions can be obtained by examining the pattern of storage over time associated with optimal storage policies determined for sets of simulated prices. Because of the costs associated with obtaining the steady state solution and its limited usefulness, we chose to investigate the two questions outlined above. The answer to the second question should provide some insight into the possible gains from holding speculative reserves to service episodic increases in demand in the future.

Another aspect of the paper is that we show that the assumption of an infinitely elastic demand curve for Australian wheat greatly simplifies the dynamic programming algorithm used in obtaining the optimal policy, thereby significantly reducing the cost and programming effort required. We also provide an economic interpretation of the solution procedure implied by Bellman's principle of optimality for inventory problems.

The Inventory Model

Dynamic programming

A concise treatment of the dynamic programming technique may be found in Bellman and Kalaba (1965) and Hastings (1973). The relationship of dynamic programming to other control theory problems is discussed extensively in Intriligator (1971) and Bellman and Dreyfus (1962). A discussion of the application of dynamic programming to problems arising in agricultural economics may be found in Throsby (1964) and Agrawal and Heady (1972). Dynamic programming was first applied to problems of grain storage by Gustafson (1958). Browning (1970) used the framework developed by Gustafson to evaluate the

[1] For a discussion, see Johnson (1975) and Alaouze et al. (1978, pp. 183-4).

costs of different methods of supplying food aid to less developed countries.

In the inventory model developed below, there is only one state variable, S_t, the amount of wheat available for export in the stage (season) with t stages of the process remaining. In the first season it is assumed that there is no carryover from the previous season, so the amount available for export, S_n, is equal to the level of deliveries to the AWB net of home consumption, X_n. The decision variable is the level of carryover, C_t. This transforms the amount available for export in the first season to $(X_n - C_n)$; the amount available for export (prior to transformation) in the following season[2] is then $S_{n-1} = (X_{n-1} + C_n)$, and after a carryover of C_{n-1}, the amount actually exported is $(S_{n-1} - C_{n-1}) = (X_{n-1} + C_n - C_{n-1})$. The optimal policy was determined for a finite time horizon ($n = 24$ seasons) using the value iteration algorithm (Hastings 1973). Due to the relatively short planning horizon, it is possible that the storage decision in the final period of the plan can bias the optimal policy. A non-zero carryover is prohibited in the last season of every set of simulated prices because any carryover in that season reduces net revenue. In a later section we present empirical evidence which suggests that any bias is probably small.

Assumptions

The major assumptions used in developing the inventory model are:

(a) The demand curve for Australian wheat exports is infinitely elastic at the world price.
(b) The area sown to wheat is fixed (at the 1968/69 level) and that the yield distribution is normal and stationary.
(c) The cost of storage function is constant over time.

The first assumption can be supported for the period before 1968/69 by the fact that Australia exported a small proportion of the commercial trade in wheat. Also, it is consistent with the models of price formation in the world wheat market developed by McCalla (1966) and Taplin (1969) in which prices were set by a co-operative duopoly involving the United States and Canada. In these models, prices were set along the residual demand curve facing the duopolists. This demand curve is obtained by subtracting the supply curves of the smaller producers (including Australia) from the world demand curve for wheat exports. The smaller exporters sold at the price, adjusted for quality, determined by the duopolists.

Elsewhere (Alaouze et al. 1978) we argued that the United States and Canada were (and are) concerned with their shares of the whole commercial export market, not just their shares of the residual market facing them. Under duopoly pricing, when the residual demand facing the United States and Canada contracted, their shares of the total com-

[2] In the model the level of deliveries to the AWB net of home consumption in the second season (X_{n-1}) is a random variable from the viewpoint of the first season.

mercial market would fall unless some compensating reduction in the exports of the smaller exporters (primarily Australia) were made. In the absence of voluntary export reductions on the part of Australia, market shares could be restored by price cutting. We argued that this was the explanation for the price war which occurred in the 1965/66 season. We therefore identify this season as one for which the assumption of an infinitely elastic demand for Australian wheat is not valid.

Furthermore, we argued that pricing in the world wheat market in the 1968/69 season and after could be explained by a market-share triopoly involving the United States, Canada and Australia. Under triopoly pricing, we showed that Australia faces an inelastic demand for wheat exports in excess of the cartel-determined share of the residual market facing the triopolists but, for exports below the cartel-determined share, the demand for Australian wheat is infinitely elastic at a price set by the price leader, usually Canada. The large wheat stocks held by the AWB in the 1968/69, 1969/70 and 1971/72 seasons indicate that Australia was facing an inelastic demand during this period. The price series simulated in this paper spanned the period 1949/50 to 1973/74. As indicated above, the assumption of infinitely elastic demand for Australian wheat cannot be supported for four of these seasons.

The assumption that the area sown to wheat is fixed was made because there are no apparent advantages in developing a comprehensive time series model for wheat sown in each season. As we show below, the occurrence of non-zero levels of carryover in the optimal policy is independent of the level of supply, and the optimal level of carryover is only affected in so far as the permissible range of carryover is affected by the probability distribution of production.[3] Statistical evidence is provided below to support the assumption that the yield distribution is stationary and normal.

