

Chapter 4. APPLICATIONS. MARKOV MODELS.

In this chapter, we look at some stopping rule problems for which the principle of optimality provides an effective method of obtaining the solution. Each of these problems has a structure that reduces the problem to a Markov decision problem with a one-dimensional state space. This allows us to show that the stopping rule given by the principle of optimality has a simple form. In fact, when the returns are functions of a Markov chain, then we may restrict attention to stopping rules that are functions of the chain.

A typical form of these problems may be described as follows. Let $\{Z_n\}_{n=1}^{\infty}$ be a sequence of random variables such that Z_n is \mathcal{F}_n -measurable, where \mathcal{F}_n denotes the σ -field generated by the observations X_1, \dots, X_n . We assume that $\{Z_n\}$ is a Markov chain in the sense that the distribution of Z_{n+1} given \mathcal{F}_n is the same as the distribution of Z_{n+1} given Z_n , and we suppose that the payoffs, Y_n , are functions of the chain, say $Y_n = u_n(Z_n)$ for some functions u_n . Then $V_n = \text{ess sup}_{N \geq n} \mathbf{E}(Y_N | \mathcal{F}_n)$ is a function of Z_n , say $V_n(Z_n)$, and the rule given by the principle of optimality has the form

$$N^* = \min\{n \geq 0 : u_n(Z_n) \geq V_n(Z_n)\}. \quad (1)$$

Thus, stopping at time n can be taken to depend only on Z_n . In the first five problems treated in this chapter, there is a time invariance that reduces this rule to a simpler form, $N^* = \min\{n \geq 0 : Z_n \geq c\}$ where c is independent of n . For the general theory of Markov stopping rule problems, see Siegmund (1967), or Chow, Robbins and Siegmund (1972), or Shiryaev (1973).

In the first section, we solve the problem of selling an asset with and without recall. This is the house-selling problem described in Example 1 of Section 1.2. It is an extension to an infinite horizon of the Cayley-Moser problem treated in Section 2.4. In the second section, we look at the problem of stopping a discounted sum. This is a version of the burglar problem of Exercise 1 of Chapter 1. In the third section, we solve the problem due to Darling, Liggett and Taylor (1976) of stopping a sum with a negative drift. In the fourth section, we treat an extension of the problem of stopping a success run mentioned in Exercise 5 of Chapter 1. The fifth section is devoted to the problem of testing simple statistical hypotheses. This is a special case of the general Bayesian statistical problem described in Example 3 of Section 1.2. In the final section, we discuss the problem of maximizing the average, presented as Example 2 of Section 1.2. It is an example of a

Markov stopping rule problem without a time invariance, and so its solution is of a higher order of difficulty.

§4.1. Selling an asset with and without recall. The house-selling problem or problem of selling an asset, as described in Chapter 1, is defined by the sequence of observations and payoffs as follows. The observations are

$$X_1, X_2, \dots \quad \text{assumed to be i.i.d. with known distribution } F(x).$$

The reward sequence depends on whether or not recall of past observations is allowed. If recall is not allowed, then

$$Y_0 = -\infty, Y_1 = X_1 - c, \dots, Y_n = X_n - nc, \dots, Y_\infty = -\infty,$$

while if recall is allowed

$$Y_0 = -\infty, Y_1 = M_1 - c, \dots, Y_n = M_n - nc, \dots, Y_\infty = -\infty,$$

where $c > 0$ is the cost per observation, and $M_n = \max\{X_1, \dots, X_n\}$. We arbitrarily put $Y_0 = -\infty$ to force you to take at least one observation. Putting $Y_\infty = -\infty$ is natural as the cost of an infinite number of observations is infinite.

The following theorem states that A1 and A2 are satisfied in problems of selling an asset with and without recall provided X_1, X_2, \dots are identically distributed with finite second moment. It is not assumed that the X_j are independent. This allows for the case where X_1, X_2, \dots are independent identically distributed given some parameter θ where θ is random with a known prior distribution. When θ is integrated out, the marginal process X_1, X_2, \dots is exchangeable (every permutation of X_1, X_2, \dots has the same distribution). The hypothesis of identical marginal distributions of the X_j , used in this theorem, is weaker than that of exchangeability. Proofs of this theorem and of a converse are contained in the appendix to this chapter.

Theorem 1. *Let X, X_1, X_2, \dots be identically distributed, let $c > 0$, and let $Y_n = X_n - nc$ or $Y_n = \max\{X_1, \dots, X_n\} - nc$.*

If $EX^+ < \infty$, then $\sup Y_n < \infty$ a.s. and $Y_n \rightarrow -\infty$ a.s.

If $E(X^+)^2 < \infty$, then $E \sup Y_n < \infty$.

We first consider the problem of selling an asset without recall: $Y_n = X_n - nc$. Suppose that X_1, X_2, \dots are independent identically distributed with finite second moment. Since A1 and A2 are satisfied, we know from Theorems 3.1 and 3.3 that an optimal stopping rule exists and is given by the principle of optimality.

Let V^* denote the expected return from an optimal stopping rule. Suppose you pay c and observe $X_1 = x_1$. Note that if you continue from this point then the x_1 is lost and the cost c has already been paid, so it is just like starting the problem over again; that is, *the problem is invariant in time*. So if you continue from this point, you can obtain

an expected return of V^* but no more. Therefore, the principle of optimality says that if $x_1 < V^*$ you should continue, and if $x_1 > V^*$ you should stop. For $x_1 = V^*$ it is immaterial what you do, but let us say you stop. This argument can be made at any stage, so the rule given by the principle of optimality is

$$N^* = \min\{n \geq 1 : X_n \geq V^*\}. \quad (2)$$

The problem now is to compute V^* . This may be done through the optimality equation,

$$V^* = E \max\{X_1, V^*\} - c = \int_{-\infty}^{V^*} V^* dF(x) + \int_{V^*}^{\infty} x dF(x) - c,$$

where F is the common distribution function of the X_i . Rearranging terms, we find

$$\int_{V^*}^{\infty} (x - V^*) dF(x) = c \quad \text{or} \quad E(X - V^*)^+ = c. \quad (3)$$

The left side is continuous in V^* and decreasing from $+\infty$ to zero. Hence, there is a unique solution for V^* for any $c > 0$.

As an example, suppose F is $\mathcal{U}(0, 1)$, the uniform distribution on the interval $(0, 1)$. For $0 \leq v \leq 1$,

$$\int_v^1 (x - v) dF(x) = (1 - v)^2/2,$$

while for $v < 0$,

$$\int_0^1 (x - v) dF(x) = 1/2 - v.$$

Equating to c , we find

$$\begin{aligned} V^* &= 1 - (2c)^{1/2} && \text{if } c \leq 1/2 \\ V^* &= -c + 1/2 && \text{if } c > 1/2. \end{aligned}$$

The optimal rule N^* calls for accepting the first offer greater than or equal to V^* .

Now consider the problem with recall, $Y_0 = Y_\infty = -\infty$, and $Y_n = M_n - nc$ where $M_n = \max\{X_1, \dots, X_n\}$. Again the rule given by the principle of optimality is optimal. Suppose at some stage you have observed $M_n = m$ and it is optimal to continue. Then at the next stage, if M_{n+1} is still m (because $X_{n+1} \leq m$), it is still optimal to continue due to the invariance of the problem in time. Thus the principle of optimality will never require you to recall an observation from an earlier stage. The best we can do among such rules is found above for the problem without recall. Thus, the same rule is optimal for both problems.

The Cayley-Moser problem has been extended to infinite horizon problems in a number of ways. Above, a cost is assessed for each observation. In another method, due to Karlin (1961), the future is discounted so that a positive return at the first stage is worth more

than the same return at later stages. These problems are treated in Exercises 2 and 3. Another possibility is to have forced stopping at a random future time. In job processing, this forced stopping time is modeled as a deadline. See Exercise 4 for an example.

