Chapter 5. MONOTONE STOPPING RULE PROBLEMS.

The ease with which the various problems of Chapter 4 were solved may be misleading. In general, stopping rule problems do not have closed form solutions and methods of finding approximate solutions must be used. Indeed, most problems without some form of Markovian structure are essentially intractable due to the high dimension of the observations involved.

In principle, it is possible to approximate a stopping rule problem by considering a truncated version of the problem. We choose a large truncation integer, T, and require stopping by stage T. More generally, for those problems for which continuing forever is a useful possibility, we would require that at stage T, the decision maker must choose between stopping at T and continuing forever, that is, between receiving Y_T and receiving $\mathrm{E}(Y_\infty|\mathcal{F}_T)$. Essentially, we replace the payoff for stopping at T by $Y_T^{(T)} = \max\{Y_T, \mathrm{E}(Y_\infty|\mathcal{F}_T)\}$, and we require stopping if stage T is reached. This is a finite horizon problem which in principle can be solved by the method of backward induction of Chapter 2. In this chapter, we find conditions under which the infinite problem may be approximated to any desired degree of accuracy by truncated problems with sufficiently large truncation points.

From a practical point of view, this method of approximating a solution by truncation is not a very good one. Solving the problem truncated at T requires computation and storage of the truncated values, $V_j^{(T)}(x_1,\ldots,x_j)$ for $j = T, T-1, \ldots, 0$ as defined by equation (1) of Chapter 2. If each x_i is allowed to assume 10 values and $T = 20$ say, this is already too large a problem for today's computers. Even if the problem can be reduced by sufficient statistics or by taking advantage of a Markovian structure, there are better methods of approximation. We consider in Section 5.1 the k-stage look-ahead rules as a simple but powerful improvement on the method of truncation. We also consider the k-time look-ahead rules.

The main topic of this chapter appears in Section 5.2. The 1- stage look-ahead rules are generally quite good; sometimes they are optimal. Under the condition that the problem be *monotone*, the 1-stage look-ahead rule is optimal for finite horizon problems. The extension of this result to infinite horizon problems requires some new conditions, detailed in Section 5.3. Essentially, these conditions are needed so that the infinite problem can be approximated by truncated problems in the sense that the limiting value of the truncated problems is the value of the original problem, $\lim_{T\to\infty} V_0^{(T)} = V^*$. Then, if the problem is monotone, the 1-stage look-ahead rule is optimal for the truncated problems, and by extension for the infinite problem as well. These ideas provide a second general method for finding simple solutions to many complex problems. (The first is the method of Chapter 4.) There are numerous applications contained in Section 5.4 and the Exercises.

§**5.1 The k-stage look-ahead rules.** For stopping rule problems, the k-stage look-ahead rule (k-sla) is described by the stopping time,

$$N_k = \min\{n \geq 0 : Y_n \geq V_n^{(n+k)}\} = \min\{n \geq 0 : Y_n \geq \mathrm{E}(V_{n+1}^{(n+k)}|\mathcal{F}_n)\}. \tag{1}$$

The k-sla is the rule which at each stage stops or continues according to whether the rule optimal among those truncated k stages ahead stops or continues. Thus at stage n, if the optimal rule among those truncated at $n+k$ continues, the k-sla continues; otherwise, the k-sla stops.

The simplest of these rules is the 1-stage look-ahead rule,

$$N_1 = \min\{n \geq 0 : Y_n \geq \mathrm{E}(Y_{n+1}|\mathcal{F}_n)\}, \tag{2}$$

In words, N_1 calls for stopping at the first n for which the return for stopping is at least as great as the expected return of continuing one stage and then stopping.

The one-stage look-ahead rule, sometimes called the myopic rule, is reasonably good, and the two- and three-stage look-ahead rules are often quite good. An important property of these rules is that *if an optimal rule exists, and if the k-sla tells you to continue, then it is optimal to continue*, for then there is at least one rule that continues and gives you at least as great an expected return as stopping at once. This property suggests a simplification of the 2-sla. Use the 1-sla until it tells you to stop, and then use the 2-sla. Similarly, the 3-sla is equivalent to : Use the 1-sla until it tells you to stop, then the 2-sla until it tells you to stop, and then use the 3-sla.

On the other hand, sometimes the 1-sla will tell you to stop, while the 2-sla, and hence the optimal rule, will tell you to continue, as examples given later will show. Therefore, it would be good to know how close to optimal the 1-sla is when it calls for stopping. Theorem 2 in Section 5.2 gives a sufficient condition for the one-stage look-ahead rule to be optimal. This theorem may be described as follows. *Suppose $V_0^{(T)} \to V^*$ as $T \to \infty$. If at some stage the 1-sla calls for stopping, and if no matter what happens in the future the 1-sla will continue to call for stopping at all future stages, then stopping immediately is optimal.* This result is true also if "1-sla" is replaced by "k-sla".

SEQUENTIAL STATISTICAL ESTIMATION. As an illustration of the computation of the 1-sla and the 2-sla, we specialize the Bayes sequential decision problems of Example 3 of Chapter 1 to the problem of statistical estimation of an unknown parameter with squared error loss. The problem is to estimate a parameter θ based on a sequentially observed sequence of random variables, $X_1, X_2, \ldots$, known to be independent and identically distributed from a distribution $F(x|\theta)$. It is assumed that θ is a real parameter and that the loss incurred when θ is estimated by a real number a is the square of the error,

$L(\theta, a) = (\theta - a)^2$. One may observe $X_1, X_2, \ldots$ for as long as desired before estimating θ at a cost of c per observation. In the Bayes approach to this problem, the prior distribution, τ, of θ is assumed known. If after observing $X_1, \ldots, X_n$, it is decided to stop and estimate θ, the estimate that minimizes the conditional expected loss, $\mathrm{E}\{(\theta - a)^2 | X_1, \ldots, X_n\}$, i.e. the Bayes estimate, is the mean of the posterior distribution of θ given $X_1, \ldots, X_n$, namely, $a = \hat{\theta}_n = \mathrm{E}(\theta | X_1, \ldots, X_n)$. The minimum expected loss is then just the conditional variance of θ, $\rho_n(X_1, \ldots, X_n) = \mathrm{Var}\,(\theta | X_1, \ldots, X_n) = \mathrm{E}\{(\theta - \hat{\theta}_n)^2 | X_1, \ldots, X_n\}$. Therefore, the total loss plus cost of stopping at stage n is

$$Y_n = \mathrm{Var}\,(\theta | X_1, \ldots, X_n) + nc. \tag{3}$$

As n increases, the posterior variance ordinarily decreases almost surely to zero, while the cost of sampling increases to ∞. The problem is to choose a stopping rule N to minimize $\mathrm{E}Y_N$.