The cost of storage function was assumed to be constant for the purposes of solving the model. This assumption was made because of the lack of a comprehensive time series on storage costs.

Analysis and simulation of the wheat price series

The price series which was analysed is the set of monthly 'basic selling prices' which are published by the AWB and the International Wheat Council. These prices were averaged over a crop-year to yield a series of annual wheat prices from 1949/50 to 1972/73.

This annual series has a mean of \$57.52 and a variance of \$140.40 and is composed of a long period of stable prices between two periods of high prices. The stable period lasted from 1953/54 to 1971/72. These observations have a mean of \$52.22 and a variance of \$6.75. A linear time trend was fitted to these observations; this was not signific-

[3] The permissible range of carryover is determined by the excess capacity of the grain storage system. This is partly determined by the probability distribution of production. The level of supply also constrains the permissible carryover when supply is less than the excess capacity.

ant at the 5 per cent level. In addition, these 19 observations were tested for serial independence using the von Neumann ratio and the serial correlation coefficient (Verhagen 1971, pp. 7-8). The more usual non-parametric tests are inapplicable for small sample sizes. The serial independence hypothesis was supported at the 5 per cent level by the von Neumann ratio test. The test based on the serial correlation coefficient also supported the hypothesis of independence.[4]

Little is known about the statistical properties of the episodic disruptions to the stable price series. For the purposes of simulation it was assumed that these events have a probability of 1/19 of occurring and that episodic prices are uniformly distributed in the interval {\$60, \$110}. The lower limit of this interval lies three standard deviations above the mean of the stable price series, and the upper limit is an estimate of the 1973/74 price.[5] The possibility of wheat prices being episodic in a downward direction was not considered because the historical evidence indicates that the stock-holding policies of the United States and Canada place a lower bound on wheat prices.

The simulation procedure involved generating normally distributed numbers with the same parameters as those of the stable period. This series was then interrupted by episodic events as described above. The episodic events were assumed to last three seasons. If the third of these uniformly distributed prices was greater than the second, the order was reversed in order to characterise the behaviour of prices declining from episodic peaks. In all, 20 sets of 24 simulated prices were used in the dynamic programming model.

The probability distribution of production and excess capacity

During the wheat harvest, storage and transportation facilities in Australia are used intensively in putting the crop under shelter in a short period. At this time storage capacity is critical because the carry-over from the previous season can interfere with the handling and storage of the current crop. Excess capacity can therefore be defined in terms of the maximum carryover that is permissible without interfering with the delivery of the current crop at a given level of probability. The excess capacity places an upper bound on the permissible level of carryover.

In this study it was assumed that a constant area is sown to wheat each year and that the total capacity of the storage system is constant. The excess capacity was determined by solving the following simple statistical inference problem:

We seek

(1) $$C': \quad p\{C' + X > (f+1)k\} = \beta$$

where C': maximum carryover or excess capacity,

[4] When the data are independent the distribution of a serial correlation coefficient tends to the normal, and is symmetrical, so that it is plausible to regard values outside a range of four standard deviations centred on the mean as indicating a departure from the hypothesis of independence. Details may be found in Verhagen (1971, pp. 7-8).

[5] This figure was a preliminary estimate provided by the AWB.

p: probability of the term in the braces,
X: quantity of deliveries to the wheat board,
k: capacity of the grain storage system,
f: proportion by which the total capacity can be exceeded without emergency storages being built.

The expansion of effective storage capacity is due to the storage of wheat in rail cars and in the holds of ships. The value of f for Victoria is about 3-4 per cent, and was assumed to be the same in all States (Victorian Grain Elevators Board, personal communication.)

The Australian average wheat yield from 1945/46 to 1968/69 has a mean of 1·148 t/ha and a variance of 0·049 t^2/ha^2.[6] A linear time trend was fitted to the Australian average yields. The regression coefficient was positive but not significant at the 5 per cent level. The yield data were also tested for normality and independence and both hypotheses were supported. It was assumed therefore that Australian average yields are stationary and normal.

Fixing the area planted at the record level of 1968/69, the mean and standard deviation of deliveries are 11·5 Mt and 2·4 Mt, respectively. These estimates include an allowance for farm retention of 92 kg/ha.

Assuming normality and using the Australian storage capacity in 1973/74 (20·5 Mt) as an estimate of k, the solution of (1) for $\beta = 0{\cdot}05$ and $f = 0{\cdot}03$ yields an estimate of 5·73 Mt for the excess capacity. This estimate is conservative in years when the wheat area is less than that of 1968/69 (10·8 million ha) and too high when a greater area is sown, other things being equal.