Is A1 needed? An interesting feature of this problem was pointed out in a paper of Robbins (1970). By an elegant direct argument based on Wald's equation and using only the assumption that $EX^+ < \infty$, Robbins shows that the rule N^* given by (2) with V^* given by (3) is optimal within the class of stopping rules, N , such that $EY_N^- > -\infty$. This provides an extension of the optimal property of the rule N^* that is valid even if $E(X^+)^2 = \infty$. Since, as is shown in the appendix, $E(X^+)^2 < \infty$ is a necessary and sufficient condition for A1 to hold in this problem, this raises the possibility of similarly extending the general theory of Chapter 3.

However, there are difficulties of interpretation that arise because of the restriction to stopping rules that satisfy $EY_N^- > -\infty$. If $EX^+ < \infty$ and $E(X^+)^2 = \infty$, there will exist rules, N , such that $EY_N^- = -\infty$ and $EY_N^+ = +\infty$ (see the appendix). Restricting attention to rules such that $EY_N^- > -\infty$ seems to say that any rule, N , with $E|Y_N| < \infty$, no matter how bad, is better than a rule whose expected payoff does not exist because $EY_N^- = -\infty$ and $EY_N^+ = +\infty$. Do you prefer a payoff of $-\$100$ (or $+\$100$) or a gamble giving you $\$X$, where X is chosen from a standard Cauchy distribution? Rather than attempting to answer such questions, we prefer the approach assuming condition A1.

§4.2. Application to stopping a discounted sum. (Dubins and Teicher (1967)) Let X_1, X_2, \dots be independent identically distributed random variables with $EX_i^+ < \infty$, and let $S_n = \sum_1^n X_i$ ($S_0 = 0$). The random variable X_n represents your return at stage n . Returns accumulate and are paid to you in a lump sum when you stop, but the future is discounted by a factor $0 < \beta < 1$, so your return for stopping at n is

$$Y_n = \beta^n S_n \quad \text{for } n = 0, 1, \dots$$

If you never stop, your return is zero, $Y_\infty = 0$, so you will never stop at n with $S_n \leq 0$. The problem is to decide how large to let S_n get before you stop.

We show that assumptions A1 and A2 are satisfied for this problem. From the strong law of large numbers $(1/n) \sum_1^n X_j^+ \rightarrow E(X^+)$ a.s., so that

$$Y_n = n\beta^n(S_n/n) \leq n\beta^n(1/n) \sum_1^n X_j^+ \simeq n\beta^n E(X^+) \rightarrow 0 \quad \text{a.s.,}$$

and A2 is satisfied. To check A1, note that

$$\begin{aligned} \sup_n Y_n &= \sup_n \beta^n S_n \leq \sup_n \beta^n \sum_1^n X_i^+ \\ &\leq \sup_n \sum_1^n \beta^i X_i^+ \leq \sum_1^\infty \beta^i X_i^+, \end{aligned}$$

so that

$$\mathbf{E} \sup_n Y_n \leq \sum_1^{\infty} \beta^i \mathbf{E} X_i^+ = \mathbf{E} X_1^+ \beta / (1 - \beta) < \infty.$$

Thus, an optimal stopping rule exists and is given by the principle of optimality.

Suppose $S_n = s$ and it is optimal to stop. Then the present return of $\beta^n s$ is at least as large as any expected future return $\mathbf{E} \beta^{n+N} (s + S_N)$. That is to say $s(1 - \mathbf{E} \beta^N) \geq \mathbf{E} \beta^N S_N$ for all stopping rules N . The same must be true for all $s' \geq s$ so that the optimal rule N^* must be of the form for some number s_0 ,

$$N^* = \min\{n \geq 0 : S_n \geq s_0\}. \quad (4)$$

That is, stop at the first n for which $S_n \geq s_0$. To find s_0 , note that if $S_n = s_0$, then we must be indifferent between stopping and continuing. The payoff for stopping, namely s_0 , must be the same as the payoff for continuing using the rule that stops the first time the sum of the future observations is positive. That is, s_0 must satisfy the equation $s_0 = \mathbf{E} \beta^T (s_0 + S_T)$ or

$$s_0 = \mathbf{E} \beta^T S_T / (1 - \mathbf{E} \beta^T), \quad (5)$$

where $T = \min\{n \geq 0 : S_n > 0\}$ is the rule: stop at the first n , if any, for which the sum of the next n observations is positive.

When X is positive with probability 1, then $T \equiv 1$ and $S_T \equiv X_1$, so that s_0 has the simple form $s_0 = \beta \mathbf{E} X / (1 - \beta)$. This gives the optimal rule N^* of (4) in explicit form. The burglar problem of Exercise 1.1 may be put in this form where β represents the probability of not getting caught on any burglary. We return to this problem in Chapters 5 and 6.

When X may take in negative values with positive probability, the right side of (5) may difficult to evaluate in general, but in special cases the computational problems can be reduced. Dubins and Teicher call the distribution of the X_i **elementary** if X_i takes on only integer values less than or equal to one. In this special case, S_T in equation (5) becomes $+1$ for $T < \infty$ and is immaterial for $T = \infty$ since β^T is zero; hence,

$$s_0 = \mathbf{E} \beta^T / (1 - \mathbf{E} \beta^T). \quad (6)$$

To find s_0 , it suffices to find $\mathbf{E} \beta^T$. This may be computed from knowledge of the generating function of X , $G(\theta) = \mathbf{E} \theta^{-X}$ as follows. The distribution of T given $X_1 = x$ is the same as the distribution of $1 + T_1 + \dots + T_{1-x}$, where T_1, T_2, \dots are independent with the same distribution as T . Hence, letting $\phi(\beta) = \mathbf{E} \beta^T$,

$$\begin{aligned} \phi(\beta) &= \mathbf{E}(\mathbf{E}\{\beta^T | X\}) = \mathbf{E}(\mathbf{E}\{\beta^{1+T_1+\dots+T_{1-x}} | X\}) \\ &= \beta \mathbf{E} \phi(\beta)^{1-X} = \beta \phi(\beta) G(\phi(\beta)). \end{aligned} \quad (7)$$

Thus, $\phi(\beta)$ can be found by solving the equation $G(\phi(\beta)) = 1/\beta$. (See Exercise 5.)

When $EX < 0$, then $S_n \rightarrow -\infty$ a.s. and the problem is still interesting even if $\beta = 1$, provided we keep $Y_\infty = 0$. However, we need a stronger assumption on the distribution of the X_i in order that A1 still hold. This problem is treated next.

§4.3. Stopping a sum with negative drift. (Darling, Liggett and Taylor (1972))
Let X_1, X_2, \dots be i.i.d. with negative mean $\mu = EX < 0$, and let $S_n = X_1 + \dots + X_n$, ($S_0 = 0$). The payoff for stopping at n is taken to be

$$Y_n = (S_n)^+ \quad \text{for } n = 0, 1, \dots, \text{ and } \quad Y_\infty = 0.$$

The law of large numbers implies that $S_n \rightarrow -\infty$ a.s. Hence, A2 is satisfied. We expect that an optimal rule may sometimes continue forever. In fact, it is clear that one need not stop at any n for which $S_n \leq 0$.

This model has application to the problem of exercising options in stock market transactions (the American version without a deadline). The owner of an option has the privilege of purchasing a fixed quantity of a stock at a fixed price, here normalized to zero, and then reselling the stock at the market price which fluctuates as a classical random walk. Continuing forever may be interpreted as not exercising the option. If a cost of waiting is taken into account, the walk can easily have negative drift. Darling, Liggett and Taylor also treat the more realistic problem in which the logarithm of the prices form a random walk and there is a discount $0 < \beta \leq 1$, so that the return function is $Y_n = \beta^n(e^{S_n} - 1)^+$, assuming $Ee^X < \beta^{-1}$.

We assume that $E(X^+)^2 < \infty$. Then the following theorem of Kiefer and Wolfowitz (1956, Theorem 5), shows that A1 is satisfied and there exists an optimal rule.

Theorem 2. *Let X, X_1, X_2, \dots be i.i.d. with finite mean $\mu < 0$ and let $M = \sup_{n \geq 0} S_n$. Then,*

$$EM < \infty \quad \text{if, and only if,} \quad E(X^+)^2 < \infty.$$

A proof is given in the appendix.

We show below that the optimal stopping rule is

$$N^* = \min\{n \geq 0 : S_n \geq EM\}, \tag{8}$$

and its expected return is $E(M - EM)^+$.