As an example, consider estimating the mean θ of a Poisson distribution

$$f(x|\theta) = e^{-\theta}\theta^x / x! \qquad \text{for } x = 0, 1, 2, \ldots,$$

based on a sequential sample, $X_1, X_2, \ldots$, with constant cost c per observation and squared error loss. Let the prior distribution of θ be gamma, $\mathcal{G}(\alpha, 1/\lambda)$, with density

$$g(\theta) = \frac{\lambda^\alpha}{\Gamma(\alpha)} e^{-\lambda\theta}\theta^{\alpha-1} \qquad \text{on } (0, \infty).$$

It is assumed that the prior parameters, α and λ, are known. The joint density of $X_1, \ldots, X_n$ and θ is the product $g(\theta)\prod_1^n f(x_i|\theta)$, and the posterior distribution of θ given $X_1, \ldots, X_n$ is proportional to this, namely, $f(\theta|x_1, \ldots, x_n) \propto e^{-(\lambda+n)\theta}\theta^{\alpha+x_1+\cdots+x_n-1}$, which is $\mathcal{G}(\alpha + S_n, 1/(\lambda + n))$ where $S_n = \sum_1^n X_i$. Thus, past information furnished by the observations may be summarized by "up-dating" the parameters of the prior from (α, λ) to $(\alpha + S_n, \lambda + n)$. Since the mean of $\mathcal{G}(\alpha, 1/\lambda)$ is α/λ, the Bayes estimate of θ based on $X_1, \ldots, X_n$ is $\hat{\theta}_n = (\alpha + S_n)/(\lambda + n)$. Since the variance of $\mathcal{G}(\alpha, 1/\lambda)$ is α/λ^2, equation (3) becomes

$$Y_n = \frac{\alpha + S_n}{(\lambda + n)^2} + nc.$$

At stage 0, the 1-sla compares the expected loss of stopping now, $Y_0 = \alpha/\lambda^2$, with the expected loss plus cost of taking one observation and stopping, $\mathrm{E}Y_1$, where, since $\mathrm{E}X_1 = \mathrm{E}(\mathrm{E}(X_1|\theta)) = \mathrm{E}(\theta) = \alpha/\lambda$,

$$\begin{aligned} \mathrm{E}Y_1 &= \mathrm{E}(\alpha + X_1)/(\lambda + 1)^2 + c \\ &= \alpha/(\lambda(\lambda + 1)) + c. \end{aligned}$$

The 1-sla calls for stopping without taking any observations if $\alpha/\lambda^2 \le \alpha/(\lambda(\lambda+1)) + c$ or equivalently if

$$\frac{\alpha}{\lambda^2(\lambda + 1)} \le c. \tag{4}$$

The general 1-sla may be obtained from this by replacing (α,λ) by the the updated parameters $(\alpha+S_n,\lambda+n)$. Hence, the 1-sla of equation (2) with the inequality reversed since we are dealing with a cost rather than a return becomes is

$$N_1=\min\{n\ge 0:\frac{\alpha+S_n}{(\lambda+n)^2(\lambda+n+1)}\le c\}.$$

To compute the 2-sla at stage 0, we compare the expected loss of stopping without any observations, $Y_0=\alpha/\lambda^2$, with the expected loss plus cost of observing X_1 and then using the optimal 1-stage procedure, namely,

$$c+\mathrm{E}\min\{(\alpha+X_1)/(\lambda+1)^2,(\alpha+X_1)/((\lambda+1)(\lambda+2))+c\}, \tag{5}$$

The 2-sla calls for stopping without taking any observations if the former is no greater than the latter. We can simplify this expression and make the comparison of the 1-sla and the 2-sla easier by subtracting $(\alpha+X_1)/(\lambda+1)^2$ from both terms in the minimum of (5) to rewrite it as

$$\begin{aligned}&c+\mathrm{E}\min\{0,-\frac{\alpha+X_1}{(\lambda+1)^2(\lambda+2)}+c\}+\mathrm{E}(\alpha+X_1)/(\lambda+1)^2\\&\qquad=c-\mathrm{E}(\frac{\alpha+X_1}{(\lambda+1)^2(\lambda+2)}-c)^++\alpha/(\lambda(\lambda+1)).\end{aligned} \tag{6}$$

from this, it follows that the 2-sla, N_2, calls for stopping at stage 0 if

$$\frac{\alpha}{\lambda^2(\lambda+1)}\le c-\mathrm{E}(\frac{\alpha+X_1}{(\lambda+1)^2(\lambda+2)}-c)^+ \tag{7}$$

The 2-sla at stage n can be obtained from this by replacing (α,λ) by $(\alpha+S_n,\lambda+n)$ where the expectation must also be taken with the updated parameters. For a problem in which the 2-sla can be found in closed form, see Exercise 1.

THE TWO-TIMER. A simple improvement on the k-stage look-ahead rule, called the **k-time look-ahead rule**, has been suggested by A. Biesterfeld (1996). The one-time look-ahead rule, which is an improvement over both the myopic rule (the 1-sla) and the hypermetropic rule of Exercise 3.2, is the rule that calls for stopping at stage n if $Y_n\ge\sup_{t>n}\mathrm{E}(Y_t|\mathcal{F}_n)$. In other words, if at stage n there is some fixed future time $t>n$ (possibly $t=\infty$) such that the conditional expected return of continuing to stage t and stopping is greater than the return of stopping immediately, then the one-time look-ahead rule continues at least to stage $n+1$. Otherwise, it stops immediately. We denote the one-time look-ahead rule by T_1. Thus,

$$T_1=\min\{n\ge 0:Y_n\ge\sup_{t>n}\mathrm{E}(Y_t|\mathcal{F}_n)\}. \tag{8}$$

Like the one-stage look-ahead rule, N_1, if T_1 calls for continuing it is optimal to continue (provided an optimal rule exists). In addition, if N_1 calls for continuing, so does T_1. From Exercise 3.3, T_1 is at least as good as N_1. It is easy to see it can be better. For example, when the Y_n are degenerate with $Y_0 = 1$, $Y_1 = 0$, $Y_2 = 2$, and $Y_3 = Y_4 = \cdots = Y_\infty = 0$, then $N_1 = 0$ with a return of 1, and $T_1 = 2$ with a return of 2.

Let us see how well T_1 does in the example of statistical sequential estimation of the mean of a Poisson distribution. The observations, $X_1, X_2, \ldots$, are i.i.d. Poisson with mean θ, and the prior distribution of θ is $\mathcal{G}(\alpha, 1/\lambda)$. To find T_1, let us find the conditions under which T_1 stops at stage zero. We compute

$$\mathrm{E}(Y_t) = \mathrm{E}(\frac{\alpha + S_t}{(\lambda+t)^2} + tc) = \frac{\alpha}{\lambda(\lambda+t)} + tc$$

since $\mathrm{E}(S_t) = \mathrm{E}(\mathrm{E}(S_t|\theta)) = \mathrm{E}(t\theta) = t\mathrm{E}(\theta) = t\alpha/\lambda$. To find T_1, we find the t at which $\mathrm{E}(Y_t)$ is a minimum. (You may check that this occurs at $t = \{\sqrt{\frac{\alpha}{c\lambda} + \frac{1}{2}} - \lambda\}$, where $\{x\}$ represents the integer closest to x.) Clearly, for some α, c and λ, this can be greater than one. So is T_1 different from N_1?

