The Annals of Probability
1974, Vol. 2, No. 1, 1-13

OPTIMAL STOPPING IN THE STOCK MARKET

BY DAVID GRIFFEATH AND J. LAURIE SNELL

Cornell University and Dartmouth College

A class of optimal stopping problems for conditioned random walk is discussed in terms of selling strategies for the stock market.

0. Introduction. The stochastic process Brownian motion was introduced in 1900 by L. Bachelier [1] as a model for fluctuation in the stock market, five years before Einstein's discussion of the process as a model for the motion of particles. There has recently been a renewed interest in stochastic models for the stock market; see for example the extensive collection of Granger and Morgenstern [6]. While it is doubtful that these studies have significantly increased the number of millionaires, they have led to interesting mathematical problems.

One such problem, introduced by Boyce [2], may be motivated as follows: Let us adopt Brownian motion as a model for stock prices. This seems to be a reasonable hypothesis according to the review articles of Cootner [3] and Samuelson [10], and also to Samuelson's paper [9]. Suppose now that we want to stop at some point within a prescribed period of time so as to maximize our expected profit. Since Brownian motion is a martingale (i.e. a fair game), we cannot guarantee a positive expected payoff under optimal stopping. Now Boyce raised the question of optimal stopping given a "prediction distribution" which was not normal with mean 0 and variance 1, but was some other normal distribution $\mathfrak{N}(\mu, \sigma^2)$. A number of rather striking results follow from this model. First, even in the face of a negative predicted trend, $\mu < 0$, it is possible to have a positive expected gain under optimal stopping. That is, one can profit from information even if it is bad news. Secondly, the optimal strategy for stopping turns out to be "sell on rallies, ride out storms" if the prediction is reliable ($\sigma^2 < 1$) and "cut losses, let profits run" if $\sigma^2 > 1$. Boyce [2] exhibited these results by computational methods using a discrete approximation to Brownian motion. Recently Föllmer [5] derived the same results working directly with Brownian motion. Related stopping problems have been studied by Shepp [11].

The purpose of this paper is to show that the results described above for Brownian motion can be demonstrated in terms of the simpler process of discrete random walk. At the same time we wish to show the use of maximum

Received February 22, 1972; revised July 5, 1973.

[1] The preparation of this paper was aided by the National Science Foundation.

AMS 1970 *subject* classifications. Primary G2L15; Secondary: 60G40, 60J15.

Key words and phrases. Optimal stopping, random walk, maximum entropy, conditional Markov chains, stock market.

entropy methods in giving the stock market problem and related prediction problems a self-contained discrete formulation.

Thus for our discrete model we assume that a man holds a unit of stock which he must sell sometime over a period of N days. For simplicity we assume that the stock price increases or decreases by one unit each day. Finally, we assume the *validity* of some predictive information abut the stock price after N days. We wish then to determine how to sell optimally taking account of this information. To obtain a convenient probability measure for the stock market process we will follow a method suggested by Jaynes [8]; namely, we choose a measure which maximizes uncertainty as measured by entropy and is consistent with our given information. For example, if we have no information this leads to a simple random walk with probability $\frac{1}{2}$ of an increase and $\frac{1}{2}$ of a decrease on each day. If we are given a predicted mean and variance of the price after N days, we shall obtain a process for which the price at time N is distributed according to a kind of discrete analogue of the normal distribution. We will also be able to describe the process using the notion of conditional Markov chains, introduced recently in the study of Markov processes and potential theory. We begin with a discussion of these two basic techniques for modifying a stochastic process.

1. Conditional Markov chains. Let Ω be a finite sample space with an assigned probability measure $\mathbf{P}$. If we are given information that an event A occurs, conditional probability in effect leads to a new measure $\hat{\mathbf{P}}$ taking A to be the sample space, while retaining the same relative probabilities of subevents of A as given by $\mathbf{P}$. In particular this is true for the individual outcomes, or atoms, of A. Each such outcome ω of A is assigned a new measure $\hat{\mathbf{P}}(\omega) = \mathbf{P}(\omega)/\mathbf{P}(A)$. More generally, if $\{A_1, A_2, \cdots, A_k\}$ is a partition of Ω; and if rather than knowing that an outcome is in one of the events of the partition we learn information which causes us to change the measure assigned to these events to $P(A_1)$, $\hat{\mathbf{P}}(A_2), \cdots, \hat{\mathbf{P}}(A_k)$ respectively, then as in the case of conditional probability we wish the new measure to keep the same relative probabilities for the outcomes within each event of the partition. This is achieved by assigning to each atom $\omega \in A_r$ the measure $\hat{\mathbf{P}}(\omega) = \mathbf{P}(\omega)\hat{\mathbf{P}}(A_r)/\mathbf{P}(A_r)$. This is called a *conditional probability measure*.