Note that V_n^* depends on past observations only through S_n . In fact, $V_n^* = V_0^*(S_n)$, where $V_0^*(s) = \sup_{N \geq 0} E(s + S_N)^+$ is the optimal return starting with initial fortune $S_0 = s$. Suppose $S_n = s > 0$ and it is optimal to stop. Then by the principle of optimality, the present return s is at least as large as future expected return $E(s + S_N)^+$. That is to say, $E\{\max(S_N, -s)\} \leq 0$ for all stopping rules N . Since this expectation is nonincreasing in s and independent of n , it must also be optimal to stop when $S_m = s'$ for any m and any $s' \geq s$. Thus, N^* must be of the form $N^* = \min\{n \geq 0 : S_n \geq s_0\}$ for some s_0 .

Let $N = N(s) = \min\{n \geq 0 : S_n \geq s\}$. We find the optimal value of s by computing the expected return EY_N . First note that the distribution of $M - S_N$ given S_N , where $S_N \geq s$, is the same as the distribution of M , because $M - S_N = \sup_n S'_n$, where $S'_n = X_{N+1} + \dots + X_{N+n}$. Hence,

$$\begin{aligned} EY_N &= E\{S_N I(S_N \geq s)\} = E\{M I(S_N \geq s)\} - E\{(M - S_N) I(S_N \geq s)\} \\ &= E\{M I(S_N \geq s)\} - EM \cdot P\{S_N \geq s\} = E\{(M - EM) I(S_N \geq s)\} \\ &= E\{(M - EM) I(M \geq s)\}. \end{aligned} \quad (9)$$

This is clearly nondecreasing in s if $s \leq EM$ and nonincreasing in s if $s \geq EM$. The optimal value of s is therefore $s = EM$ and the optimal expected return is $E(M - EM)^+$.

§4.4. Rebounding from failures. (Ferguson (1976)) Let Z_1, Z_2, \dots be i.i.d. random variables, and let $\epsilon_1, \epsilon_2, \dots$ be i.i.d. Bernoulli random variables, independent of Z_1, Z_2, \dots , with probability p of success, $P(\epsilon = 1) = p = 1 - P(\epsilon = 0)$, $0 < p < 1$. In a given economic system, Z_i represents your return and ϵ_i represents the indicator function of success in time period i . As long as successes occur consecutively, your returns accumulate, but when there is a failure, your accumulated return drops to zero. Failure represents a failure of your enterprise due to mismanagement, or a general calamity on the stock market, or a revolution in the country, etc. A failure does not remove you from the system as it does the burglar. You are allowed to accumulate future returns until the next failure drops you to zero again, and so on. If $X_0 = x$ denotes your initial fortune, then your fortune at the end of the n th time period is X_n , where

$$X_n = \epsilon_n(X_{n-1} + Z_n), \quad \text{for } n = 1, 2, \dots \quad (10)$$

You want to retire to a safer, more comfortable environment while you are still young enough to enjoy it. Thus, if $c > 0$ represents the cost of the passing of time, your return if you stop at stage n is

$$Y_n = X_n - nc \quad (11)$$

for finite n . We take $Y_\infty = -\infty$ and note that A2 is satisfied. The problem is to choose a stopping rule N to maximize your expected net return, EY_N .

In the special case in which the Z_i are identically one, this is the problem of stopping a success run described in Exercise 5.

The following theorem states that A1 is satisfied if the distribution of the Z 's has finite second moment. A proof is given in the appendix.

Theorem 3. $E \sup_n (X_n - nc) < \infty$ if, and only if, $E(Z^+)^2 < \infty$.

We note that the invariance of the problem in time implies that V_n^* depends on X_1, \dots, X_n only through the value of X_n , and that if $V_0^*(x)$ represents the optimal return with initial fortune x , then $V_n^*(X_n) = V_0^*(X_n) - nc$.

Suppose it is optimal to stop with $X_n = x$; then $x - nc \geq V_0^*(x) - nc$, or equivalently $x \geq \text{E}Y_N(x)$ for all stopping rules N , where $Y_n(x)$ represents the payoff as a function of the initial fortune, x . For any $x' > x$, the difference $Y_n(x') - Y_n(x)$ is 0 or $x' - x$ according to whether the first failure has or has not occurred by time n . Thus, $\text{E}Y_N(x') - \text{E}Y_N(x) = (x' - x)\text{P}(N < K) \leq x' - x$, where K is the time of the first failure. This implies that for all stopping rules N ,

$$x' \geq \text{E}Y_N(x') - \text{E}Y_N(x) + x \geq \text{E}Y_N(x'),$$

so that it is optimal to stop with $X_m = x'$ for any m and any $x' \geq x$. Hence, the optimal rule N^* given by the principle of optimality has the form

$$N(s) = \min\{n \geq 0 : X_n \geq s\} \quad (12)$$

for some s .

As an illustration, we take the distribution of the Z 's to be exponential with density $f(z) = (1/\mu)\exp\{-z/\mu\}\mathbf{I}(z > 0)$. To find the optimal value of s , let us compute the expected payoff for $N = N(s)$, for $s > 0$ and for $x = 0$, namely,

$$\text{E}Y_N = \text{E}(X_N - Nc) = \text{E}X_N - c\text{E}N.$$

The lack of memory property of the exponential distribution implies that $X_N - s$ has the same distribution as Z . Hence, $\text{E}X_N = s + \mu$. To find $\text{E}N$, let

$$N' = N'(s) = \min\{n \geq 0 : S_n \geq s\}$$

where $S_n = Z_1 + \dots + Z_n$. The sequence of points, S_1, S_2, \dots , forms a Poisson point process with points occurring at rate $1/\mu$, so that $N' - 1$, the number of points in $(0, s)$, has a Poisson distribution with mean $\lambda = s/\mu$. Let K represent the time of the first failure, so that K has the geometric distribution, $\text{P}(K = k) = (1 - p)p^{k-1}$ for $k = 1, 2, \dots$. Then,

$$\begin{aligned} \text{E}N &= \text{E}N\mathbf{I}(N' < K) + \text{E}N\mathbf{I}(N' \geq K) \\ &= \text{E}N'\mathbf{I}(N' < K) + \text{E}(K + \text{E}N)\mathbf{I}(N' \geq K) \\ &= \text{E}\min\{N', K\} + (\text{E}N)\text{P}(N' \geq K) \end{aligned}$$

so that $\text{E}N = \text{E}\min\{N', K\}/\text{P}(N' < K)$. This may be computed from

$$\text{P}(N' < K) = \text{E}\{\text{P}(N' < K|N')\} = \text{E}p^{N'} = p \exp\{-\lambda(1 - p)\},$$

and

$$\begin{aligned} \text{E}\min\{n, K - 1\} &= \sum_{k=1}^n k(1 - p)p^k + n \sum_{k=n+1}^{\infty} (1 - p)p^k \\ &= \sum_{k=1}^n kp^k - \sum_{k=1}^n kp^{k+1} + np^{n+1} \\ &= \sum_{k=1}^n p^k \\ &= p(1 - p^n)/(1 - p), \end{aligned}$$

and

$$\begin{aligned} E \min\{N', K\} &= E \min\{N' - 1, K - 1\} + 1 \\ &= (p/(1-p))E(1 - p^{N'-1}) + 1 \\ &= (p/(1-p))(1 - e^{-\lambda(1-p)}) + 1 \\ &= (1/(1-p))[1 - pe^{-\lambda(1-p)}]. \end{aligned}$$

We search for the value of s that maximizes

$$EY_N = s + \mu - (c/(p(1-p)))[e^{s(1-p)/\mu} - p].$$

Taking a derivative with respect to s , setting the result to zero and solving, gives the optimal value of s as

$$s = (\mu/(1-p)) \log(p\mu/c).$$

provided $p\mu \geq c$. If $p\mu < c$, then EY_N is a decreasing function of $s \geq 0$, and so the optimal value of s is $s = 0$, that is, stop without taking any observations.