The probability distribution of export supply was obtained from the probability distribution of deliveries to the AWB by subtracting Australian consumption for the 1973/74 crop year from the mean of the probability distribution of deliveries to the AWB.[7] The range of four standard deviations on each side of the mean was split into twenty equal intervals. The mid-points of the intervals formed the vector of supply, and the area under the normal curve for each interval is the probability associated with each mid-point. The area under this range of the normal distribution accounts for 0·999 936 of the total area: the remaining area was added to the two intervals on either side of the mean.[8]

[6] Although average yields were available to 1972/73, the later observations were not used because it is possible that the introduction of wheat quotas in the 1969/70 season resulted in some substitution of land for other inputs in the seasons immediately following the introduction of quotas. Furthermore, the reduced wheat area per farm implied by the quotas meant that the land base on which wheat was grown was probably different from that of the preceding seasons. In particular, the expansion into marginal lands was arrested by the introduction of quotas.

[7] This approach assumes that consumption is fairly stable and that it is effectively estimated by AWB domestic sales in 1973/74. This was a year in which there was little 'over-the-border' trading because export prices were higher than domestic prices, hence most purchases were made from the AWB.

[8] Because of the iterative nature of the solution procedure, large errors can occur in the return associated with the optimal policy if the probabilities do not add precisely to one.

Storage costs and the cost of storage function

The maximum permissible level of carryover in the inventory model is the excess capacity of the Australian grain storage system (C') estimated above. This carryover was partitioned into 301 equally spaced amounts, ranging from zero to C'. The storage cost associated with each of these amounts of carryover determined the cost of storage function. The excess capacity is due to storage facilities constructed by the AWB to overcome the storage problem associated with the 'wheat crisis' of the late 'sixties. Because these storages are used in seasons of little carryover to facilitate the delivery of the wheat crop, we have attributed only the variable costs of storage to the carryover.

The variable cost of storage function was estimated for the 1973/74 season. Estimates of the 1973/74 average variable costs of storing wheat in commercial facilities were obtained by updating the estimates published by Freebairn (1967). His estimates were for each type of silo filled to capacity for New South Wales in 1965.

Assuming that New South Wales costs closely reflect those of other States, the cost of storage function was obtained by imposing the following order of filling on the grain elevator system:

(1) terminals,
(2) subterminals,
(3) country depots.

A stopover charge was imposed on wheat storage in the subterminals. Overall, estimated storage costs are probably higher than the 'true' variable costs of storage. The average variable costs of storage ranged between 20 c/t for the low levels of carryover to 60 c/t, depending on the quantity stored. For a detailed discussion of the calculation of the cost of storage function, see Alaouze (1975, pp. 38-45).

Restrictions of the carryover

Three restrictions are imposed on the carryover:

(a) the carryover must be non-negative,
(b) the carryover cannot exceed the quantity available for export,
(c) the carryover cannot exceed the excess capacity of the grain handling system.

These restraints can be written:

(2) $$0 \leqslant C_t \leqslant S_t, \quad C_t \leqslant C'.$$

These two restraints can be condensed by defining the function:

(3) $$S^*_t = \min\{S_t, C'\}.$$

The restrictions on the carryover can now be written:

(4) $$0 \leqslant C_t \leqslant S^*_t.$$

Formulation and Solution of the Model

The remaining variables and functions in the model are defined as follows:

X_t is wheat production net of farm retention and home consumption in a period with t stages of the process remaining. This variable is stochastic from the viewpoint of preceding stages in the process.

P_t is the wheat price in the period with t stages of the process remaining.

$K(C_t)$ is the cost of storage function. This is a monotonically increasing function of C_t.

$\theta_t(S_t)$ is a function which explicitly describes the way C_t depends upon S_t; that is $C_t = \theta_t(S_t)$. Providing $t \neq n$ (the first stage), this can also be written : $C_t = \theta_t(X_t + C_{t+1})$. Following Gustafson (1958), the function θ_t will be referred to as a storage rule.

W_t is the net revenue function associated with the export of $(S_t - C_t)$. This can be written: $W_t = P_t \,.\, (S_t - C_t) - K(C_t)$. Providing $t \neq n$, this can also be written: $W_t = P_t \,.\, (X_t + C_{t+1} - C_t) - K(C_t)$.

V_n is the return function associated with the n-stage process. This is the sum of the net revenues associated with every stage of the process discounted back to the first stage.

α is the discount factor.

The return function can be written:

(5) $$V_n = W_n + \alpha W_{n-1} + \alpha^2 W_{n-2} + \ldots + \alpha^{n-1} W_1.$$

Because the state variable is stochastic from the point of view of the first stage, we seek to maximise the expected value of the return function:

(6a) $$E[V_n] = E[W_n + \alpha W_{n-1} + \alpha^2 W_{n-2} + \ldots + \alpha^{n-1} W_1],$$

(6b) $$= W_n + \alpha E[W_{n-1} + \alpha^2 W_{n-2} + \ldots + \alpha^{n-2} W_1],$$

(6c) $$= W_n + \alpha E[V_{n-1}].$$

The problem can now be formally stated: we seek a set of storage rules, $\theta_n, \theta_{n-1}, \ldots, \theta_1$, (one rule for each stage of the process) which maximises the expected value of the return function. This set, denoted by $\theta'_n, \theta'_{n-1}, \ldots, \theta'_1$ is termed the optimal policy. Using equation (6c), the recurrence relationship associated with this problem can be written:

(7) $$E[V_n] = \begin{cases} W_n + \alpha E[V_{n-1}] & n \neq 1, \\ W_n & n = 1. \end{cases}$$

Denoting the maximised value of $E[V_n]$ as $E[V^*_n]$, the relationship we seek to maximise can be written:

(8) $$E[V^*_n] = \max_{0 \leqslant C_t \leqslant S^*_t} \{W_n + E[V_{n-1}]\} \quad (n \neq 1).$$

When the problem is expressed in the form of equation (8), the solution procedure implied by Bellman's principle of optimality can be used to find the optimal policy (Hastings 1973). In this case the multistage process associated with (8) has a finite time horizon. Therefore, the solution can be obtained using the value iteration algorithm (Hastings 1973).

We have mentioned that one of the important implications of the assumption that Australia is a price taker in the world wheat market is that the storage rules are independent of the level of supply. This result simplifies the solution procedure with a consequent saving in programming effort. Because there may be other dynamic programming problems for which similar simplifying assumptions can be made, we shall consider the solution procedure in some detail.

Bellman's principle of optimality states that, for all the possible values of the state variable which can result from the decision associated with the first stage of the process, the remaining decisions must be optimal. This implies that, for operational purposes, (8) may also be written:

(8a) $$E[V^*_n] = \max_{0 \leqslant C_t \leqslant S^*_t} \{W_n + \alpha E[V^*_{n-1}]\},$$

where $E[V^*_{n+1}]$ is the maximised discounted expected revenue associated with the $(n—1)$ stage process beginning with the second stage of the original (n) stage process.

The solution procedure for the process when there are t stages remaining involves finding the storage rule, $C^*_t = \theta'_t(S_t)$ which satisfies the recurrence relationship

(9) $$E[V^*_t] = \max_{0 \leqslant C_t \leqslant S^*_t} \{W_t + \alpha E[V^*_{t-1}]\}.$$

Substituting for W_t, (9) can also be written:

(9a) $$E[V^*_t] = \{P_t \,.\, (S_t - C^*_t) - K(C^*_t) + \alpha E[V^*_{t-1}]\},$$

or

(9b) $$E[V^*_t] = \{P_t \,.\, (X_t + C_{t+1} - C^*_t) - K(C^*_t) + \alpha E[V^*_{t-1}]\},$$

where C^*_t satisfies (9). That is C^*_t is obtained from the storage rule.

From the viewpoint of the preceding stage $(t + 1)$, X_t is a random variable, hence:

(10) $$E[V^*_t] = E[P_t \,.\, (X_t + C_{t+1} - C^*_t) - K(C^*_t) + \alpha E(V^*_{t-1})],$$
(10a) $$= P_t E[X_t] + P_t \,.\, (C_{t+1} - C^*_t) - K(C^*_t) + \alpha E[V^*_{t-1}],$$
(10b) $$= g_{t+1}\,(C_{t+1}).$$

That is, the expected maximised discounted net revenue for the t-stage process is dependent upon the decision (the levels of carryover) in the preceding stage (Gustafson 1958). From this result it follows that $E[V^*_{t-1}] = g_t(C_t)$; therefore (9) can be written:

(11) $$E[V^*_t] = \max_{0 \leqslant C_t \leqslant S^*_t} \{P_t \,.\, (S_t - C_t) - K(C_t) + \alpha g_t(C_t)\}.$$

In practice, the problem of finding the storage rule for each stage of the process (and hence the storage policy) is approached by solving (11) using discrete values of the various functions corresponding to selected finite values of their respective arguments. This is necessary because it is impossible to store all the possible values which the variables can take (Bellman and Kalaba 1965, p. 59).

Examining (11), we see that there are only two relevant variables, S_t and C_t. If we consider only l equally spaced points along the domain of S_t, and m equally spaced points along the domain of C_t (beginning with zero and ending with the excess capacity C'), representing all the possible values which these variables can take, that is, S_{it} for $i = 1, \ldots, l$ and C_{jt} for $j = 1, \ldots, m$, we can write (11) as follows:

$$(12) \quad E[V^*_{it}] = \max_{0 \leqslant C_{jt} \leqslant S^*_{it}} \{P_t \,.\, (S_{it} - C_{jt}) - K(C_{jt}) + \alpha g(C_{jt})\}.$$

The storage rule $C^*_{jt} = \theta'_t(S_{it})$ is then obtained by finding the l row maxima associated with the matrix:

$$(13) \quad \{P_t \,.\, (S_{it} - C_{jt}) - K(C_{jt}) + \alpha g(C_{jt})\},$$

subject to $0 \leqslant C_{jt} \leqslant S^*_{it}$ for $i = 1, \ldots, l$ and $j = 1, \ldots, m$. Each row corresponds to a value of the state variable, and each column to a value of the decision variable.