Next consider a finite Markov chain with transition matrix $P = \{P_{ij}\}$. This chain is said to be *absorbing* if it has at least one absorbing state, or trap, and if from every other state it is possible to reach an absorbing state. If an absorbing chain is started in state i, it will eventually reach an absorbing state a almost surely. We denote by ν_{ia_r} the probability, starting in i, that the process ends in a_r.

A function h on the state space of a Markov chain is said to be *regular* if $h = Ph$ (writing h as a column vector), and *superregular* if $h \geqq Ph$.

Suppose now that we have a finite absorbing Markov chain which starts in state 0 and ends at one of the absorbing states $a_1, a_2, \cdots, a_k$. Let A_r denote the set of all finite paths leading from 0 to a_r, and let $\omega = (0, i_1, \cdots, a_r)$ be a typical such path. Then except for a set of measure 0, $\{A_1, A_2, \cdots, A_k\}$ is a partition of the probability space Ω of all paths of the chain. The original Markov chain assigns a probability to each event of this partition. Assume now that a prediction causes us to change these probabilities from $\mathbf{P}(A_r)$ to $\hat{\mathbf{P}}(A_r)$. Then, in accordance with the above discussion, we change the measure of ω from

$$\mathbf{P}(\omega) = P_{0i_1} P_{i_1 i_2} \cdot \cdots \cdot P_{i_s a_r}$$

to

$$\hat{\mathbf{P}}(\omega) = [P_{0i_1} P_{i_1 i_2} \cdot \cdots \cdot P_{i_s a_r}]\hat{\mathbf{P}}(A_r)/\mathbf{P}(A_r) . \tag{1}$$

Conveniently, the distribution (1) can be achieved by defining a new Markov chain as follows: First observe that if h is any nonnegative regular function, then the matrix $\hat{P}$ defined by $\hat{P}_{ij} = P_{ij} h(j)/h(i)$ is again a transition matrix. The resulting Markov chain is called an *h-chain*. For such chains, the probability of a sequence ω becomes

$$\begin{aligned}\hat{\mathbf{P}}(\omega) &= P_{0i_1} h(i_i)/h(0) \cdot P_{i_1 i_2} h(i_2)/h(i_1) \cdot \cdots \cdot P_{i_s a_r} h(a_r)/h(i_s) \\ &= [P_{0i_1} P_{i_1 i_2} \cdot \cdots \cdot P_{i_s a_r}] h(a_r)/h(0) .\end{aligned}$$

We now choose an appropriate h. Note first that for any a_r, $h(i) = \nu_{ia_r}$ is a regular function. The same is true of the function $h(i) = \nu_{ia_r}\hat{\mathbf{P}}(A_r)/\mathbf{P}(A_r)$ since multiples of regular functions are regular. Finally, we obtain the desired h by summing,

$$h(i) = \textstyle\sum_r \nu_{ia_r}\hat{\mathbf{P}}(A_r)/\mathbf{P}(A_r) . \tag{2}$$

For this choice $h(0) = 1$ and $h(a_r) = \hat{\mathbf{P}}(A_r)/\mathbf{P}(A_r)$, so that $\hat{\mathbf{P}}(\omega)$ takes on the form (1) as desired. Therefore the predicted chain becomes an *h-chain* with h as given by (2).

2. Optimal stopping for a space-time chain. With a given Markov chain it is often convenient to associate a new process called the *space-time chain*. For this chain we take the states to be pairs (m, i), where i is a state of the original chain and m represents time, so that the transition matrix for the new chain is given by

$$P_{(m,i)(m+1,j)} = P_{ij} .$$

According to this formulation the process is in a state at most once.

Suppose now that we have a space-time chain which we watch for a finite number N steps, the states (N, r) being made absorbing, and also that we have a function $\pi(m, i)$ defined on the state space which designates our payoff if the process is in state (m, i) and we decide to stop. For such a process, by standard dynamic programming techniques it follows that the *value* $g(m, i)$ of being in state (m, i) and being allowed optimal stopping satisfies the boundary condition

$$g(N, r) = \pi(N, r) \tag{3}$$

and the recursion equation

$$(4) \qquad g(m, i) = \max\{\pi(m, i), \textstyle\sum_j P_{(m,i)(m+1,j)} g(m+1, j)\} \qquad \text{for } m < N .$$

An optimal stopping rule consists in stopping the first time $g(m; i) = \pi(m, i)$. The value function g may also be characterized as the least superregular function dominating the payoff function π (see [4]).