§4.5 Testing Simple Statistical Hypotheses. The Bayes approach to general sequential statistical problems was discussed in Example 3 of Chapter 1. Here we specialize to the problem of testing simple statistical hypotheses. In this problem, there are two hypotheses, H_0 and H_1 , and one distribution corresponding to each hypothesis. In the notation of Example 1.3, the parameter space is a two-point set, $\Theta = \{H_0, H_1\}$, and the observations, X_1, X_2, \dots , are assumed to be i.i.d. according to a density $f_0(x)$ if H_0 is true, and density $f_1(x)$ if H_1 is true, where $f_0(x)$ and $f_1(x)$ are distinct as distributions. We must decide which hypothesis to accept. Thus, the action space is also a two-point set, $\mathcal{A} = \{a_0, a_1\}$, where a_0 (resp. a_1) represents the action “accept H_0 ” (resp. “accept H_1 ”). We lose nothing if we accept the true hypothesis, but if we accept the wrong hypothesis we lose an amount depending on which hypothesis is true; thus,

$$L(H_i, a_j) = \begin{cases} 0 & \text{if } i = j, \\ L_i & \text{if } i \neq j, \end{cases}$$

where L_0 and L_1 are given positive numbers.

We are given the prior probability, τ_0 , that H_1 is the true hypothesis. After observing X_1, \dots, X_n , the posterior probability that H_1 is the true hypothesis becomes, according to Bayes rule,

$$\tau_n(X_1, \dots, X_n) = \frac{\tau_0 \prod_1^n f_1(X_i)}{\tau_0 \prod_1^n f_1(X_i) + (1 - \tau_0) \prod_1^n f_0(X_i)}. \quad (13)$$

The likelihood ratio is $\lambda(x) = f_1(x)/f_0(x)$ with the understanding that $\lambda(x) = 0$ if $f_1(x) = 0$ and $f_0(x) > 0$, $\lambda(x) = \infty$ if $f_1 > 0$ and $f_0(x) = 0$, and $\lambda(x)$ is undefined if $f_1(x) = 0$ and $f_0(x) = 0$. Using this, we may rewrite τ_n as

$$\tau_n(X_1, \dots, X_n) = \frac{\tau_0 \prod_1^n \lambda(X_i)}{\tau_0 \prod_1^n \lambda(X_i) + (1 - \tau_0)} = \frac{\tau_0 \lambda_n}{\tau_0 \lambda_n + (1 - \tau_0)} \quad (14)$$

where $\lambda_n = \lambda_n(X_1, \dots, X_n)$ denotes the likelihood ratio, or probability ratio of the first n observations, $\lambda_n = \prod_{i=1}^n f_1(X_i)/f_0(X_i)$ and $\lambda_0 \equiv 1$.

Suppose it is decided to stop and the probability of H_1 is τ . If H_1 is accepted, the expected loss is $(1 - \tau)L_0$, while if H_0 is accepted, the expected loss is τL_1 . Thus, it is optimal to accept H_1 if $(1 - \tau)L_0 < \tau L_1$ and to accept H_0 otherwise, incurring an expected loss of

$$\rho(\tau) = \min\{\tau L_1, (1 - \tau)L_0\} \quad (15)$$

Therefore, if we stop at stage n having observed X_1, \dots, X_n , we would accept H_1 if $\tau_n(X_1, \dots, X_n)L_1 < (1 - \tau_n(X_1, \dots, X_n))L_0$, and the expected terminal loss would then be $\rho(\tau_n(X_1, \dots, X_n))$.

There is a cost of $c > 0$ for each observation taken, so the total loss plus cost of stopping at stage n after observing X_1, \dots, X_n is

$$Y_n = \rho(\tau_n(X_1, \dots, X_n)) + nc \quad \text{for } n = 0, 1, 2, \dots, \quad (16)$$

and $Y_\infty = +\infty$ if we never stop. The problem is to find a stopping rule N to minimize EY_N .

Now that the stopping rule problem has been defined, let us check conditions A1 and A2. Since this is a minimization problem, A1 and A2 must be replaced by

A1. $E\{\inf_n Y_n\} > -\infty$.

A2. $\liminf_{n \rightarrow \infty} Y_n \geq Y_\infty$ a.s.

We note that A1 follows from $Y_n \geq 0$, and A2 follows since $Y_n \geq nc \rightarrow \infty = Y_\infty$. Thus there is an optimal rule, N^* , given by the principle of optimality. Let $V_0^*(\tau_0)$ denote the expected loss plus cost using this rule, as a function of the prior probability τ_0 . There is an invariance in time; at stage n after observing X_1, \dots, X_n , the probability distribution of future payoffs and costs is the same as it was at stage 0 except that the prior has been changed to $\tau_n(X_1, \dots, X_n)$ and the cost of the first n observations must be paid. Thus,

$$V_n^*(X_1, \dots, X_n) = V_0^*(\tau_n(X_1, \dots, X_n)) + nc, \quad (17)$$

where $V_n^*(X_1, \dots, X_n)$ is the conditional minimum expected loss plus cost having observed X_1, \dots, X_n . The rule given by the principle of optimality reduces to

$$\begin{aligned} N^* &= \min\{n \geq 0 : Y_n = V_n^*(\tau_n(X_1, \dots, X_n))\} \\ &= \min\{n \geq 0 : \rho(\tau_n(X_1, \dots, X_n)) = V_0^*(\tau_n(X_1, \dots, X_n))\} \end{aligned} \quad (18)$$

From this it follows that the optimal decision to stop may be based on the value of $\tau_n(X_1, \dots, X_n)$.

We now note that $V_0^*(\tau)$ is a concave function of $\tau \in [0, 1]$. Let $\alpha, \tau, \tau' \in [0, 1]$; we are to show

$$\alpha V_0^*(\tau) + (1 - \alpha)V_0^*(\tau') \leq V_0^*(\alpha\tau + (1 - \alpha)\tau'). \quad (19)$$

Suppose before stage 0 in the above decision problem, a coin with probability α of heads is tossed and if the coin comes up heads the prior $\tau_0 = \tau$ is used, while if the coin comes up tails the prior $\tau_0 = \tau'$ is used. The left side of (19) is the minimum value of this stopping rule problem when the information on the outcome of the toss may be used. The right side is the minimum value when this information may not be used. Since the class of stopping rules that ignore this information is a subset of the class that may use it, the inequality follows.

In addition, we note that $V_0^*(0) = 0 = \rho(0)$ and $V_0^*(1) = 0 = \rho(1)$. This together with concavity and $V_0^*(\tau) \leq \rho(\tau)$ implies that there are numbers a and b with $0 \leq a \leq L^* \leq b \leq 1$ such that the optimal rule (18) has the form,

$$N^* = \min\{n \geq 0 : \tau_n(X_1, \dots, X_n) \leq a \quad \text{or} \quad \tau_n(X_1, \dots, X_n) \geq b\}.$$

where $L^* = L_0/(L_0 + L_1)$. Writing the inequalities in this expression in terms of the likelihood ratio, $\lambda_n = \prod_1^n f_1(X_i)/f_0(X_i)$, we find

$$N^* = \min\{n \geq 0 : \lambda_n \leq \frac{(1 - \tau_0)a}{\tau_0(1 - a)} \quad \text{or} \quad \lambda_n \geq \frac{(1 - \tau_0)b}{\tau_0(1 - b)}\}.$$

This is the Wald sequential probability ratio test: Sample as long as the probability ratio, λ_n lies between two preassigned numbers, $\alpha < \lambda_n < \beta$; if λ_n falls below α , stop and accept H_0 , if λ_n rises above β , stop and accept H_1 .

The problem of finding the values of a and b of the optimal rule N^* usually requires approximation. There are standard methods of approximation that originate with Wald. See the book of Siegmund (1985) for general methods and applications to statistical problems. Numerical approximation on a computer may also be used. This involves evaluating EY_N for an arbitrary rule of the form $N_{\alpha, \beta} = \min\{n \geq 0 : \lambda_n \leq \alpha \text{ or } \lambda_n \geq \beta\}$, and then searching for α and β to minimize this quantity. In a few cases, this may be done explicitly.