Surprisingly, the answer is no. Here is why. As noted before, if N_1 says continue, then T_1 also says continue. Suppose N_1 says stop. Then $\alpha \le c\lambda^2(\lambda+1)$ from (4). This implies that for $t \ge 1$, $t\alpha \le tc\lambda^2(\lambda+t)$, or equivalently,

$$\frac{\alpha}{\lambda^2} \le \frac{\alpha}{\lambda(\lambda+t)} + tc.$$

Thus T_1 calls for stopping also.

It is somewhat disappointing that the one-time look-ahead rule does not improve upon the one-stage look-ahead rule for this example. Moreover, this holds true for rather general distributions. It holds whenever $\mathrm{E}(\mathrm{Var}\,(\theta|X_1, \ldots, X_n))$ is a convex function of n.

Therefore, to get an improvement we must look at the k-time look-ahead rule for $k > 1$. This rule may be described as follows. At stage n, choose k fixed times, $n < t_1 < \cdots < t_k$, and consider the best sequential rule among those that stop only at these times. If there exists a set of $t_1, \ldots, t_k$ for which this rule gives smaller expected loss than stopping at n then continue; otherwise stop.

Consider the two-time rule (called simply *the two-timer*) at stage $n = 0$. We choose two times, $0 < t_1 < t_2$. Then we consider the sequential rule that looks first at $X_1, \ldots, X_{t_1}$ and decides whether to stop or to continue to t_2 and stop. If we stop, we pay Y_{t_1}, and if we continue we expect to pay $\mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1})$. We stop at t_1 if the former is less than the latter. The expected loss using such a rule is $\mathrm{E}(\min\{Y_{t_1}, \mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1})\})$. Therefore, the two-timer stops without taking any observations if

$$Y_0 \le \mathrm{E}(\min\{Y_{t_1}, \mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1})\}) \qquad \text{for all } 0 < t_1 < t_2. \tag{9}$$

Let us compute the expectation on the right side of (9) in our example. From

$$\mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1}) = \frac{\alpha + S_{t_1}}{(\lambda+t_1)(\lambda+t_2)} + t_2 c,$$

we find

$$\begin{aligned}\min\{Y_{t_1}, \mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1})\} &= \min\{\frac{\alpha+S_{t_1}}{(\lambda+t_1)^2}+t_1c, \frac{\alpha+S_{t_1}}{(\lambda+t_1)(\lambda+t_2)}+t_2c\} \\ &= \frac{\alpha+S_{t_1}}{(\lambda+t_1)^2}+t_1c+\min\{0, -\frac{(t_2-t_1)(\alpha+S_{t_1})}{(\lambda+t_1)^2(\lambda+t_2)}+(t_2-t_1)c\} \\ &= \frac{\alpha+S_{t_1}}{(\lambda+t_1)^2}+t_1c-(t_2-t_1)\left(\frac{\alpha+S_{t_1}}{(\lambda+t_1)^2(\lambda+t_2)}-c\right)^+,\end{aligned}$$

from which we may compute the expectation in (9) as

$$\mathrm{E}(\min\{Y_{t_1}, \mathrm{E}(Y_{t_2}|\mathcal{F}_{t_1})\}) = \frac{\alpha}{\lambda(\lambda+t_1)}+t_1c-(t_2-t_1)\mathrm{E}\left(\frac{\alpha+S_{t_1}}{(\lambda+t_1)^2(\lambda+t_2)}-c\right)^+.$$

Thus the two-timer stops without taking observations if for all $0<t_1<t_2$, $Y_0=\alpha/\lambda^2$ is less than or equal to this, that is, if

$$\frac{\alpha t_1}{\lambda^2(\lambda+t_1)} \le t_1c-(t_2-t_1)\mathrm{E}\left(\frac{\alpha+S_{t_1}}{(\lambda+t_1)^2(\lambda+t_2)}-c\right)^+$$

or, equivalently,

$$\frac{\alpha}{\lambda^2} \le c(\lambda+t_1)-\frac{(t_2-t_1)}{t_1(\lambda+t_1)(\lambda+t_2)}\mathrm{E}\left(\alpha+S_{t_1}-(\lambda+t_1)^2(\lambda+t_2)c\right)^+. \qquad (10)$$

for all $0<t_1<t_2$. To compute the expectation on the right, note that the marginal distribution of S_{t_1} is negative binomial:

$$\begin{aligned}\mathrm{P}(S_t=x) &= \mathrm{E}(\mathrm{P}(S_t=x|\theta)) = \mathrm{E}(\frac{1}{x!}e^{-t\theta}(t\theta)^x) \\ &= \frac{t^x}{x!}\frac{\lambda^\alpha}{\Gamma(\alpha)}\int_0^\infty e^{-(\lambda+t)\theta}\theta^{\alpha+x-1}\,d\theta \\ &= \frac{t^x}{x!}\frac{\lambda^\alpha}{\Gamma(\alpha)}\frac{\Gamma(\alpha+x)}{(\lambda+t)^{\alpha+x}}\end{aligned}$$

for $x=0,1,2,\ldots$.

For $\alpha=1$, this is the geometric distribution, $\mathcal{G}(t/(\lambda+t))$. If $X\in\mathcal{G}(p)$, it is easy to compute $\mathrm{E}(X-k)^+$ for $k>-1$ as follows.

$$\begin{aligned}\mathrm{E}(X-k)^+ &= \sum_{x=\lceil k\rceil}^\infty (x-k)p^x(1-p) = p^{\lceil k\rceil}\sum_{x=\lceil k\rceil}^\infty (x-k)p^{x-\lceil k\rceil}(1-p) \\ &= p^{\lceil k\rceil}\sum_{y=0}^\infty (y+\lceil k\rceil-k)p^y(1-p) = p^{\lceil k\rceil}(\frac{p}{1-p}+\lceil k\rceil-k).\end{aligned} \qquad (11)$$

Now assume $\alpha = 1$ in (10) and put $p = t_1/(\lambda + t_1)$ and $k = (\lambda + t_1)^2(\lambda + t_2)c - 1$ into (11). We find that the two-timer calls for stopping at stage 0 if for all $0 < t_1 < t_2$,

$$\frac{1}{\lambda^2} \le c(\lambda + t_1) - \frac{(t_2 - t_1)}{t_1(\lambda + t_1)(\lambda + t_2)} \left(\frac{t_1}{\lambda + t_1}\right)^{\lceil k \rceil} (\frac{t_1}{\lambda} + \lceil k \rceil - k).$$