Our approach is to show that the row maximum is unique, and then we prove that it is independent of the row.

Consider the elements of the i-th row (corresponding to a value of the state variable S_{it}). Each element can be written as the sum of two components.

$$(14a) \quad T_{1ij} = P_t(S_{it} - C_{jt}) - K(C_{jt}),$$

$$(14b) \quad T_{2ij} = \alpha g(C_{jt}).$$

The cost function $K(C_{jt})$ increases monotonically with C_{jt}, hence (T_{1ij}) is a monotonically decreasing function of C_{jt}. The second term (T_{2ij}) is a monotonically increasing function of C_{jt}. Hence, as C_{jt} increases across a row (with j), each element can be written as the sum of two functions which vary monotonically with C_{jt}, and hence each element is unique.

It is interesting to note that when the row maximum corresponds to an interior maximum, that is $0 < C^*_{jt} < S^*_{it}$, the C^*_{jt} is a numerical approximation to the C_t which satisfies the following marginal condition in the continuous case:

$$(14c) \quad -\,\partial T_{1i}/\partial C_t = \partial T_{2i}/\partial C_t.$$

The C^*_t which satisfies (14c) equates the marginal opportunity cost of the carryover $-\partial T_{1i}/\partial C_t$ to the marginal maximised discounted expected value of the carryover $\partial T_{2i}/\partial C_t$. This is an economic interpretation of the solution procedure implied by Bellman's principle of optimality for inventory problems.

We shall now show that the row maxima of (13) are independent of the row and hence the value of the state variable.

For simplicity, we rewrite the matrix (13):

$$(15) \quad \{\gamma_i + \beta_j\}, \qquad i = 1, \ldots, l \text{ and } j = 1, \ldots, m,$$

where $\gamma_i = P_t \,.\, S_{it}$ and $\beta_j = -P_t \,.\, C_{jt} - K(C_{jt}) + \alpha g(C_{jt})$. Consider the i-th row of (15):

(16) $\{\gamma_i + \beta_1, \gamma_i + \beta_2, \ldots, \gamma_i + \beta_h, \ldots, \gamma_i + \beta_m\}$.

Let $\gamma_i + \beta_h$ be the row maximum, that is:

(17) $\max\{\gamma_i + \beta_1, \gamma_i + \beta_2, \ldots, \gamma_i + \beta_h, \ldots, \gamma_i + \beta_m\} = \gamma_i + \beta_h$.

Subtracting γ_i from both sides of (17), we obtain:

(18) $\max\{\beta_1, \beta_2, \ldots, \beta_h, \ldots, \beta_m\} = \beta_h$.

That is, we have shown that the row maximum is independent of the row. Therefore (12) can be written:

(19) $E[V^*_{it}] = P_t \,.\, (S_{it} - C_{ht}) - K(C_{ht}) + \alpha g(C_{ht})$

for all i subject to: $0 \leqslant C_{ht} \leqslant S^*_{it}$, where $S^*_{it} = \min\{S_{it}, C'\}$ as described above (equation (3)).

Let us summarise our progress so far. We have shown that the row maximum of (17) is independent of the row. The column in which the row maximum is found corresponds to a specific level of carryover C_{ht}. Although this carryover is associated with the row maximum, and by construction must be less than or equal to the excess capacity, it is not feasible for all values of S_{it} because negative storage is not permitted. When $C_{ht} \geqslant S_{it}$, the optimal level of carryover C^*_t, is equal to S_{it}. Hence, as written, equation (19) does not represent the maximised expected discounted net revenue associated with values of the state variable S_{it} when $S_{it} < C_{ht}$. For completeness, the equation corresponding to (19) for this case is:

(20) $E[V^*_{it}] = -K(S_{it}) + \alpha g(S_{it})$,

for all $S_{it} < C_{ht}$.

Equations (19) and (20) imply that the storage rule can be written as follows:

(21) $$C^*_t = \begin{cases} S_{it} & \text{for } C_{ht} > S^*_{it} \\ C_{ht} & \text{for } C_{ht} \leqslant S^*_{it}. \end{cases}$$

The three possible forms of storage rules implied by (21) are shown in Figure 1. The curve OAB corresponds to the case when the maximum optimal level of carryover, C_{ht}, is less than the excess capacity. The segment OA corresponds to the part of the rule where $C^*_t = S_t$, that is, carryover is constrained by the available supply; the segment AB corresponds to the case where the available supply, S_t, exceeds the maximum optimal carryover C_{ht}. This level of carryover corresponds to an interior maximum and approximately satisfies equation (14c). The curve OCD illustrates the case when the maximum optimal carryover is constrained by the excess capacity. This curve is composed of the following two boundary cases: the segment OC corresponds to $C^*_t = S_t$, and the segment CD corresponds to $C_h = C'$. The final case corresponds to a maximum carryover of zero, that is $C_h = 0$. In this case the curve coincides with the S_t axis.