§4.6. Application to maximizing the average. Let X_1, X_2, \dots be i.i.d. random variables with mean $\mu = EX$. For $n \geq 1$, let $Y_n = S_n/n$, and let $Y_0 = -\infty$ and $Y_\infty = \mu$. By the law of large numbers, $Y_n \rightarrow \mu$ a.s. so A2 is satisfied. Therefore an optimal stopping rule will exist if A1 is satisfied; that is, if $E \sup_n (S_n/n) < \infty$. The following theorem states that A1 holds if $EX \log^+(X) < \infty$, where $\log^+(x)$ is defined to be $\log(x)$ if $x > 1$, and 0 otherwise.

Theorem 4. *Let X_1, X_2, \dots be i.i.d. with finite mean μ . Then*

$$E \sup_n (S_n/n) < \infty \quad \text{if, and only if,} \quad EX \log^+(X) < \infty.$$

The sufficiency of this condition is due essentially to Marcinkiewitz and Zygmund (1937). The necessity is due to Burkholder (1962). See the appendix for a proof.

The lack of a time invariance in this problem makes it more difficult than the other problems in this chapter. The problem still has a Markov structure so the decision to stop at stage n can be made on the basis of S_n alone. The principle of optimality may be used to show that the optimal rule is given by a monotone sequence of constants $a_1 \geq a_2 \geq \dots$ with stopping at stage n if $S_n/n \geq a_n$. However, even in specific simple cases, such as the Bernoulli case, the values of the optimal a_n are quite difficult to evaluate.

An interesting question arises: Does the optimal rule for maximizing the average stop with probability one? Chow and Robbins (1965) show that it does in the Bernoulli case, and Dvoretzky (1967) and Teicher and Wolfowitz (1966) show that it does provided $EX^2 < \infty$. The general problem is treated in Klass (1973) where it is seen that the optimal rule stops with probability one provided $E(X^+)^{\alpha} < \infty$ for some $\alpha > 1$, and examples are given for which the optimal rule does not stop with probability one.

§4.7 Exercises.

1. Let X_1, X_2, \dots be a sample from the negative exponential distribution with density

$$f(x|\sigma) = \sigma \exp\{-\sigma x\}I\{x > 0\},$$

where $\sigma > 0$ is known. Find the optimal stopping rule for the problem $Y_n = X_n - nc$ (or $Y_n = \max\{X_1, \dots, X_n\} - nc$).

2. *Selling an asset without recall and without cost, but with discounted future.* Let X_1, X_2, \dots be independent identically distributed with finite first moment, let $0 < \beta < 1$, and let $Y_n = \beta^n X_n$ with $Y_0 = Y_{\infty} = 0$.

- (a) Show that A1 and A2 are satisfied. (Hint: $\sup_n Y_n \leq \sum_n \beta^n |X_n|$.)
- (b) Show that it is optimal to stop after the first n for which $X_n \geq V^*$, where V^* is the optimal expected return and the unique solution of the equation $V^* = \beta E \max\{X, V^*\}$.
- (c) Specialize to the cases $\mathcal{U}(0, 1)$ and $\mathcal{U}(-1, 1)$.

3. Solve the discounted version of the house-selling problem for sampling with recall, where $Y_n = \beta^n \max\{X_1, \dots, X_n\}$. (Hint for (a): $Y_n \leq \max\{\beta|X_1|, \dots, \beta^n|X_n|\}$.)

4. *Job Processing.* A given job must be assigned to a processor before a (continuous time) deadline, D , which is unknown (until it occurs) but has an exponential distribution at rate $\alpha > 0$: $P(D > d) = \exp\{-\alpha d\}$ for $d > 0$. Job processors arrive at times, T_1, T_2, \dots , given by a renewal process independent of D ; that is, the interarrival times, $T_1, T_2 - T_1, T_3 - T_2, \dots$ are i.i.d. according to a known distribution G on the positive real line and independent of D . As each processor arrives, its value, that is to say the return it provides if the job is assigned to it, becomes known. These values, denoted by Z_1, Z_2, \dots , are assumed to be i.i.d. according to a known distribution, F , and independent of the arrival times, T_1, T_2, \dots , and deadline, D . If the deadline occurs before the job is assigned to a processor, the return is some constant, say μ . Therefore, the payoff, Y_n , for assigning the job to the n th processor to arrive is taken to be $Y_0 = \mu$, $Y_{\infty} = \mu$, and

$$Y_n = Z_n I(T_n < D) + \mu I(T_n \geq D).$$

Let $K = \min\{n \geq 1 : T_n > D\}$ denote the index of the first processor to arrive after the deadline, and let $\beta = P(D > T_1) = E(\exp\{-\alpha T_1\})$. Then K is independent of Z_1, Z_2, \dots and has a geometric distribution, $P(K \geq k) = \beta^{k-1}$ for $k = 1, 2, \dots$. The payoff may be written in an equivalent form,

$$Y_n = Z_n I(n < K) + \mu I(n \geq K).$$

In this form, we see that the problem depends on α and the distribution G only through the constant β .

(a) Assume that $EZ^+ < \infty$. Show that A1 and A2 are satisfied.

(b) Show that it is optimal to stop at the first n for which $Z_n \geq V^*$, where V^* is the optimal expected return and the unique solution of the equation

$$V^* = \mu + (\beta/(1 - \beta))E((Z - V^*)^+).$$

(c) Take $\mu = 0$ and specialize to the cases $\mathcal{U}(0, 1)$ and $\mathcal{U}(-1, 1)$.

5. In the problem of Section 4.2, suppose the distribution of X gives probability p to $+1$ and probability $1 - p$ to -1 .

(a) Find the optimal stopping rule, N^* .

(b) Find $P(N^* < \infty)$.

(c) For what values of $p < 1/2$ is it true that N^* requires you to stop at the first time that $S_n = 1$ no matter what β is?

6. In the problem of stopping a sum with negative drift, assume that the distribution of the X_i is elementary.

(a) Show that the distribution of M is geometric with probability of success $P(T < \infty)$ where $T = \min\{n \geq 0 : S_n > 0\}$.

(b) Show how to find $P(T < \infty)$ and hence EM from knowledge of the generating function, $G(\theta) = E\theta^{-X}$.

(c) Specialize to the case $P(X = 1) = p$ and $P(X = -1) = 1 - p$, where $p < 1/2$.

7. *Setting a record* (Ferguson and MacQueen (1992).) Let X_1, X_2, \dots be i.i.d. with nonpositive mean, let $S_n = X_1 + \dots + X_n$, and let $M_n = \max\{M_0, S_1, \dots, S_n\}$ with $M_0 = 0$. Let $c > 0$ and let $Y_n = M_n - nc$, with $Y_\infty = -\infty$. The problem of choosing a stopping rule to maximize EY_N is the problem of deciding when to give up trying to set a new record, the return being the value of the record.

(a) Show that A1 and A2 are satisfied if $E(X^+)^2 < \infty$.

(b) Show that the optimal rule has the form, $N = \min\{n \geq 0 : M_n - S_n \geq \gamma\}$ for some $\gamma \geq 0$.

(c) Suppose X_1, X_2, \dots are i.i.d. with $P(X = 1) = P(X = -1) = 1/2$. Show the stopping rule of (b) for γ an integer has expected return, $EY_N = \gamma - \gamma(\gamma + 1)c$. Find the optimal rule.

8. *Attaining a goal* (Ferguson and MacQueen) Let X_1, X_2, \dots be i.i.d. and let $S_n = X_0 + X_1 + \dots + X_n$, where X_0 is a given number. Let $a > 0$, $c > 0$, let $Y_n = I(S_n \geq a) - nc$ and let $Y_\infty = -\infty$. This is the problem of choosing a stopping rule to maximize the

probability of attaining a goal when there is a cost of time.

(a) Show that A1 and A2 are satisfied.

(b) Show that the optimal rule has the form, $N = \min\{n \geq 0 : S_n < \gamma \text{ or } S_n \geq a\}$ for some γ .

(c) Suppose that X_1, X_2, \dots are i.i.d. with $P(X = 1) = P(X = -1) = 1/2$. Show the stopping rule of (b) for γ an integer has expected return, for $\gamma < X_0 < a$, of $EY_N = (X_0 - \gamma)/(a - \gamma) - c(X_0 - \gamma)(a - X_0)$. Find the optimal rule.