Now suppose λ is large, say $\lambda = 100$, and c is small, say $c = .000001$. Then $1/(\lambda^2(\lambda + 1)) < c$, so the 1-sla and the one-timer call for stopping. Then for small values of t_1 and t_2 such that $\lceil k \rceil = 1$ (i.e., $(1 + .01t_1)^2(1 + .01t_2) < 2$), the above inequality becomes

$$\begin{aligned}\frac{1}{\lambda^2} &\le c(\lambda + t_1) - \frac{(t_2 - t_1)}{(\lambda + t_1)^2(\lambda + t_2)}(\frac{t_1}{\lambda} + 1 - k) \\ &= c(\lambda + t_2) - \frac{(t_2 - t_1)}{(\lambda + t_1)^2(\lambda + t_2)}(\frac{t_1}{\lambda} + 2).\end{aligned}$$

For fixed t_2, the right side is increasing in t_1, and so the inequality is sharpest when t_1 is smallest, namely $t_1 = 1$. At $t_1 = 1$, the inequality becomes

$$1 \le c(\lambda + t_2)\lambda^2 - \frac{(t_2 - 1)\lambda}{(\lambda + 1)^2(\lambda + t_2)}(1 + 2\lambda).$$

At $t_2 = 2$, this inequality is $1 \le 1.00068$, so the 2-sla also calls for stopping. Yet the right side of this inequality decreases in t_2 until $t_2 = 41$ (value there $= .851023$), and then increases. Therefore the two-timer calls for taking the first observation.

In this example, the two-timer provides an improvement over the 2-stage look-ahead rule with no real increase in the cost of computation. But in continuous-time problems the k-time look-ahead rule plays a more fundamental role. In such problems, stopping can occur at any real time, not just at integer times; stopping does not occur in stages so the k-stage look-ahead approximations to the optimal rule are not available. The analog of the one-stage look-ahead rule in continuous-time problems is the infinitesimal look-ahead rule of Ross (1971). (See Example 5 in Section 5.4.) Although analogs to the k-stage look-ahead rules are not meaningful, the analogs to the k-time look-ahead rules exist and generally provide improvements over the infinitesimal look-ahead rule. The two-timer will generally look infinitesimally ahead as well as at a second time a positive amount into the future. See Biesterfeld (1996) for an example.

§**5.2 Monotone Stopping Rule Problems.** We seek conditions under which it is optimal to stop when the 1-sla calls for stopping. The basic condition is that the problem be monotone, a notion due to Chow and Robbins (1961). In finite horizon monotone stopping rule problems, the 1-sla is optimal.

Definition. *Let A_n denote the event $\{Y_n \ge \mathrm{E}(Y_{n+1}|\mathcal{F}_n)\}$. We say that the stopping rule problem is* **monotone** *if*

$$A_0 \subset A_1 \subset A_2 \subset \ldots \quad a.s. \qquad (12)$$

Monotone problems may be described as follows. The set A_n is the set on which the 1-sla calls for stopping at stage n (given that stage is reached). The condition $A_n \subset A_{n+1}$ means that if the 1-sla calls for stopping at stage n, then it will also call for stopping at stage $n+1$ no matter what X_{n+1} happens to be (a.s.). Similarly, $A_n \subset A_{n+1} \subset A_{n+2} \subset \ldots$ means that if the 1-sla calls for stopping at stage n, then it will call for stopping at all future stages no matter what the future observations turn out to be (a.s.).

Theorem 1. *In a finite horizon monotone stopping rule problem, the one-stage look-ahead rule is optimal.*

Proof. Suppose the horizon is J. One optimal rule is

$$N^* = \min\{n \geq 0 : Y_n \geq \mathrm{E}(V_{n+1}^{(J)}|\mathcal{F}_n)\}$$

where $V_{J+1}^{(J)} = -\infty, V_J^{(J)} = Y_J$, and by backward induction,

$$V_n^{(J)} = \max\{Y_n, \mathrm{E}(V_{n+1}^{(J)}|\mathcal{F}_n)\} \quad \text{for} \quad n = 0, 1, \ldots, J-1.$$

Fix $n < J$. If N_1 calls for continuing at n, then, since $V_{n+1}^{(J)} \geq Y_{n+1}$ a.s., N^* calls for continuing at n also. Suppose N_1 calls for stopping at n, that is, suppose A_n holds. Then, since the problem is monotone $A_{n+1}, \ldots, A_{J-1}$ also hold. Thus

$$\begin{aligned}
Y_{J-1} &\geq \mathrm{E}(Y_J|\mathcal{F}_{J-1}) = \mathrm{E}(V_J^{(J)}|\mathcal{F}_{J-1}). \quad \text{Hence} \quad V_{J-1}^{(J)} = Y_{J-1}. \\
Y_{J-2} &\geq \mathrm{E}(Y_{J-1}|\mathcal{F}_{J-2}) = \mathrm{E}(V_{J-1}^{(J)}|\mathcal{F}_{J-2}). \quad \text{Hence} \quad V_{J-2}^{(J)} = Y_{J-2}. \\
&\vdots \\
Y_n &\geq \mathrm{E}(Y_{n+1}|\mathcal{F}_n) = \mathrm{E}(V_{n+1}^{(J)}|\mathcal{F}_n). \quad \text{Hence} \quad V_n^{(J)} = Y_n.
\end{aligned}$$

Thus, N^* also calls for stopping. ∎

In particular, for a monotone stopping rule problem, the k-sla is no better than the 1-sla for any $k > 1$.

Corollary 1. *For a monotone stopping rule problem, for all $k > 1$, the k-stage look-ahead rule is equivalent to the 1-stage look-ahead rule.*

There are various problems associated with the extension of Theorem 1 to the infinite horizon case. Consider Examples 1 and 2 of Chapter 3. In both of these examples the problem is monotone because the 1-sla always tells you to continue. Yet the 1-sla is not optimal; in fact it is the worst of all stopping rules since it has you continue forever and receive nothing. Even if we assume A1 and A2, the 1-sla still might not be optimal for a monotone problem as the following two counterexamples show.

Counterexample 1. If $Y_n = 1/(n+1)$ for $n = 0, 1, \ldots,$ and $Y_\infty = 2$, then the unique optimal stopping rule is $N = \infty$ with return $V = 2$. The 1-sla always calls for stopping so the problem is monotone. But stopping at stage 0 has return $V = 1$ so the 1-sla is not optimal. ■

Clearly, it is suboptimal to stop if continuing forever gives a greater expected return, that is, if the hypermetropic rule calls for continuing. In other words, any stopping rule is improved by replacing any decision to stop by the decision to continue forever if that gives a greater expected payoff. Hence, we may replace Y_n by $\max\{Y_n, \mathrm{E}(Y_\infty|\mathcal{F}_n)\}$ without changing the problem; this merely rules out some suboptimal stopping rules. Thus, we may assume

$$Y_n \geq \mathrm{E}(Y_\infty|\mathcal{F}_n) \quad \text{a.s.}$$

without loss of generality. Given that A1 and A2 hold, we then would have the following strengthened form of A2,

$$\text{A3:}\ \lim_{n\to\infty} Y_n = Y_\infty \ \text{a.s.} \tag{13}$$

since by the martingale convergence theorem $\mathrm{E}(Y_\infty|\mathcal{F}_n) \to \mathrm{E}(Y_\infty|\mathcal{F}_\infty) = Y_\infty$ a.s. (See, for example, Chow Robbins and Siegmund (1971), p. 18.) We have assumed that Y_∞ is $\mathcal{F}_\infty$-measurable; that is, we have replaced Y_∞ by its expectation given $\mathcal{F}_\infty$.