3. The conditioned stock market problem. We now formulate the stock market problem in terms of a space-time Markov chain. In the notation of Section 2, the simple random walk model for an N-day stock market process starting at (0, 0) is described by the transition matrix P with

$$P_{(N,r)(N,r)} = 1 ; \qquad \text{and} \qquad P_{(m,i)(m+1,i+1)} = P_{(m,1)(m+1,i-1)} = \tfrac{1}{2} \qquad \text{for } m < N .$$

The absorption probabilities are

$$\nu_{(0,0)(N,r)} = \binom{N}{(N+r)/2} \Big/ 2^N \qquad (r = -N, -N+2, \cdots, N-2, N) .$$

The effect of prediction is to assign a new set of absorption probabilities $\hat{p}_r = \hat{\nu}_{(0,0)(N,r)}$ to the stock process. We have seen that this can be done by forming an h-chain according to (2), so that

$$h(m, i) = \textstyle\sum_r \nu_{(m,i)(N,r)} \hat{p}_r / \nu_{(0,0)(N,r)} .$$

The values $\nu_{(m,i)(N,r)}$ are easily computed as binomial probabilities, and hence it is a simple matter to determine h. The new transition matrix for the h-process is then

$$(5) \qquad \begin{aligned} \hat{P}_{(m,i)(m+1,i+1)} &= h(m+1, i+1)/2h(m, i) \\ \hat{P}_{(m,i)(m+1,i-1)} &= h(m+1, i-1)/2h(m, i) \qquad \text{for } m < N, \end{aligned}$$

with, of course, $\hat{P}_{(N,r)(N,r)} = 1$.

We shall denote the process which represents the daily price by $V = \{V_m ; 0 \leqq m \leqq N\}$ with V_0 normalized to 0, as we are only interested in net gains or losses.

4. Maximum entropy measures. Consider a finite probability space $(\Omega, \mathbf{P})$ with elements ω. Suppose we wish to assign a probability measure to Ω on the basis of partial information about that measure. Suppose in fact that we only know the expected values of certain random variables $X_0, X_1, \cdots, X_k$ defined on Ω:

$$\mathbf{E}(X_0) = \mu_0 ; \quad \mathbf{E}(X_1) = \mu_1 ; \quad \cdots ; \quad \mathbf{E}(X_k) = \mu_k .$$

A standard information theoretic approach is to assign to Ω the measure which has maximum uncertainty while satisfying the given information as expressed by constraints of the above form. Uncertainty is measured by the *entropy* of $\mathbf{P}$:

$$H(\mathbf{P}) = -\textstyle\sum_\omega \mathbf{P}(\omega) \ln \mathbf{P}(\omega) .$$

The maximum entropy measure $\bar{\mathbf{P}}$ is typically unique and of the form

$$\bar{\mathbf{P}}(\omega) = c\lambda_0^{X_0(\omega)}\lambda_1^{X_1(\omega)} \cdot \cdots \cdot \lambda_k^{X_k(\omega)} ,$$

where $c, \lambda_0, \cdots, \lambda_k$ are constants determined by the given constraints and the additional condition that the $\bar{\mathbf{P}}(\omega)$ sum to 1. See [8], [7], for a discussion of this problem.

If there are no constraints except that $\bar{\mathbf{P}}$ should be a probability measure, then the maximum entropy solution assigns equal weight to all outcomes. If the possible outcomes ω are integers and the only constraint consists in giving the expected value of the outcome, then the maximum entropy solution is

$$\bar{\mathbf{P}}(r) = c\lambda_1^r ,$$

a geometric distribution. If the first two moments are given, then the solution has the form

$$\bar{\mathbf{P}}(r) = c\lambda_1^r\lambda_2^{(r^2)} . \tag{6}$$

In this latter case we have a discrete analogue of the normal distribution. For continuous distributions with densities, the normal solution maximizes entropy when the first two moments, or equivalently the mean and variance, are known.

We now apply these ideas to the stock market problem. The basic probability space is the space of sequences $\omega = (\omega_1, \omega_2, \cdots, \omega_N)$, where ω_m is 1 or -1 depending on whether the stock rises or drops one unit on the mth day. Thus the corresponding price $V_m = \sum_{s=1}^{m} \omega_s$. In the absence of any information the maximum entropy method suggests that we assign an equal weight to each possible sequence ω. This leads to the binomial distribution for V_N, and our a priori simple random walk measure $\mathbf{P}$ on Ω.

Suppose now that a new distribution for V_N is predicted. Then we can write the new distribution as a set of predictions in the form $\mathbf{E}(X_r) = \mu_r$, where $X_r = 1_{A_r}$ is the indicator function of $A_r = \{\omega\colon V_N(\omega) = r\}$, and $\mu_r = \hat{p}_r$ is the predicted probability of absorption at (N, r). (V_m is defined, according to our model, for both positive and negative prices i.) By applying the method of maximum entropy to the path space Ω we therefore obtain

$$\bar{\mathbf{P}}(\omega) = \hat{p}_r \Big/ \binom{N}{(N+r)/2} \qquad \text{for } \omega \in A_r \tag{7}$$

(assuming, for convenience, that N is even). In other words, the maximum entropy measure assigns the same probability $\bar{\mathbf{P}}(\omega)$ to all sequences in A_r for fixed r. Looking back at the way we assigned a measure by means of the h-chain, we see that the same is true. Hence both the maximum entropy method and the method of h-chains lead to the same basic measure to take into account the prediction of our example.