9. (Ferguson and MacQueen) Let X_1, X_2, \dots be i.i.d. with a distribution that is symmetric about zero. Let $S_n = X_1 + \dots + X_n$, let $c > 0$ and let $Y_n = |S_n| - nc$ for n finite and $Y_\infty = -\infty$.

(a) Show that A1 and A2 are satisfied if $EX^2 < \infty$.

(b) Show that an optimal rule has the form $N = \min\{n \geq 0 : |S_n| \geq \gamma\}$ for some $\gamma \geq 0$.

(c) Suppose that X_1, X_2, \dots are i.i.d. with $P(X = 1) = P(X = -1) = 1/2$. Show that the return of the stopping rule of (b) is $EY_N = \gamma - \gamma^2 c$. Find an optimal rule.

10. *A change-point repair model.* (This model is due essentially to Girshick and Rubin (1952).) Let T denote an unobservable change-point, and assume the distribution of T is geometric with known parameter π , $P(T = t) = (1 - \pi)\pi^{t-1}$ for $t = 1, 2, \dots$. Given $T = t$, the observations, X_1, X_2, \dots , are independent with X_1, \dots, X_{t-1} i.i.d. having density $f_0(x)$ with mean $\mu_0 \geq 0$, and X_t, X_{t+1}, \dots i.i.d. having density $f_1(x)$ with mean $\mu_1 < 0$ and finite variance. The observations represent the daily returns from operating a machine. Let the return for stopping at n be $Y_n = S_n - aI(T \leq n) - c$, where $S_n = \sum_1^n X_j$, $c > 0$ represents the cost of shutdown for repair, and $a > 0$ is the excess cost of repair when the machine is in the poor state. Since T is unobservable, it is preferable to work with $Y_n = S_n - aQ_n - c$, where $Q_n = P\{T \leq n | X_1, \dots, X_n\}$ for finite n ($Q_0 = 0$) and $Y_\infty = -\infty$.

(a) Show that A1 and A2 are satisfied.

(b) Find Q_{n+1} as a function of Q_n and $\lambda(X_{n+1})$, where $\lambda(x)$ is the likelihood ratio, $\lambda(x) = f_1(x)/f_0(x)$.

(c) Show there is an optimal rule of the form, $N = \min\{n \geq 0 : Q_n \geq \gamma\}$ for some constant $\gamma \geq 0$.

11. *Selling two assets.* (Bruss and Ferguson, (1997).) You want to buy Christmas presents for your two children. After deciding which two presents to buy, you go to various stores. With two presents to buy, you can be a little more choosy. If the price of one of the gifts is clearly too high, you know you will have to go to another store anyway, so you will reject a borderline price for the other gift.

We restate this problem as a selling problem to be able to use the formulas of §4.1. You have two objects to sell, x and y . Offers come in daily for these objects, $(X_1, Y_1), (X_2, Y_2), \dots$, assumed i.i.d. with finite second moments. There is a cost of $c > 0$ to observe each vector. At each stage you may sell none, one or both of the objects. If just one of the objects is sold, you must continue until the other object is sold. Your payoff is the sum of the selling prices minus c times the number of vectors observed.

Once one object is sold, the problem reduces to the standard problem of §4.1. Let V_x and V_y denote the optimal values of selling the x -object and y -object separately. These

values are the unique solutions of the Optimality Equations

$$E(X - V_x)^+ = c \quad \text{and} \quad E(Y - V_y)^+ = c.$$

Therefore, one can consider the two-asset selling problem as a stopping rule problem in which stopping at stage n with offers (X_n, Y_n) yields payoff $W_n = \max\{X_n + Y_n, X_n + V_y, Y_n + V_x\}$.

(a) Find the Optimality Equation for V_{xy} , the value of the two-asset problem, and describe the optimal stopping rule.

(b) Suppose the (X_n, Y_n) are a sample from (X, Y) with X and Y independent having uniform distributions on the interval $(0, 1)$. Take $c = 1/8$ (so that $V_x = V_y = 1/2$), and find the optimal rule.

(c) In (b), suppose $c = 1/2$ (so that $V_x = V_y = 0$), and find the optimal rule.

APPENDIX TO CHAPTER 4

The following proof of Theorem 1 is taken partly from DeGroot (1970) pp. 350-352, where it is attributed to Bramblett (1965). For this theorem and the others in this appendix, the following inequalities are basic.

$$\sum_{n=1}^{\infty} P(Z > n) \leq E(Z^+) = \int_0^{\infty} P(Z > z) dz \leq \sum_{n=0}^{\infty} P(Z > n)$$

so that

$$E(Z^+) < \infty \quad \text{if and only if} \quad \sum_n P(Z > n) < \infty.$$

Similarly,

$$E(Z^+)^2 = 2 \int_0^{\infty} z P(Z > z) dz < \infty \quad \text{if and only if} \quad \sum_n n P(Z > n) < \infty.$$

and

$$E(Z^+)^2 = 2 \int_0^{\infty} \int_0^{\infty} P(Z > z+u) du dz < \infty \quad \text{if and only if} \quad \sum_k \sum_n P(Z > n+k) < \infty.$$

Theorem 1. *Let X, X_1, X_2, \dots be identically distributed, let $c > 0$, and let $Y_n = X_n - nc$ or $Y_n = M_n - nc$, where $M_n = \max\{X_1, \dots, X_n\}$. If $E(X^+) < \infty$, then $\sup Y_n < \infty$ a.s. and $Y_n \rightarrow -\infty$ a.s. If $E(X^+)^2 < \infty$, then $E \sup_{n \geq 1} Y_n \leq E(X^+)^2 / (2c)$.*

Proof. Since $M_n - nc = \max(X_1 - nc, \dots, X_n - nc) \leq \max(X_1 - c, \dots, X_n - nc)$, one sees that $\sup(X_n - nc) = \sup(M_n - nc)$. This implies that in the statements about $\sup Y_n$ it does not matter which definition of Y_n we take; so let us take $Y_n = X_n - nc$. Suppose $E(X^+) < \infty$. Then,

$$\begin{aligned} P(\sup_{n \geq 1} Y_n > z) &\leq \sum_{n=1}^{\infty} P(Y_n > z) = \sum_{n=1}^{\infty} P(X > z + nc) \\ &= \sum_{n=1}^{\infty} P((X - z)/c > n) \leq E((X - z)^+/c) \rightarrow 0 \end{aligned}$$

as $z \rightarrow \infty$. Thus, $\sup Y_n < \infty$ a.s. Moreover, $Y_n \leq M_n - nc = (M_n - nc/2) - nc/2 \leq U - nc/2 \rightarrow -\infty$, where $U = \sup(M_n - nc/2)$.

Now assume that $E(X^+)^2 < \infty$.

$$\begin{aligned}
E \sup_{n \geq 1} Y_n &\leq \int_0^\infty P(\sup_{n \geq 1} Y_n > z) dz \\
&\leq \int_0^\infty E((X - z)^+ / c) dz \\
&= \int_0^\infty \int_z^\infty (x - z) dF(x) dz / c \\
&= \int_0^\infty \int_0^x (x - z) dz dF(x) / c \\
&= \int_0^\infty (x^2 / 2) dF(x) / c = E(X^+)^2 / (2c). \blacksquare
\end{aligned}$$

For the converse, the variables are required to be independent.