Counterexample 2. Let K have a geometric distribution $\mathrm{P}(K = k) = 1/2^k$ for $k = 1, 2, \ldots$ and define $X_n = I(n \neq K)$ for $n = 1, 2, \ldots$. For a fixed ϵ, $0 < \epsilon < 1$, let $Y_0 = -1 + \epsilon$, $Y_n = (-2^n + \epsilon)X_n$ for $n = 1, 2, \ldots$ and $Y_\infty = -\infty$. Then A1 and A3 are satisfied and the optimal rule is obviously $N = \min\{n \geq 0 : X_n = 0\}$ having return 0. However, at stage n with $K > n$, if we continue one stage and stop, we expect $(1/2)0 + (1/2)(-2^{n+1} + \epsilon) = -2^n + \epsilon/2$ compared with $-2^n + \epsilon$ for stopping immediately. Hence, the 1-sla always calls for stopping so the problem is monotone. The 1-sla stops without taking any observations and has the return $-1 + \epsilon$. ■

In spite of these counterexamples, it is usually true that the 1-sla is optimal for monotone infinite horizon problems. What is needed is a condition to ensure that the infinite horizon problem can be approximated well by finite horizon problems in the sense that $V_0^{(J)} \to V_0^{(\infty)}$ as $J \to \infty$, where $V_0^{(J)}$ denotes the optimal return for the problem truncated at J, and $V_0^{(\infty)}$ denotes V^*, the optimal return for the infinite horizon problem. The following theorem states this formally. A similar approach is used in Bayes sequential statistical problems. (See, for example, Theorem 7.2.5 in Ferguson (1967).)

Theorem 2. *Suppose A1 and A2 are satisfied and suppose the problem is monotone. If $V_0^{(J)} \to V_0^{(\infty)}$ as $J \to \infty$, then the one-stage look-ahead rule is optimal.*

Proof. Let N^* denote the 1-sla and let N_j be the 1-sla truncated at j, $N_j = \min\{N^*, j\}$. Then N_j is the 1-sla for the problem truncated at j and so by Theorem 1, N_j is optimal for the problem truncated at j, so that $\mathrm{E}Y_{N_j} = V_0^{(j)}$. Note that N_j is an increasing sequence of stopping rules converging to N^*. Thus, as in the proof of Theorem 3.1,

$$V_0^{(\infty)} = \lim \mathrm{E}Y_{N_j} \leq \mathrm{E}\limsup Y_{N_j} \leq \mathrm{E}Y_{N^*},$$

showing that N^* is optimal. ■

§5.3 Approximation of the Infinite Problem by Finite Horizon Problems. This brings up the problem of the approximation of optimal rules by truncated rules. Counterexample 2 indicates that we need some sort of lower bounds for the Y_n. Under the extra condition that $T_n = \sup_{j\geq n}(Y_j - Y_n)$ be uniformly integrable, the infinite horizon problem can be approximated by the finite horizon problems.

Definition. *A set of random variables $\{T_n\}$ is said to be* **uniformly integrable** *(u.i.) if*

$$\sup_n \mathrm{E}\{|T_n| I(|T_n| > a)\} \to 0 \quad as \quad a \to \infty. \tag{14}$$

Note 1. If $\mathrm{E}(|T_n|) \to 0$ as $n \to \infty$, then T_n is u.i.
(*Proof.* Let $\epsilon > 0$. Find N such that $n > N$ implies $\mathrm{E}(|T_n|) < \epsilon$. For each n, find a_n such that $\mathrm{E}\{|T_n| I(|T_n| > a_n)\} < \epsilon$. Let $A = \max_{1\leq n\leq N} a_n$. Then $a > A$ implies that $\mathrm{E}\{|T_n| I(|T_n| > a)\} < \epsilon$, for all n, so that $\sup_n \mathrm{E}\{|T_n| I(|T_n| > a)\} \leq \epsilon$.)

Note 2. If $\limsup \mathrm{E}(|T_n|) = \infty$, then T_n is not u.i.
(*Proof.* Fix a. $\mathrm{E}(|T_n|) = \mathrm{E}\{|T_n| I(|T_n| \leq a)\} + \mathrm{E}\{|T_n| I(|T_n| > a)\} \leq a + \mathrm{E}\{|T_n| I(|T_n| > a)\}$, so $\sup_n \mathrm{E}\{|T_n| I(|T_n| > a)\} \geq \sup_n \mathrm{E}(|T_n|) - a = \infty$.)

Note 3. If $\mathrm{E}(|T_n|)$ stays bounded away from zero and infinity, then T_n may or may not be u.i. For example, if T_n is i.i.d. with $\mathrm{E}(|T_n|) = 1$, then T_n is u.i., but if $T_n = n$ w.p. $1/n$ and $T_n = 0$ otherwise, then $\mathrm{E}(|T_n|) = 1$ for all n, yet T_n is not u.i.

Lemma 1. *If $\{T_n\}$ is uniformly integrable and if A_n are any events such that $\mathrm{P}(A_n) \to 0$ as $n \to \infty$, then $\mathrm{E}I(A_n)|T_n| \to 0$ as $n \to \infty$.*

Proof.

$$\begin{aligned}\mathrm{E}I(A_n)|T_n| &= \mathrm{E}I(A_n)I(|T_n| \leq a)|T_n| + \mathrm{E}I(A_n)I(|T_n| > a)|T_n| \\ &\leq a\mathrm{P}(A_n) + \mathrm{E}I(|T_n| > a)|T_n|.\end{aligned}$$

From uniform integrability, for fixed $\epsilon > 0$ there exists an a such that the second term is less than ϵ for all n. Now let $n \to \infty$ and the first term tends to zero. ∎

Theorem 3. *Assume A1 and A3 and let $T_n = \sup_{j\geq n}\{Y_j - Y_n\}$. If the T_n are uniformly integrable, then $V_0^{(n)} \to V_0^{(\infty)}$ as $n \to \infty$.*

Proof. Let N denote an optimal stopping rule, and let $N(n) = \min\{N, n\}$ for $n = 1, 2, \ldots$. Then,

$$\begin{aligned}0 \leq V_0^{(\infty)} - V_0^{(n)} &\leq \mathrm{E}Y_N - \mathrm{E}Y_{N(n)} \\ &= \mathrm{E}I(n < N < \infty)(Y_N - Y_n) + \mathrm{E}I(N = \infty)(Y_\infty - Y_n) \\ &\leq \mathrm{E}I(n < N < \infty)T_n + \mathrm{E}(Y_\infty - Y_n)^+.\end{aligned}$$