This will not be true for a general absorbing Markov chain. Moreover, the maximum entropy theory is not available for the continuous time model as the

h-process theory is. The maximum entropy method can, however, be used in certain problems where the h-chain theory does not suffice. For example, suppose that only the mean and variance for the price of our stock after N days are predicted. Then there are many distributions which agree with these predictions, giving different h-chains and corresponding measures on the path space Ω. The principle of maximum entropy provides a rationale for choosing one particular measure, which in turn determines a specific h-chain to which we may apply our optimal stopping results. We shall use this technique in the next section, where we treat the optimal stopping problem for the stock market example.

5. Solution of the stock market problem. We are now prepared to give a discrete formulation of the stock market problem of Boyce [2]. Following the methods and notation of the preceding sections, consider an N-day stock market process with states (m, i) where m denotes the day and i the relative price on that day. As a matter of notational convenience we will always take N *even*; only trivial differences arise in the analysis for N odd. Let the predicted $\hat{\nu}_{(0,0)(N,r)}$ be denoted $\hat{p}_r$, and the corresponding final h-value h_r $(=h(N, r))$. The h-value at state (m, i) will be written $h(m, i)$, the optimal value $g(m, i)$ and the conditioned transition probability from (m, i) to $(m + 1, j)$—where $j = i \pm 1$—as $\hat{P}_{(m,i)(m+1,j)}$. The payoff $\pi(m, i)$ if we choose to sell our stock on the mth day is simply i, the price of the stock at that time.

For any prediction probabilities whatever, the optimal strategy may be obtained by the backward iteration method of Section 2, making use of (3) and (4). Specifically, $g(N, i) = i$, and for $m < N$,

$$g(m, i) = \max\{i, \hat{P}_{(m,i)(m+1,i+1)}g(m + 1, i + 1) + \hat{P}_{(m,i)(m+1,i-1)}g(m + 1, i - 1)\}\,. \tag{8}$$

An optimal strategy consists in stopping (selling our stock) if i is greater than or equal to the second term of the max, and continuing otherwise. Of course, from (8) we have $g(m, i) \geqq i$ for all (m, i) since we can always sell immediately. It follows that a sufficient condition for continuing at (m, i) is

$$\hat{P}_{(m,1)(m+1,i+1)} > \tfrac{1}{2} > \hat{P}_{(m,i)(m+1,i-1)} \qquad (m \leqq N - 1)\,,$$

or equivalently by (5),

$$h(m + 1, i + 1) > h(m + 1, i - 1) \qquad (m \leqq N - 1)\,. \tag{9}$$

If we stop at $(m + 1, i + 1)$ and $(m + 1, i - 1)$ so that $g(m + 1, i + 1) = i + 1$ and $g(m + 1, i - 1) = i - 1$ (in particular, if $m = N - 1$), then the above conditions are also *necessary* for continuing at (m, i) when we adopt the optimal strategy just described.

Arbitrary prediction distributions will in general lead to highly complex strategies, so that if a family of "nice" strategies is desired, some restriction must be placed on the $\hat{p}_r$. One such restriction, which we believe to be an appropriate analogue of the continuous time normal formulation of [2], is that the prediction probabilities *for the path space* Ω be a "discrete normal" distribution of given mean and variance, determined by the method of maximum entropy described in Section 4. More precisely, for each $\omega \in A_r$ choose $\bar{\mathbf{P}}(\omega)$ to have the measure $\bar{\mathbf{P}}(r)$ of (6). Then by (7), the prediction probability p_r of ending at price r is given by

$$\hat{p}_r = \binom{N}{(N+r)/2} \bar{\mathbf{P}}(\omega) = c \binom{N}{(N+r)/2} \lambda_1{}^r \lambda_2{}^{(r^2)} , \tag{10}$$

with

$$h_r = c 2^N \lambda_1{}^r \lambda_2{}^{(r^2)} , \tag{11}$$

where c, λ_1 and λ_2 are uniquely determined by the mean μ and variance σ^2 (and the fact that the $\hat{p}_r$ constitute a probability distribution).

This formulation gives the stock market problem a self sufficient discrete version, whereas that of Boyce relies on approximation to the continuous normal distribution. It is worth noting that while our discrete model is not that of Boyce, it does agree with his in the extreme case where the predictor designates with certainty the outcome after N days. In this case the process reduces to the following simple model studied by Shepp [11]: An urn contains R red balls and B blue balls. You are allowed to draw without replacement as long as there are any balls in the urn. Each time you draw a red ball you receive one dollar and each time you draw a blue ball you lose one dollar. As Shepp and Boyce have shown, finding the optimal value and strategy, even for this simpler process, provides challenging problems.