Theorem 1'. *If X, X_1, X_2, \dots are i.i.d. and if $E \sup(X_n - nc) < \infty$, then $E(X^+)^2 < \infty$.*

Proof. Take $c = 1$ without loss of generality and suppose $E \sup_{n \geq 0} (X_n - n) < \infty$. Then,

$$\begin{aligned}
P(\sup_n (X_n - n) > z) &= 1 - \prod_{n=1}^\infty P(X_n - n \leq z) \\
&= 1 - \prod_{n=1}^\infty (1 - P(X > z + n)) \geq 1 - \exp\{-\sum_{n=1}^\infty P(X > z + n)\}.
\end{aligned}$$

Since $\sum_z P(\sup_n (X_n - n) > z) < \infty$, we have $P(\sup_n (X_n - n) > z) \rightarrow 0$ as $z \rightarrow \infty$, which in turn implies that $\sum_{n=1}^\infty P(X > z + n) \rightarrow 0$, so that for sufficiently large z ,

$$P(\sup_n (X_n - n) > z) \geq \sum_{n=1}^\infty P(X > z + n) / 2$$

(using $1 - \exp\{-x\} \geq x/2$ for x sufficiently small.) Hence,

$$\begin{aligned}
\sum_z P(\sup_n (X_n - n) > z) < \infty &\text{ implies} \\
\sum_z \sum_n P(X > z + n) < \infty &\text{ which implies} \\
E(X^+)^2 < \infty. &\blacksquare
\end{aligned}$$

As an alternate proof of Theorem 1', we show that the rule $N = \min\{n \geq 1 : X_n \geq 2cn\}$ gives $E(X_N - Nc)^+ = \infty$ when $EX^+ < \infty$ and $E(X^+)^2 = \infty$.

$$\begin{aligned} E(X_N - Nc)^+ &= \sum_{n=1}^{\infty} E(X_n - nc)I(N = n) \\ &\geq \sum_{n=1}^{\infty} cnP(N = n) \\ &= \sum_{n=1}^{\infty} cnP(N > n - 1)P(X_n > 2cn) \\ &\geq \sum_{n=1}^{\infty} cnP(N = \infty)P(X_n > 2cn) = \infty. \end{aligned}$$

since $E(X^+)^2 = \infty$ implies $\sum_{n=1}^{\infty} nP(X_n > 2cn) = \infty$, and

$$\begin{aligned} P(N = \infty) &= P(X_n < 2cn \text{ for all } n) \\ &= \prod_{n=1}^{\infty} P(X_n < 2nc) = \prod_{n=1}^{\infty} (1 - P(X_n \geq 2nc)) \\ &\sim \exp\left\{-\sum_{n=1}^{\infty} P(X_n > 2cn)\right\} \sim \exp\{-EX^+/2c\} > 0. \end{aligned}$$

The following proof of the result of Kiefer and Wolfowitz (1956) is a modification, due to Thomas Liggett, of a computation of Kingman (1962). We assume a finite variance for X and derive an upper bound for EM . The theorem of Section 4.3, assuming only that $E(X^+)^2 < \infty$, may be deduced from Theorem 2 and 2' below by truncating the distribution of X below at $-B$ where B is chosen large enough so that if $X' = \max\{X, -B\}$, then EX' is still negative. Let $X'_j = \max\{X_j, -B\}$, $S'_n = \sum_1^n X'_j$, and $M' = \sup_{n \geq 0} S'_n$. Then $M \leq M'$, so that from Theorem 2 below, $EM \leq EM' \leq \text{Var}(X')/(2|EX'|) < \infty$.

Theorem 2. *Let X, X_1, X_2, \dots be i.i.d. with $\mu = EX < 0$ and $\sigma^2 = \text{Var} X < \infty$. Let $S_n = \sum_1^n X_j$, $S_0 = 0$ and $M = \sup_{n \geq 0} S_n$. Then*

$$EM \leq \sigma^2/(2|\mu|).$$

Proof. Let $M_n = \max_{0 \leq j \leq n} S_j$. Note that $M_{n+1} = \max(0, X_1 + \max_{1 \leq j \leq n+1} (S_j - X_1))$ so that the distribution of M_{n+1} is the same as the distribution of $(M_n + X)^+$. Then, writing $M_n + X = (M_n + X)^+ - (M_n + X)^-$, and noting that $(M_n + X)^2 = ((M_n + X)^+)^2 + ((M_n + X)^-)^2$, we find,

$$\begin{aligned} E(M_n + X)^- &= E(M_{n+1}) - E(M_n) - E(X) \quad \text{and} \\ E((M_n + X)^-)^2 &= -E(M_{n+1}^2) + E(M_n^2) + 2\mu E(M_n) + E(X^2) \end{aligned}$$

so that

$$\begin{aligned} 0 &\leq \text{Var}((M_n + X)^-) \\ &= 2\mu\text{E}(M_{n+1}) + \sigma^2 - (\text{E}M_{n+1} - \text{E}M_n)^2 - (\text{E}(M_{n+1}^2) - \text{E}(M_n^2)) \\ &\leq 2\mu\text{E}(M_{n+1}) + \sigma^2. \end{aligned}$$

Hence, $\text{E}(M_{n+1}) \leq \sigma^2/(2|\mu|)$ for all n and the result now follows by passing to the limit using monotone convergence. ■

The proof of the following converse to Theorem 2 is due to Michael Klass. In the proof it is seen that if $\text{E}(X^+)^2 = \infty$, then the simple rule $T = \min\{n \geq 0 : S_n > 0\}$ has infinite expected return, $\text{E}S_T^+ = \infty$.

Theorem 2'. *Let X, X_1, X_2, \dots be i.i.d. with finite first moment $\mu < 0$. Then, $\text{E}M < \infty$ implies $\text{E}(X^+)^2 < \infty$.*

Proof. Let $T = \min\{n \geq 0 : S_n > 0\}$ with $T = \infty$ if $S_n \leq 0$ for all n . Then $\text{E}M < \infty$ implies that $\text{E}S_{T\text{I}}(T < \infty) < \infty$. But,

$$\begin{aligned} \text{E}S_{T\text{I}}(T < \infty) &= \sum_n \text{E}S_{T\text{I}}(T = n) \\ &\geq \sum_n \text{E}S_{T\text{I}}(T = n, S_{n-1} > 2n\mu, X_n > -3n\mu) \\ &\geq \sum_n n|\mu|\text{P}(T = n, S_{n-1} > 2n\mu, X_n > -3n\mu) \\ &= \sum_n n|\mu|\text{P}(n \leq T \leq \infty, S_{n-1} > 2n\mu, X_n > -3n\mu) \\ &= \sum_n n|\mu|\text{P}(n \leq T \leq \infty, S_{n-1} > 2n\mu)\text{P}(X_n > -3n\mu). \end{aligned}$$

But $\text{P}(n \leq T \leq \infty, S_{n-1} > 2n\mu) \rightarrow \text{P}(T = \infty) > 0$ as $n \rightarrow \infty$, so that

$$\sum_n n\text{P}(X_n > -3n\mu) < \infty,$$

which implies that $\text{E}(X^+)^2 < \infty$. ■

In the restatement of Theorem 3, we put $c = 1$ without loss of generality. The proof of the “only if” part of the theorem was suggested by Thomas Liggett.

Theorem 3. *Let Z, Z_1, Z_2, \dots be i.i.d., let $\epsilon, \epsilon_1, \epsilon_2, \dots$ be i.i.d. Bernoulli with $p = \text{P}(\epsilon = 1) = 1 - \text{P}(\epsilon = 0)$, with $0 < p < 1$. Let the $\{Z_j\}$ and $\{\epsilon_j\}$ be independent, and let $X_0 = 0$ and $X_n = \epsilon_n(X_{n-1} + Z_n)$ for $n = 1, 2, \dots$. Then,*

$$\text{E} \sup_n (X_n - n) < \infty \quad \text{if and only if} \quad \text{E}(Z^+)^2 < \infty.$$

Proof. First, suppose that $E(Z^+)^2 < \infty$. Since X_n is the sum of the Z_j since the last $\epsilon = 0$, it has the same distribution as $\sum_1^{\min(K,n)} Z_j$, where K represents the distance back from n to the most recent failure if any. We may take $K \geq 0$ to have a geometric distribution with success probability p and to be independent of the Z_j . This latter sum is less than $Q = \sum_1^K Z_j^+$.

$$\begin{aligned} E \sup_n (X_n - n) &\leq \sum_{x=0}^{\infty} P(\sup_n (X_n - n) \geq x) \\ &\leq \sum_{x=0}^{\infty} \sum_{n=0}^{\infty} P(X_n - n \geq x) \\ &\leq \sum_{x=0}^{\infty} \sum_{n=0}^{\infty} P(Q \geq x + n). \end{aligned}$$

This last sum is finite since the variance of Q is finite:

$$\begin{aligned} E(Q^2) &= E(E(Q^2|K)) = E(E(\sum_{i=1}^K \sum_{j=1}^K Z_i^+ Z_j^+ | K)) \\ &= EK E(Z^+)^2 + EK(K-1)(EZ^+)^2 < \infty. \end{aligned}$$