The first term tends to zero by Lemma 1. If we let ϵ be an arbitrary positive number, the second term may be written as

$$\begin{aligned}\mathrm{E}I((Y_\infty - Y_n) \leq \epsilon)(Y_\infty - Y_n)^+ &+ \mathrm{E}I((Y_\infty - Y_n) > \epsilon)(Y_\infty - Y_n)^+ \\ &\leq \epsilon + \mathrm{E}I((Y_\infty - Y_n) > \epsilon)T_n.\end{aligned}$$

Since $\mathrm{P}((Y_\infty - Y_n)^+ > \epsilon) \to 0$ as $n \to \infty$ by A3, the final expectation tends to zero from Lemma 1, so the whole can be made less than 2ϵ for n sufficiently large. ■

The proof of this theorem shows more. Namely, under the conditions of that theorem, we have $\mathrm{E}Y_{N(n)} \to V_0^{(\infty)}$, i.e., one loses very little in truncating an optimal rule at sufficiently large n. In addition, it follows from this theorem that $\mathrm{E}Y_{N_k} \to V_0^{(\infty)}$, where N_k is the k-stage look-ahead rule.

The following corollary gives a simple sufficient condition for the uniform integrability of $\{T_n\}$ that is easily checked in many cases in which the payoff is a reward depending on the observations minus a constant cost per observation.

Corollary 2. *Assume A3 and suppose that $Y_n = Z_n - W_n$ where $\mathrm{E}\sup_n |Z_n| < \infty$ and W_n is nonnegative and nondecreasing a.s. Then, A1 holds and $V_0^{(J)} \to V_0^{(\infty)}$ as $J \to \infty$.*

Proof. A1 is satisfied since $Y_n \le Z_n$ is bounded above by an integrable function. Moreover, for $j > n$, $Y_j - Y_n \le Z_j - Z_n$ so that $0 \le T_n = \sup_{j \ge n}(Y_j - Y_n) \le 2\sup_n |Z_n| = Z'$, say, which has finite expectation. This implies uniform integrability since for all n $\mathrm{E}I(|T_n| > a)|T_n| \le \mathrm{E}I(|Z'| > a)Z' \to 0$ as $a \to \infty$. ■

In particular, putting $W_n \equiv 0$, we see that the strengthened form of A1, $\mathrm{E}(\sup_n |Y_n|) < \infty$, together with A3 implies that $V_0^{(J)} \to V_0^{(\infty)}$. It is worth remarking that for a cost problem, where we are trying to minimize $\mathrm{E}Y_N$, the only change required in Corollary 2 is to replace $Y_n = Z_n - W_n$ by $Y_n = Z_n + W_n$.

Consider Counterexample 2. One finds that $T_n = (2^n - \epsilon)I(K > n)$. The T_n cannot be u.i. since otherwise the 1-sla would be optimal. Indeed, one can check directly that the T_n are not u.i. For fixed a, just find n such that $2^n - \epsilon > a$; then $\mathrm{E}|T_n|I(|T_n| > a) = (2^n - \epsilon)\mathrm{P}(K > n) = 1 - (\epsilon/2^n)$.

If ϵ in Counterexample 2 were negative, then the 1-sla continues until $n = K$ and so is optimal. However, its optimality is not implied by the theorems of this section. First of all, the problem is not monotone since if you continue beyond $n = K$, the 1-sla will again have you continue. But even if we modify the problem to make it monotone by changing X_n to be $I(n < K)$, the $\{T_n\}$ are still as in the previous paragraph and so are not u.i. This provides an example of a monotone problem in which the 1-sla is optimal even though $V_0^{(J)}$ does not converge to $V_0^{(\infty)}$; that is, it provides a counterexample to the converse of Theorem 2.

§**5.4 Examples.** 1. THE BURGLAR PROBLEM. Suppose that the returns from the burglaries are i.i.d. non-negative random variables, $X_1, X_2, \ldots$ with known distribution having finite mean, μ. Let $Z_1, Z_2, \ldots$ be the random variables that indicate whether the burglar is caught, with $Z_n = 1$ denoting that the nth burglary is successful and $Z_n = 0$ indicating that he is caught during the nth burglary, and assume that the Z_n are i.i.d. with $\mathrm{P}(Z_n = 1) = \beta$ and $\mathrm{P}(Z_n = 0) = 1 - \beta$ where β is known and $0 < \beta < 1$. The payoff for stopping after the nth burglary is

$$Y_n = (\prod_1^n Z_i) \sum_1^n X_i,$$

for $n = 0, 1, \ldots,$ and we take $Y_\infty = 0$ so that A3 is satisfied. Letting $\mathcal{F}_n$ denote the σ-field generated by both $X_1, \ldots, X_n$ and $Z_1, \ldots, Z_n$, we compute on $\{\prod_1^n Z_i = 1\}$,

$$\mathrm{E}(Y_{n+1}|\mathcal{F}_n) = \mathrm{E}(Z_{n+1} \sum_1^{n+1} X_i|\mathcal{F}_n) = \beta(\sum_1^n X_i + \mu).$$

Hence, the one-stage look-ahead rule is

$$\begin{aligned} N_1 &= \min\{n \geq 0 : \sum_1^n X_i \geq \beta(\sum_1^n X_i + \mu)\} \\ &= \min\{n \geq 0 : \sum_1^n X_i \geq \beta\mu/(1-\beta)\}, \end{aligned}$$

that is, stop at the first n at which your accumulated gain is at least $\beta\mu/(1-\beta)$. Because the X_i are assumed to be non-negative, one sees that the problem is monotone. To check optimality, we note that the Y_n are bounded below by 0 and above by $\sup_n Y_n = \sum_1^M X_i$, where M is the number of successful burglaries, $M = \max\{n \geq 0 : \prod_1^n Z_i = 1\}$. Since $\mathrm{E}(\sum_1^M X_i) = \mu\mathrm{E}(M) = \mu/(1-\beta)$, the conditions of Corollary 2, with $W_n \equiv 0$, are satisfied and so N_1 is optimal.

This problem is just a version of the discounted stopping of a sum problem of Dubins and Teicher where the returns are nonnegative. Instead of a discount β, there is a probability β that your fortune is set forever to zero. Since you have no control over the Z_n and they are independent of everything else in the problem, you may as well replace them by their expected values. When you do, you find that $Y_n = \beta^n \sum_1^n X_i$, as in the problem of Dubins and Teicher.

In §4.2, this problem was seen to have an optimal rule of the simple form, $N = \min\{n > 0 : S_n \geq s_0\}$ for some s_0, whether or not the X_i are assumed to be non-negative. If the X_i may assume negative values, then the problem is not monotone and the methods of the present chapter are not capable of finding the solution directly. However, if the X_i are non-negative, then these methods may sometimes be used to solve more general problems where the X_i are dependent. (See Exercises 2 and 3.)