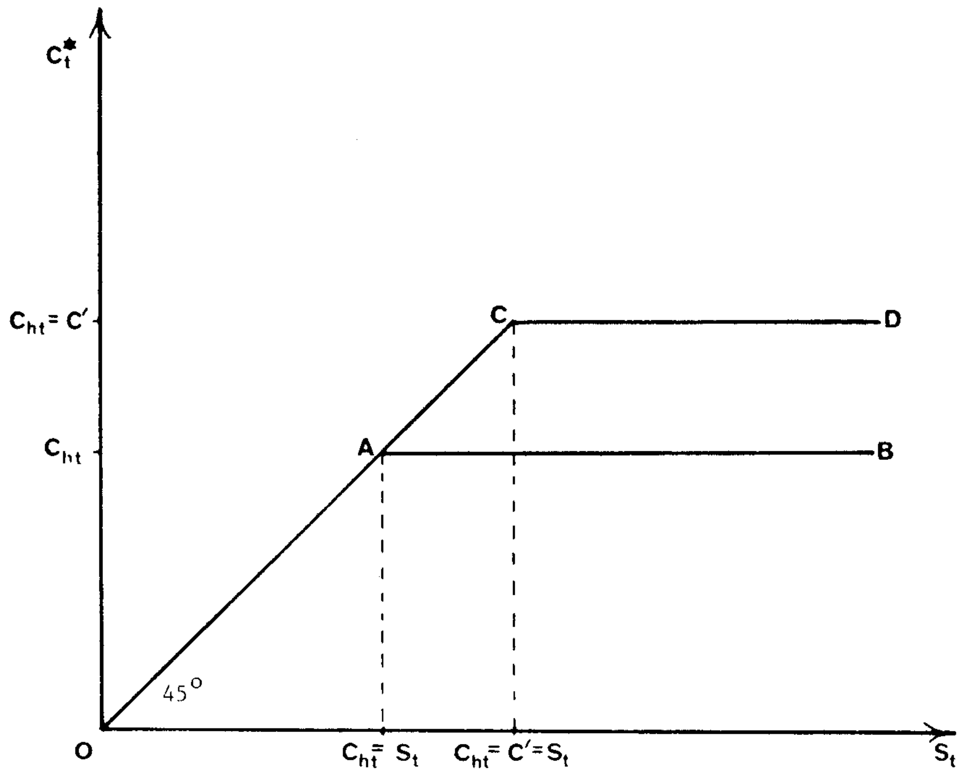

FIGURE 1—Forms of the storage rules.

The preceding results show that, under the assumption of infinitely elastic demand for Australian wheat exports, the solution algorithm involves finding only one row maximum for each stage of the process. The computational saving involved means that problems of much larger dimensionality may be solved approximately with ease and with a greater level of accuracy than would be feasible otherwise. An important implication of this is that the tradeoff between feasibility and realism will have to be considered carefully in larger inventory problems which do not quite satisfy the assumption of infinitely elastic demand.

Analysis

Preliminary discussion

The inventory model was analysed using twenty sets of simulated prices and discount factors of 1·0, 0·9524, 0·9091 and 0·8696, which correspond to interest rates of zero, 5 per cent, 10 per cent and 15 per cent, respectively. The interest rate used in inventory models for discounting purposes should reflect the opportunity cost of the money tied up in stock. In the context of this model, interest rates should reflect the fact that we have assumed stationary prices and costs. Therefore we should not include an allowance for inflation in the interest rate which should reflect only risk and the cost of time. A high discount factor would reflect a situation where future prices are considered highly uncertain. For the purposes of long-term policy, a discount factor reflecting both risk and time costs of 0·9524 may be appropriate. On the other hand, lower discount factors may be appropriate for some farmers who place a high value on liquidity.

Even though the storage costs used in the model are estimated from 1973/74 costs, the solutions for the 5 per cent interest rate should give a reasonable reflection of optimal inventory policy for the period which was simulated, because the storage costs are small compared with the interest cost. This is especially true for the costs associated with the smaller amounts of carryover.

Results

Two representative sets of simulated prices with their associated storage rules for the four interest rates are shown in Tables 1 and 2. Only the maximum carryover for each stage of the process is shown. Referring to Figure 1, the parts of the storage rules shown in these tables correspond to the horizontal segments of the diagram. The ex-

TABLE 1

Optimal Storage Rules for a Simulated Price Sequence (Run No. 2)

Season	Simulated prices $A Mean: 57·83 Variance: 237·92	Storage rules—maximum carryover (Mt)			
		Interest rate			
		0·0	0·05	0·10	0·15
1	55·86	0	0	0	0
2	52·28	1·91	0	0	0
3	52·82	0	0	0	0
4	52.64	0	0	0	0
5	51·75	5·73	0	0	0
6	54·17	0	0	0	0
7	53·28	0	0	0	0
8	53·04	5·73	5·40	0	0
9	56·75	0	0	0	0
10	55·02	0	0	0	0
11	50·46	5·42	0	0	0
12	51·85	5·42	0	0	0
13	53·33	0	0	0	0
14	52·64	0	0	0	0
15	47·98	5·73	5·73	0	0
16	52·39	5·73	5·42	0	0
17	56·17	0	0	0	0
18	48·43	5·42	0	0	0
19	49·92	0	0	0	0
20	45·96	5·73	5·73	5·73	5·73
21	101·49	5·73	0	0	0
22	103·43	0	0	0	0
23	84·35	0	0	0	0
24	51·95	0	0	0	0
Expected return to storage ($ million)		402·37	42·24	16·56	6·81