Conversely, note that

$$\sup_n (X_n - n) \geq \sup_n (I(\epsilon_n = 1, \epsilon_{n-1} = 0)(Z_n - n)).$$

This latter sup is equal in distribution to $\sup_n (Z_n - K_n)$, where K_n is the time of the n th appearance of the pattern $\epsilon_{j-1} = 0, \epsilon_j = 1, K_n$ starting with $\epsilon_0 = 0$. Therefore, the differences $K_2 - K_1, K_3 - K_2, \dots$ are i.i.d. with finite second moment. For an arbitrary positive number c' , we have

$$\sup(Z_n - c'n) \leq \sup(Z_n - K_n) + \sup(K_n - c'n).$$

The expectation of the first term on the right is finite since it is bounded above by $E \sup_n (X_n - n) < \infty$. The expectation of the second term is finite from Theorem 2, provided c' is chosen large enough, say $c' \geq 2\mu$, where $\mu = E(K_2 - K_1)$. Hence, $E \sup(Z_n - c'n) < \infty$ so that from the converse part of Theorem 1, $E(Z^+)^2 < \infty$. ■

We precede the proof of Theorem 4 by two lemmas. Let X_1, X_2, \dots be i.i.d. with finite mean and let $S_n = \sum_1^n X_j$.

Lemma 1. For $j \leq n$, $E\{S_j^+/j | S_n, S_{n+1}, \dots\} \geq S_n^+/n$.

Proof. By symmetry, $E\{X_j | S_n, S_{n+1}, \dots\}$ is the same for all $j \leq n$, and since the sum is S_n , we must have $E\{X_j | S_n, S_{n+1}, \dots\} = S_n/n$. Hence, $E\{S_j/j | S_n, S_{n+1}, \dots\} = S_n/n$, and hence $S_n^+/n \leq E\{S_j^+/j | S_n, S_{n+1}, \dots\}$ (since $(EX)^+ \leq EX^+$). ■

Lemma 2. Let $A_n = \{\sup_{j \leq n} S_j^+ / j \geq \lambda\}$ where $\lambda > 0$. Then $E\{X_1^+ I(A_n)\} \geq \lambda P\{A_n\}$.

Proof. Let $A_{nj} = \{S_j^+ / j \geq \lambda \text{ and } S_k^+ / k < \lambda \text{ for } k = j+1, \dots, n\}$. Then $A_n = \bigcup_{j=1}^n A_{nj}$ and $X_1^+ = S_1^+$, so

$$\begin{aligned} E\{X_1^+ I(A_n)\} &= \sum_{j=1}^n E\{S_1^+ I(A_{nj})\} \\ &= \sum_{j=1}^n E[E\{S_1^+ I(A_{nj}) | S_j, S_{j+1}, \dots\}] \\ &= \sum_{j=1}^n E[I(A_{nj}) E\{S_1^+ | S_j, S_{j+1}, \dots\}] \\ &\geq \sum_{j=1}^n E[I(A_{nj}) S_j^+ / j] \\ &\geq \lambda \sum_{j=1}^n E\{I(A_{nj})\} = \lambda P(A_n). \blacksquare \end{aligned}$$

Theorem 4. If $E\{X_1 \log^+ X_1\} < \infty$, then $E \sup_n S_n / n < \infty$.

Proof. (Doob (1953) p. 317) Let $Z_n = \sup_{j \leq n} S_j^+ / j$. Then $E Z_n < \infty$ (since $Z_n \leq \sum_1^n S_j^+ / j$).

$$\begin{aligned} E Z_n &= \int_0^\infty P(Z_n \geq z) dz \leq 1 + \int_1^\infty P(Z_n \geq z) dz \\ &\leq 1 + \int_1^\infty (1/z) E\{X_1^+ I(Z_n \geq z)\} dz \quad (\text{Lemma 2}) \\ &= 1 + E\{X_1^+ \int_1^{Z_n} (1/z) dz I(Z_n \geq 1)\} = 1 + E X_1^+ \log^+ Z_n. \end{aligned}$$

Now, note that $\log x \leq x/e$ for all $x > 0$. (There is equality at $x = 1/e$, the slopes are equal there, and $\log x$ is concave.) Replacing x by Z_n / X_1^+ , we find $X_1^+ \log^+ Z_n \leq X_1^+ \log^+ X_1 + Z_n/e$. Therefore,

$$E Z_n \leq 1 + E X_1^+ \log^+ X_1 + E Z_n / e,$$

so that

$$E Z_n \leq (e/(1-e))(1 + E X_1 \log^+ X_1).$$

But Z_n converges monotonically to $\sup_j S_j^+ / j$, so that $E \sup_j S_j / j \leq E \sup_j S_j^+ / j < \infty$. \blacksquare

The above proof uses the fact that S_n is a sum of i.i.d. random variables only through Lemma 1. Thus, the theorem is valid for sequences S_n^+ / n satisfying that lemma (a non-negative reverse supermartingale).

The following converse, due to McCabe and Shepp (1970) and Davis (1971), explicitly exhibits a simple stopping rule with $E S_N / N = \infty$.

Theorem 4'. Let X_1, X_2, \dots be i.i.d. with finite first moment and suppose $EX_1 \log^+ X_1 = \infty$. Let $c > 0$ be such that $P(X < c) > 0$, and let N denote the stopping rule $N = \min\{n \geq 1 : X_n \geq nc\}$. Then $ES_N/N = \infty$.

Proof. Without loss of generality, we take $c = 1$ since we could work as well with the sequence $X_1/c, X_2/c, \dots$. First note that $P(N = \infty) > 0$, since $P(X_1 < 1) > 0$ and

$$\begin{aligned} P(N = \infty) &= P(X_n < n \text{ for all } n) = \prod_{n=1}^{\infty} P(X_n < n) \\ &\geq \left[\prod_{n=1}^{m-1} P(X_n < n) \right] \left(1 - \sum_{n=m}^{\infty} P(X_n \geq n) \right), \end{aligned}$$

and

$$\begin{aligned} \sum_{n=m}^{\infty} P(X_n \geq n) &= \sum_{n=m}^{\infty} \int_n^{\infty} dF(x) \\ &\leq \int_m^{\infty} x dF(x) \rightarrow 0 \text{ as } m \rightarrow \infty, \end{aligned}$$

where F denotes the distribution function of X_1 . Now, choose m so that the latter sum is less than 1.

Second, note that $EX_N/N = \infty$, since

$$\begin{aligned} E((X_N/N)I\{N < \infty\}) &= \sum_1^{\infty} P(N = n)n^{-1}E\{X_n|N = n\} \\ &= \sum_1^{\infty} P(N = n)n^{-1}E\{X_n|X_n \geq n\} \\ \text{(since the } X_j \text{ are independent)} & \\ &= \sum_1^{\infty} P(N \geq n)n^{-1} \int_n^{\infty} x dF(x) \\ \text{(since } P(N = n) = P(N \geq n)P(X_n \geq n)) & \\ &\geq P(N = \infty) \sum_1^{\infty} n^{-1} \int_n^{\infty} x dF(x) \\ &= P(N = \infty) \int_1^{\infty} \sum_1^{[x]} n^{-1} x dF(x) \\ &\geq P(N = \infty) \int_1^{\infty} x \log(x) dF(x) = \infty. \end{aligned}$$

Finally, note that $E(S_N/N) = E(X_N/N) + E(\sum_{k=1}^{N-1} X_k/N)$. But,

$$\begin{aligned} E\left(\sum_{k=1}^{N-1} X_k/N\right) &= \sum_{n=1}^{\infty} P(N = n)n^{-1} \sum_{k=1}^{n-1} E(X_k|N = n) \\ &= \sum_{n=1}^{\infty} P(N = n)n^{-1} \sum_{k=1}^{n-1} E(X_k|X_k < k) \\ &\geq \sum_{n=1}^{\infty} P(N = n)n^{-1} E(X_1|X_1 < 1)(n-1) > -\infty, \end{aligned}$$

since $E(X|X < k)$ is nondecreasing in k . Thus, $E(S_N/N) = \infty$. ■