Let us consider, then, prediction distributions of the form (10), which we shall refer to in the sequel as *discrete normal* (λ_1, λ_2)*-distributions*. For the corresponding h-process, note that by (5) only the ratio of h-values determines the transition probabilities, so that in (11) the normalizing constant c and the factor of 2^N are irrelevant to the determination of g. It is the relation between μ and σ^2, or alternatively between λ_1 and λ_2, which determines the nature of the optimal strategy, and hence it is this relation in which we are particularly interested. We will soon restrict our attention to the sub-family of mean-0 discrete normal distributions, where the essential features of the stock market problem are most readily demonstrated. In the continuous model Föllmer [5] has shown that the non-mean-0 situation can be represented as a mean-0 problem with a "drift," while in our case the interrelation of discrete normal stock market problems is indicated by the equivalence of the N day (λ_1, λ_2)-process starting at (m, i) and an $N - m$ day $(\tilde{\lambda}_1, \lambda_2)$-process starting at $(0, 0)$, where $\tilde{\lambda}_1 = \lambda_1 \lambda_2{}^{2i}$. Before limiting ourselves to the mean-0 situation, however, we mention a result which may be considered analogous to one obtained in the continuous case by Shepp, using Brownian motion transformations.

PROPOSITION 1. *If the prediction distribution* $\{\hat{p}_r\}$ *is binomial, then an optimal* *selling strategy for the stock market problem consists in holding our stock until day N for $\mu > 0$, and selling immediately if $\mu \leqq 0$.*

PROOF. If the $\hat{p}_r$ are binomially distributed, it is easily verified that they are $(\lambda_1, 1)$-discrete normal. For $\mu > 0$ we have $\lambda_1 > 1$; for $\mu = 0$, $\lambda_1 = 1$ and $\hat{\mathbf{P}}$ reduces to the a priori measure $\mathbf{P}$; and if $\mu < 0$, $\lambda_1 < 1$. Hence according to (11) the h_r are monotone increasing as r increases for $\mu > 0$, and monotone decreasing for $\mu < 0$. Now, in (5) note that

$$h(m, i) = [h(m + 1, i + 1) + h(m + 1, i - 1)]/2 .$$

By backward induction, this averaging property preserves the final ordering for all $h(m, i)$ with m fixed. Hence according to condition (9) if $\mu > 0$ an optimal strategy consists in continuing at *all* (m, i) with $m < N$. For $\mu \leqq 0$, on the other hand, (9) is violated at *all* (m, i), whence a backward induction shows that we stop at every state.

For the remainder of this paper we will consider only the mean-0, or equivalently, the $(1, \lambda_2)$-stock market problem. As will be seen, the characterization of optimal strategies in this case depends fundamentally on the size of σ^2. For $\sigma^2 < N$ (a "reliable" prediction) we play when our stock is down and sell when it is up, while for $\sigma^2 > N$ (an "unreliable" prediction) we play when the stock is up and sell when it is down. Some insight into the nature of this "flip" is obtained by considering the extremal mean-0 distributions: $\sigma^2 = 0$ and $\sigma^2 = N^2$. In the former case, where we know with certainty that we will end at $(N, 0)$, it is clearly unwise to sell a stock which is down, and advisable to sell at certain states where it is ahead (since it will eventually drop to 0). For $\sigma^2 = N^2$, where we end at (N, N) with probability $\frac{1}{2}$ and $(N, -N)$ with probability $\frac{1}{2}$, we should sell immediately if the stock drops on the first day since it will certainly continue dropping to $(N, -N)$. But if the stock gains on the first day we should continue playing since it will with certainty climb to (N, N). In order to see that the qualitative flip suggested here occurs at $\sigma^2 = N$, and to describe more fully the optimal strategy governing discrete normal distributions, it is necessary to take a closer look at the underlying h-process.

Consider, then, the $(1, \lambda_2)$ stock market problem with predicted probabilities and h-process given by

$$\hat{p}_r = c\binom{N}{(N + r)/2}\lambda_2^{(r^2)} \tag{12}$$

and

$$h_r = c2^N\lambda_2^{(r^2)} , \tag{13}$$

where $0 \leqq \lambda_2 \leqq \infty$, and λ_2 satisfies

$$\theta(\lambda_2) \equiv \frac{\sum_r r^2\binom{N}{(N + r)/2}\lambda_2^{(r^2)}}{\sum_r \binom{N}{(N + r)/2}\lambda_2^{(r^2)}} = \sigma^2 \tag{14}$$

$$(r = -N, -N + 2, \cdots, N - 2, N) .$$

In (12) we see that $\hat{p}_r = \hat{p}_{-r}$ so that μ is clearly 0. Also, since the a priori binomial distribution for V_m and the $\hat{p}_r$ are both symmetric about 0, it follows that $\hat{P}_{(m,i)(m+1,j)} = \hat{P}_{(m,-i)(m+1,-j)}$ and that $h(m, i) = h(m, -i)$.