2. SELLING AN ASSET. Let $X_1, X_2, \ldots$ (with X_n representing the n-th offer) be i.i.d. with known distribution F having finite variance, and consider the problem of selling an asset with constant cost $c > 0$ and with recall so that $Y_n = M_n - nc$, where $M_n = \max\{X_1, \ldots, X_n\}$, $Y_0 = -\infty$ and $Y_\infty = -\infty$. Since

$$\begin{aligned} \mathrm{E}(Y_{n+1}|\mathcal{F}_n) - Y_n &= \mathrm{E}(\max\{M_n, X_{n+1}\}|\mathcal{F}_n) - M_n - c \\ &= \mathrm{E}((X_{n+1} - M_n)^+|\mathcal{F}_n) - c \\ &= \int (x - M_n)^+ \, dF(x) - c, \end{aligned}$$

the 1-sla may be written

$$N_1 = \min\{n \ge 1 : \int (x - M_n)^+ dF(x) \le c\}$$

Since $\int (x-v)^+ \, dF(x)$ is a nonincreasing function of v, this reduces to $N_1 = \min\{n \ge 1 : M_n \ge v\}$, where v satisfies $\int (x-v)^+ \, dF(x) = c$, the stopping rule already seen to be optimal in §4.1. Let us check optimality using the theorems of this section. The problem is clearly monotone. Note that Corollary 2 does not apply, since $\sup_n M_n = \infty$ if the distribution of X is unbounded. We check the condition of Theorem 3. Write

$$T_n = \sup_{j \ge n}(Y_j - Y_n) = \sup_{j \ge n}(M_j - M_n - (j-n)c)$$

and note that conditional on M_n, the right side of this equality has the same distribution as $\sup_{j>0}(M'_j - jc)$ where M'_j is the maximum of a sample of size j from the distribution of $(X - M_n)^+$. Thus, the $\mathrm{E}(T_n | M_n)$ are a.s. decreasing to zero, and since they have finite mean by Theorem 4.1, $\mathrm{E}(T_n) \to 0$ by monotone convergence, and thus the T_n are uniformly integrable.

We point out that the problem of selling an asset without recall, $Y_n = X_n - nc$, is not monotone. This is because the 1-sla can call for stopping at a very high value of X_n, but if you continue and observe a very low value of X_{n+1}, the 1-sla would certainly call for continuing. The 1-sla is

$$N'_1 = \min\{n \ge 0 : X_n > \mathrm{E}X - c\},$$

quite different from N_1 above. In fact, as noted in Chapter 4, N_1 is optimal for sampling without recall as well, since for all n, $X_n \le M_n$, and this implies that for all stopping rules, N,

$$\mathrm{E}(X_N - Nc) \le \mathrm{E}(M_N - Nc) \le \mathrm{E}(M_{N_1} - N_1 c) = \mathrm{E}(X_{N_1} - N_1 c).$$

3. THE PARKING PROBLEM. The formulation of the parking problem given in Chapter 2 does not produce a monotone stopping rule problem. This is because if the 1-sla calls for stopping at some open parking place before your destination and you continue one step, the next parking place may be filled so the 1-sla would call for continuing. This is somewhat artificial since you are not allowed to stop if the parking place is filled; a different formulation does lead to a monotone problem. We describe this formulation and at the same time give an extension due to Tamaki (1982) to the case where the destination is not known precisely. This problem does not have a finite horizon and so is not amenable to the methods of Chapter 2.

You are at the origin driving to a target destination $T > 0$. The distribution of T is known and has finite mean but the value of T is not known until you reach it. Parking places occur at random along the street at points $Z_1, Z_2, \ldots$ chosen according to a Poisson process independent of T with intensity $\lambda > 0$; that is, $Z_1, Z_2 - Z_1, Z_3 - Z_2, \ldots$ are i.i.d. with an exponential distribution of mean $1/\lambda$. Your loss is the distance you have to walk to your destination,

$$Y_n = |T - Z_n| \quad \text{for} \quad n = 0, 1, 2, \ldots \quad \text{and} \quad Y_\infty = \infty,$$

where Z_0 is defined to be 0 and $Y_0 = |T|$ represents the loss if you walk all the way.

The information available at stage n is the set of values of $Z_1, \ldots, Z_n$ and the information as to whether $Z_n < T$ or $Z_n \geq T$, and in the latter case the exact value of T. On $\{Z_n < T\}$, Y_n may be replaced by

$$\mathrm{E}\{Y_n | \mathcal{F}_n\} = \mathrm{E}\{T | Z_n, Z_n < T\} - Z_n.$$

If $Z_n \geq T$, it is clearly optimal to stop; so let us assume $Z_n < T$ and look at the 1-sla. To compute the 1-sla, we find on $\{Z_n < T\}$

$$\mathrm{E}\{Y_{n+1} | \mathcal{F}_n\} = \mathrm{E}\{|T - Z_n - Z| | Z_n, Z_n < T\},$$

where Z is exponential with mean $1/\lambda$. For arbitrary $u > 0$,

$$\begin{aligned} \mathrm{E}|u - Z| &= \mathrm{E}(u - Z) + 2\mathrm{E}(Z - u)^+ \\ &= (u - 1/\lambda) + 2P(Z > u)\mathrm{E}\{Z - u | Z > u\} \\ &= (u - 1/\lambda) + 2e^{-\lambda u}(1/\lambda), \end{aligned}$$

so that on $\{Z_n < T\}$

$$\mathrm{E}\{|T - Z_n - Z| | \mathcal{F}_n\} = \mathrm{E}\{(T - Z_n) - 1/\lambda + (2/\lambda)e^{-\lambda(T - Z_n)} | \mathcal{F}_n\}.$$

Comparing with Y_n and letting $g(x) = \mathrm{E}\{e^{-\lambda(T-x)} | T > x\}$, we see that the 1-sla is

$$N_1 = \min\{n \geq 0 : g(Z_n) \geq 1/2\}.$$

It is evident from this that the problem is monotone if $g(x)$ is monotonically nondecreasing, since the Z_n are a.s. increasing. To check the conditions of Corollary 2, modified to reflect the fact that this is a minimum problem, write $Y_n = Z'_n + W'_n$ where $Z'_n = (T - Z_n)^+$ is bounded below by zero and above by T and $W'_n = (Z_n - T)^+$ is non-decreasing a.s. Thus, the conditions of Corollary 2 are satisfied, so that the 1-sla is optimal if $g(x)$ is nondecreasing.

Consider the case where T has an exponential distribution. The lack of memory property of the exponential distribution implies that $g(x)$ is a constant, and hence nondecreasing. Thus, the exponential distribution is the borderline case. For distributions with thinner tails than the exponential, i.e. distributions with increasing failure rate, $g(x)$ will be nondecreasing; for distributions with decreasing failure rates, $g(x)$ will be nonincreasing.