TABLE 2

Optimal Storage Rules for a Simulated Price Sequence (Run No. 18)

Season	Simulated prices $A Mean: 64·93 Variance: 466·45	Storage rules—maximum carryover (Mt) Interest rate			
		0·0	0·05	0·10	0·15
1	52·65	5·42	0	0	0
2	53·76	0	0	0	0
3	52·55	5·73	5·73	1·83	0
4	57·78	5·73	5·73	5·73	5·73
5	98·18	5·73	5·40	0	0
6	103·91	0	0	0	0
7	102·86	0	0	0	0
8	94·26	5·73	5·73	5·73	0
9	108·55	0	0	0	0
10	99·21	0	0	0	0
11	56·83	0	0	0	0
12	50·61	5·73	0	0	0
13	53·07	0	0	0	0
14	48·68	5·73	5·73	1·81	0
15	54·05	0	0	0	0
16	50·90	5·73	0	0	0
17	52·60	0	0	0	0
18	49·85	5·73	0	0	0
19	51·63	5·42	0	0	0
20	52·81	5·73	0	0	0
21	54·65	0	0	0	0
22	52·96	0	0	0	0
23	49·70	5·73	5·73	5·42	0
24	56·41	0	0	0	0
Expected return to storage ($ million)		461·42	232·66	144·41	100·34

cess capacity is 5·73 Mt and when this number appears in the tables, the storage rule requires the maximum possible to be stored. Storage of 5·73 Mt, and storage of less than 5·73 Mt corresponds to the segments *CD* and *AB* in Figure 1, respectively. The expected return to storage (Tables 1 and 2) is defined as the difference between the maximised discounted expected net revenue with optimal storage and the discounted expected net revenue when the carryover in each period is zero (Gustafson 1958, p. 55).

Tables 1 and 2 show that the optimal policy associated with each set of simulated prices is markedly affected by the choice of interest rate. For example, in Table 1 the optimal policy requires non-zero storage in 10 out of a possible 23 seasons at an interest rate of zero. Non-zero storage is required in four seasons at an interest rate of 5 per cent and one season at interest rates of 10 and 15 per cent. This pattern is also reflected in the expected return to storage associated with each interest rate.

We shall now discuss the pattern of storage rules associated with episodic price increases. In Table 1 these prices occur in seasons 21, 22 and 23. As expected, the optimal policy requires the maximum possible storage in season 20 at all interest rates. However, the optimal policy in season 19 required zero storage at all interest rates. This implies that the $4 drop in prices between seasons 19 and 20 was enough to offset any gain to storage from holding a carryover in season 19 in anticipation of a twofold increase in price in season 21. Because this effect was observed at a zero interest rate, it can be concluded that an important determinant of storage (apart from the interest rate) in any season is the price in the following season, with prices and storage in other future seasons having little effect. This suggests that any bias in the optimal policy due to compulsory zero storage in the final period of the simulation horizon is probably small.[9]

The pattern of storage in Table 2 illustrates the structure of the optimal policy when the simulated prices in the two seasons immediately preceding an episodic increase in price (season 5) form an increasing sequence. In this example (at the zero interest rate) we find that the price fall of $1.21 between seasons 2 and 3 is enough to offset any gains from storage in the second season, even though the sequence of prices for seasons 3, 4 and 5 is increasing, culminating with a high price in season 5. Storage in season 3 can be attributed to the $5 increase in price between seasons 3 and 4. This conclusion is reached because it is clear from the pattern of storage rules (at the zero interest rate) that a small positive price difference can result in storing the maximum possible, as can be seen from the storage in season 16 for the $1.70 price increase between seasons 16 and 17. The preceding discussion indicates that holding (speculative) carryover to service episodic increases in prices is in general not an optimising strategy except for the season immediately preceding an episodic price increase. The other conclusion that can be drawn is that a major determinant of storage (apart from interest rates) in a particular season is the price in the following season.

We shall now compare the policy of storing wheat in seasons of below average prices with the optimal policy when prices are stable. There were 110 seasons in which non-zero levels of storage were required by the optimal policies for the 5 per cent interest rate in the 20 sets of simulated prices. This number is small compared to the 460 simulated seasons for which storage rules can appear (after allowing for the fact that storage is prohibited in the last season of each set). These storage rules may be separated into:

(a) those not immediately preceding episodic price increases (76), and

(b) those immediately preceding episodic price increases (34).