Observe in (14) that $\theta(\lambda_2)$ is monotone increasing on $[0, \infty)$ ($\theta' \geqq 0$), with $\theta(0) = 0$, $\theta(1) = N$, and $\lim_{\lambda_2 \to \infty} \theta(\lambda_2) = N^2$. It follows that λ_2 is uniquely defined in terms of σ^2, with

$$(15) \qquad \begin{aligned} \lambda_2 &< 1 \quad \text{for} \quad \sigma^2 < N, \\ \lambda_2 &= 1 \quad \text{for} \quad \sigma^2 = N, \\ \lambda_2 &> 1 \quad \text{for} \quad \sigma^2 > N, \end{aligned}$$

Now using (13), for the $(1, \lambda_2)$-process at $m = N - 1$, condition (9) becomes

$$\lambda_2^{(i+1)^2} > \lambda_2^{(i-1)^2},$$

which by (15) holds for $i < 0$ when $\sigma^2 < N$ and $i > 0$ when $\sigma^2 > N$. (For $\sigma^2 = N$, $\hat{\mathbf{P}}$ reduces to the uniform measure $\mathbf{P}$ on Ω, and our process has optimal value 0).Since, as already noted, (9) is necessary and sufficient for continuing at $(N - 1, i)$, it follows that

(a) for $\sigma^2 < N$ we stop at $i > 0$, continue at $i < 0$;
(b) for $\sigma^2 > N$ we continue at $i < 0$, stop at $i > 0$.

We proceed to show that for $\sigma^2 < N$ we play at *all* (m, i) where $i < 0$ $(m < N)$. According to (9), it is enough to show that

$$h(m, i + 2) > h(m, i) \qquad \text{for all } m \text{ and all } i\colon -2 \geqq i \geqq -m.$$

We know that $h_{r+3} > h_{r+1} > h_{r-1}$ for $r \leqq -3$, and by the averaging property of the h-process, each preceding h-value is an average of the h-values at that time. Also, $h(m, -1) = h(m, 1)$ for m odd; $h(m, 0) = [h(m + 1, -1) + h(m + 1, 1)]/2 = h(m + 1, -1)$ for m even. By backward induction, the order is therefore preserved. A strictly analogous argument establishes that for $\sigma^2 > N$ we continue to play at all (m, i) where $i > 0$. Finally, note that for $\sigma^2 < N$, $g(m, -1) > -1$ and $g(m, 1) \geqq 1$ whereas $\hat{P}_{(m,0)(m+1,-1)} = \hat{P}_{(m,0)(m+1,1)} = \frac{1}{2}$, so that by continuing at $i = 0$ we guarantee $g(m, 0) > 0$. Similarly, for $\sigma^2 > N$ we continue at $(m, 0)$.

The flip in optimal strategies just demonstrated is illustrated in Figs. 1 and 2. We summarize our results in the form of a proposition:

PROPOSITION 2. *According to the optimal strategy for the mean-0 stock market problem, we continue to play at all (m, i) where $m < N$ and $i \leqq 0$ if $\sigma^2 > N$, and we continue to play at all (m, i) where $m < N$ and $i \geqq 0$ for $\sigma^2 > N$.*

Optimal strategy description for $i > 0$ when $\sigma^2 < N$, and for $i < 0$ when $\sigma^2 > N$, is a far more difficult problem. No tractable explicit characterization for either the discrete or continuous model is known at this time; only approximate and asymptotic results are available. As already mentioned, any particular finite stock scheme is soluable by the backward iteration method.

For $\sigma^2 < N$, by means of an involved and not particularly enlightening induction argument, we are able to show that the mean-0 discrete normal optimal selling strategy is described by a stopping curve $s = \{(m, s_m)\}$ $(m < N)$ below which we continue, and at (or above) which we sell. Moreover,

$$s_{m-1} = s_m \pm 1 \quad \text{and} \quad s_{m-2} = s_m \quad \text{or} \quad s_m + 2\,,$$

so that s is "almost monotone increasing" from $m = N$ backwards. The stopping curve s is shown for $N = 10$, $\sigma^2 = 8$ in Fig. 1; an optimal strategy starting at (0, 0) consists in holding onto our stock until we reach s or $m = N$. Unfortunately, the induction argument does not extend to $\sigma^2 > N$, and simulations suggest that s may not always be monotone decreasing from N back for this case. Being tedious and of limited applicability, the proof of our partial result is therefore omitted.

A few more remarks regarding the optimal strategies and stopping curves are of interest. First, note that although the mean-0 predictions given in Figs. 1 and 2 are rather close to the martingale case where $\sigma^2 = N$, the optimal strategy ensures a mean profit of about 15 percent over the ten day period. For more

FIG. 1.

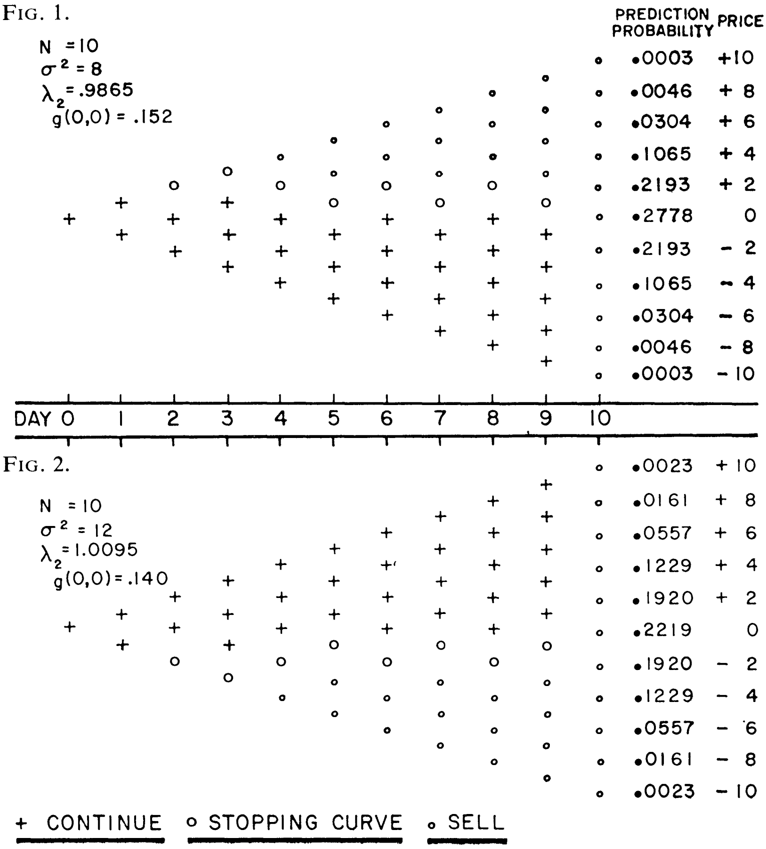

FIG. 2.

peaked distributions, i.e. for σ^2 very large or small compared to N, by adopting an optimal strategy we may achieve significant expected gains in cases where $\mu < 0$. As an extreme but illustrative 10-day process, consider the limiting case where our stock climbs to (10, 10) with probability $\frac{2}{5}$ and drops to (10, -10) with probability $\frac{3}{5}$; for this example we have $\mu = -2$, but

$$g(0, 0) = \tfrac{2}{5}(10) + \tfrac{3}{5}(-1) = 3.4 .$$

Needless to say less deterministic instances can be constructed, yielding the perhaps paradoxical result that, according to our model, it is even possible to profit from a declining market.

Simulations suggest that as the variance of a discrete normal distribution increases from σ^2 to σ'^2, where $0 \leqq \sigma^2 < \sigma'^2 < N$, the associated stopping curve decreases in the sense that $s_m(\sigma'^2) \leqq s_m(\sigma^2)$ for all m. In [2] Boyce conjectures (for the continuous model) that s does not decrease to $0+$ for σ^2 approaching N as we might expect, but rather to some limiting curve. We are able to verify this for the discrete model, by means of our final result.

PROPOSITION 3. *Adopting an optimal strategy for $\sigma^2 < N$, the last day on which we continue with $V_m = 1$ is $N - k - 1$, where k is the smallest even integer such that*

$$(16) \qquad (h_k - h_{k+2}) - \sum_r \frac{r(k-r-1)}{2k-r} \binom{k}{(r+2)/2} (h_{k-r-2} - h_{k-r}) < 0 ,$$
$$(r = 2, 4, \cdots, k-2;\ k \geqq 4) .$$

PROOF. A necessary and sufficient condition for continuing at $(N-k-1, 1)$ as given is $\hat{P}_{(N-k-1,1)(N-k,2)} \cdot 2 + \hat{P}_{(N-k-1,1)(N-k,0)} g(N-k, 0) > 1$, or equivalently,

$$(17) \qquad \hat{P}_{(N-k-1,1)(N-k,2)} + \hat{P}_{(N-k-1,1)(N-k,0)}[g(N-k, 0) - 1] > 0 .$$

By the averaging property of the h-process we have

$$\hat{P}_{(N-k-1,1)(N-k,2)} = Z^{-1}\left\{h_{k+2} + h_k + \sum_{r=2}^{k-2}\left[\binom{k}{(k+r-2)/2} + \binom{k}{(k+r-6)/2}\right] h_r + \binom{k}{(k-2)/2} h_0\right\} ;$$

and

$$\hat{P}_{(N-k-1,1)(N-k,0)} = Z^{-1}\left\{2 \sum_{r=2}^{k} \binom{k}{(k+r)/2} h_r + \binom{k}{k/2} h_0\right\} .$$

Now $g(N-k, 0) = \sum_r r\hat{\mathbf{P}}(V_m < 1 \text{ for all } m > N-k \text{ and } V_N = r) + \hat{\mathbf{P}}(V_m = 1$ for some $m\colon N-k < m < N)$. By a classical ballot theorem we are able to evaluate the given probabilities, and obtain

$$g(N-k, 0) = \left[\frac{\sum_{r=0}^{k} \phi(r)(-r)}{Z\hat{P}_{(N-k-1,1)(N-k,0)}}\right] + \left[1 - \frac{\sum_{r=0}^{k} \phi(r)}{Z\hat{P}_{(N-k-1,1)(N-k,0)}}\right] \cdot 1 ,$$

$$\text{where} \quad \phi(r) = \frac{r+1}{k+1} \binom{k+1}{(k+r+2)/2} h_r .$$

Using these relations the condition (17) can be written entirely in terms of the h-process. After combinatorial simplifications we derive (16).