[9] Because the price in season 20 (in Table 1) is less than that in season 19, any return to storage in season 19 occurs because the carryover from season 19 allows more wheat to be stored when supply is less than the excess capacity in season 20. In the model as structured, this phenomenon is not strong enough to offset the opportunity cost associated with the carryover.

TABLE 3

Summary of Results

Interest rate	0·0	0·05	0·10	0·15
Average expected return to storage ($ million)	555·57	161·55	80·48	47·69
Range of the expected return to storage ($ million)	150·61 to 699·19	29·19 to 293·98	0·93 to 237·54	0 to 214·00
Number of storage rules associated with episodic events	35	34	32	30
Number of storage rules not associated with episodic events	184	76	23	3
Number of storage rules not associated with episodic events which require the maximum possible carryover	68	32	8	0
Number of storage rules not associated with episodic events which require less than the maximum possible carryover	116	44	15	3

Of the storage rules which do not immediately precede episodic price increases, 65 occur in seasons with below average prices.

The number of seasons in which these storage rules appear is small compared to the 171 seasons for which below average prices prevailed. This does not include seasons in which below average prices occurred immediately before episodic price increases, nor does it include below average prices which occurred in the last season of each set of simulated prices. This indicates that, even if the AWB knew the mean of the price series with certainty and based its storage policy on the occurrence of below average prices, this policy would have been wrong more than half the time. Moreover, any rule of thumb approach to the storage problem which is based on deviations of prices from the mean or trend requires an accurate estimate of the mean or trend. For example, if the estimate of the mean is higher than the true mean, stocks will tend to accumulate indefinitely. This argument is similar to the one advanced by Gustafson (1958, p. 8) against basing storage rules on the deviation of crop size from normal.

Summary, Conclusions and Further Work

In this paper we developed an inventory model based on the simplifying assumption that the demand for Australian wheat exports is infinitely elastic at the world price. This model is valid for analysing the inventory problem for commodities which are traded under competitive conditions or where the exporting country is in a price-taking position, for example, a fringe-supplier in a cartel dominated market.

The major conclusions of this paper are as follows.

(a) Apart from interest rate, the most important factor affecting storage in any season is the price in the following season.

(b) The holding of a speculative reserve to be sold in seasons of episodic price increases (events which have a low frequency) is generally unwarranted.

(c) The optimal policies associated with simulations of the historical price series observed for the period 1953/54 to 1971/72 (when Australian wheat prices had a stable mean and a low variance) indicate that a storage policy based on storing wheat in seasons of below average prices would have been wrong more than half the time.

The model we have developed here has limitations in analysing Australian wheat storage. These arise because of the complex nature of the demand facing Australian exports. We do not know how our conclusions would change if we relaxed the assumption that the demand for Australian wheat exports is infinitely elastic at the world price; for this reason alone our results require cautious interpretation.

References

Agrawal, R. C. and Heady, E. O. (1972), *Operations Research Methods for Agricultural Decisions*, Iowa State University Press, Ames.

Alaouze, C. M. (1975), Australian wheat storage: an application of dynamic programming, unpublished M.Agr.Sc. thesis, University of Melbourne.

———, Watson, A. S. and Sturgess, N. H. (1978), 'Oligopoly pricing in the world wheat market', *American Journal of Agricultural Economics* 60 (2), 173-85.

Australian Wheat Board (1974), *Annual Report: Season 1972-73*.

Bellman, R. E. and Dreyfus, S. (1957), *Applied Dynamic Programming*, Princeton University Press.

Bellman, R. and Kalaba, R. (1965), *Dynamic Programming and Modern Control Theory*, Academic Press, New York.

Browning, S. C. (1970), Analysis of optimum carryover level for grains in relation to foreign aid requirements, unpublished Ph.D. thesis, University of Connecticut.

Freebairn, J. W. (1967), 'Wheat storage costs in New South Wales', *Review of Marketing and Agricultural Economics* 35(1), 9-42.

Gustafson, R. L. (1958), *Carryover Levels for Grains—A Method for Determining Amounts that are Optimal Under Specified Conditions*, USDA, Technical Bulletin 1178, U.S. Government Printing Office.

Hastings, N. (1973), *Dynamic Programming with Management Applications*, Butterworth, Seven Oaks, Kent.

Intriligator, M. D. (1971), *Mathematical Optimization and Economic Theory*, Prentice Hall, Englewood Cliffs.

Johnson, D. G. (1975), 'World agriculture, commodity policy, and price variability', *American Journal of Agricultural Economics* 57(5), 823-8.

McCalla, A. F. (1966), 'A duopoly model of world wheat pricing', *Journal of Farm Economics* 48(3), 711-27.

Taplin, J. H. (1969), Demand in the world wheat market and the export policies of the United States, Canada and Australia, unpublished Ph.D. thesis, Cornell University.

Throsby, C. D. (1964), 'Some dynamic programming models for farm management research', *Journal of Agricultural Economics* 16(1), 98-110.

Verhagen, A. M. W. (1971), *The Significance of Association in Two Dimentions*, Technical Paper No. 34, CSIRO Division of Mathematical Statistics, Melbourne.