An easy calculation making use of (16) for $k = 6$ now shows that if $N \geqq 8$, we continue at $(N - 7, 1)$ no matter how close to N the variance. Using the partial characterization of s stated above, it follows that $s_m \geqq 1$ for all $m \leqq N - 7$ whenever $\sigma^2 < N$, which proves Boyce's conjecture.

Observe that the analogous argument applied to $\sigma^2 > N$ yields the condition $\hat{P}_{(N-k-1,-1)(N-k,-2)} - \hat{P}_{(N-k-1,-1)(N-k,0)}[1 + g(N - k, 0)] > 0$. By taking λ_2 sufficiently large the first probability exceeds $1 - \varepsilon$, so a necessary condition for continuing at $(N - k - 1, -1)$ becomes

$$g(N - k, 0) > \frac{1 - 2\varepsilon}{\varepsilon}.$$

Taking $\varepsilon < \frac{1}{2} + N$, we see that the last condition is *never* satisfied. Hence for λ_2 sufficiently large, or equivalently for σ^2 sufficiently close to N^2, the stopping curve s becomes identically 0—no matter how large N.

That such a complex relation as (16) determines the simplest possible stopping curve demonstrates graphically the intricacy of the stock market problem for even those prediction distributions which define reasonable strategies. This suggests that the most profitable extension of our results might be the characterization of a larger family of prediction distributions for which Proposition 2 holds. Boyce's normal predictions in the continuous time model and our (λ, λ_2) family should probably be viewed as convenient, relatively tractable examples from a wider class of "nice" measures. For example, using (9) it is a simple matter to show that Proposition 2 holds for *any* family of symmetric mean-0, variance σ^2 distributions $\{\hat{p}_r(\sigma^2)\}$ such that the corresponding final h-values $h_r(\sigma^2)$ satisfy:

(a) $\{h_r(N)\}$ is uniformly distributed on $\{-N, -N + 2, \cdots, N - 2, N\}$;

(b) $\{h_{|r|}(\sigma^2)\}$ is monotone decreasing on $\{0, 2, \cdots, N\}$ for $\sigma^2 < N$, and monotone increasing for $\sigma^2 > N$.

Of special interest would be the derivation of general conditions under which the flip occurs for $\mu \neq 0$, as this phenomenon is certainly the most intriguing aspect of our problem.

REFERENCES

[1] BACHELIER, L. (1900). Théorie de la spéculation. *Ann. Sci. Ecole Norm. Sup.* **17** 21-86.

[2] BOYCE, W. M. (1970). Stopping rules for selling bonds. *Bell J. Econ. Man. Sci.* **1** 27-53.

[3] COOTNER, P. (1967). *The Random Character of Stock Market Prices.* M.I.P. Press, Cambridge.

[4] DYNKIN, E. B. (1963). Optimal selection of a stopping time for a Markov process. *Dokl. Ak. Nauk USSR* **150** 238-240.

[5] FÖLLMER, H. (1972). Optimal stopping of constrained Brownian motion. *J. Appl. Probability* **9** 557-571.

[6] GRANGER, C. W., and MORGENSTERN, O. (1970). *Predictability of Stock Market Prices*. Heath Lexington.

[7] GRIFFEATH, D. S. (1972). Computer solution of the discrete maximum entropy problem. *Technometrics* **14** 891–897.

[8] JAYNES, E. T. (1957). Information theory and statistical mechanics. *Phys. Rev.* **106** 620–630.

[9] SAMUELSON, P. A. (1965). A proof that probably anticipated prices fluctuate randomly. *Industrial Man. Rev.* **6** 41–49.

[10] SAMUELSON, P. A. (1973). Mathematics of speculative price. *SIAM Rev.* **15** 1–42.

[11] SHEPP, L. A. (1969). Explicit solutions to some problems of optimal stopping. *Ann. Math. Statist.* **40** 993–1010.

DAVID GRIFFEATH
DEPARTMENT OF MATHEMATICS
WHITE HALL
CORNELL UNIVERSITY
ITHACA, NEW YORK

J. LAURIE SNELL
DEPARTMENT OF MATHEMATICS
DARTMOUTH COLLEGE
HANOVER, NEW HAMPSHIRE