4. PROOFREADING. Let us find the 1-stage look-ahead rule for Example 1.4, the proofreading problem. For this problem, the number of misprints M and the numbers of misprints detected on subsequent consecutive proofreadings $X_1, X_2, \ldots$ have a known joint distribution such that $X_j \geq 0$ and $\sum X_j \leq M$ a.s. and $\mathrm{E}M < \infty$. The cost for stopping after n proofreadings is

$$Y_n = nc_1 + (M - \sum_1^n X_j)c_2 \quad \text{and} \quad Y_\infty = \infty,$$

where $c_1 > 0$ is the cost of each proofreading, and $c_2 > 0$ is the cost of each undetected misprint.

If you stop at stage n you expect to lose

$$\mathrm{E}\{Y_n | X_1, \ldots, X_n\} = nc_1 + [\mathrm{E}\{M | X_1, \ldots, X_n\} - \sum_1^n X_j]c_2.$$

If you continue one stage and stop, you expect to lose

$$\mathrm{E}\{Y_{n+1} | X_1, \ldots, X_n\} = (n+1)c_1 + [\mathrm{E}\{M | X_1, \ldots, X_n\} - \sum_1^n X_j - \mathrm{E}\{X_{n+1} | X_1, \ldots, X_n\}]c_2.$$

The 1-sla calls for stopping if the former is not greater than the latter,

$$N_1 = \min(n \geq 0 : \mathrm{E}\{X_{n+1} | X_1, \ldots, X_n\} \leq c_1/c_2).$$

Clearly, the problem is monotone if and only if $\mathrm{E}\{X_{n+1} | X_1, \ldots, X_n\}$ is nonincreasing a.s. Assuming the problem is monotone let us check Corollary 2 for optimality of the 1-sla. Y_∞ was chosen so that A3 is satisfied. Moreover, $Y_n = Z_n + W_n$, where $W_n = nc_1$ is nondecreasing and $Z_n = (M - \sum_1^n X_j)c_2$ is bounded in absolute value by Mc_2 which has finite expectation by assumption. Hence, the 1-sla is optimal if the problem is monotone.

One case which leads to a monotone problem is mentioned in Yang, Wackerly and Rosalsky (1982). (See also the correction in Chow and Schechner (1985).) Let M have a Poisson distribution with known mean $\lambda > 0$ and for $n = 0, 1, \ldots$ let X_{n+1} have the binomial distribution with sample size $M - \sum_1^n X_j$ and known success probability p, $0 < p < 1$. Let $M_n = M - \sum_1^n X_j$ denote the number of misprints remaining after n proofreadings. Then the posterior distribution of M_n given $X_1, \ldots, X_n$ is Poisson with mean $\lambda(1-p)^n$ independent of $X_1, \ldots, X_n$ as is easily checked. Since

$$\begin{aligned} \mathrm{E}\{X_{n+1} | X_1, \ldots, X_n\} &= \mathrm{E}\{M_n p | X_1, \ldots, X_n\} \\ &= \lambda(1-p)^n p \end{aligned}$$

is a decreasing function of n, the problem is monotone. The corresponding optimal rule,

$$N_1 = \min\{n \geq 0 : \lambda p(1-p)^n \leq c_1/c_2\},$$

is a fixed sample size rule, i.e. a rule that stops at a fixed predetermined number of observations.

5. FISHING. (Starr, Wardrop and Woodroofe (1976)) Consider a lake with n fish and let $T_1,\ldots,T_n$ denote the capture times, assumed to be i.i.d. according to a given distribution function $F(t)$ on $(0,t)$. For $t>0$, let $X(t)$ denote the number of fish caught by time t, i.e. $X(t)$ is the number of T_j less than or equal to t. The payoff if you stop at time t is $Y(t)=X(t)-ct$, where $c>0$ is the cost of time. If stopping is allowed at all times, it may be optimal to stop between catches. However, as mentioned in Exercise 2 of Chapter 1, Starr and Woodroofe (1974) have shown that if F has increasing failure rate, then there is an optimal rule that stops only at catch times. In spite of this, the easy case turns out to be the case of decreasing failure rate! If F has decreasing failure rate, then it may be optimal to stop between catch times, but there is an analogue of the 1-sla for continuous time problems that is easy to compute and is optimal for this problem, namely, the infinitesimal look-ahead rule. The theory for this rule is developed by Ross (1971) based on the theory of Markov processes. Here, we find an approximation by discretizing the time axis, finding the ordinary 1-sla and passing to the limit.

Let $F(t)$ be an absolutely continuous distribution function with density $f(t)$ on the interval $(0,\infty)$. The *failure rate* (or *hazard rate*) of F at t is defined to be $h(t)=f(t)/(1-F(t))$, and may be interpreted as the instantaneous death rate of those individuals who have reached age t. Conversely, from the failure rate, $h(t)$, one can obtain the distribution function by the formula, $1-F(t)=\exp\{-\int_0^t h(s)\,ds\}$. The distribution with a constant failure rate equal to λ on $(0,\infty)$ is the exponential distribution, $F(t)=1-\exp\{-\lambda t\}$.

We say that F has *decreasing failure rate* (DFR) if the failure rate, $h(t)$, is nonincreasing, and *increasing failure rate* (IFR) if $h(t)$ is nondecreasing. Monotone failure rate distributions may be characterized in terms of the stochastic order of the residual lifetimes at t. Specifically, if F has DFR (resp. IFR), then the distribution of the residual life time at t,

$$P(T<t+\epsilon|T>t)=1-\exp\{-\int_t^{t+\epsilon}h(s)\,ds\},$$

is a nonincreasing (resp. nondecreasing) function of t, for every $\epsilon>0$.

Let ϵ be a small fixed positive number, and consider the return of stopping at time t, $Y(t)=X(t)-tc$, compared to the conditional expected return of continuing to time $t+\epsilon$ and stopping, namely

$$\mathrm{E}(Y(t+\epsilon)|\mathcal{F}_t)=X(t)+(n-X(t))\mathrm{P}(T<t+\epsilon|T>t)-(t+\epsilon)c.$$

The former is greater than or equal to the latter if and only if

$$(n-X(t))P(T<t+\epsilon|T>t)\le c\epsilon.$$

The first factor on the left, $n-X(t)$, is nonnegative and nonincreasing a.s., and the second is nonincreasing in t, provided F has DFR. Thus, in the DFR case, once this inequality becomes satisfied at some time t, it stays satisfied at all future times. This is a version of the monotonicity property.

Suppose that stopping is allowed only at times $t = k\epsilon$, for $k = 0, 1, 2, \ldots$. The 1-sla would be

$$N_1 = \min\{k\epsilon \geq 0 : (n - X(k\epsilon))\mathrm{P}(T < (k+1)\epsilon | T > k\epsilon) \leq c\epsilon\}.$$

As argued above, if the distribution of T has DFR, then the problem is monotone, and so from Corollary 2, with $Z_k = X(k\epsilon)$ bounded and $W_k = k\epsilon c$ nondecreasing, the 1-sla is optimal. This is true for all $\epsilon > 0$.

Now, noting that $\mathrm{P}(T < t+\epsilon | T > t)/\epsilon \to h(t)$ as $\epsilon \to 0$, we find as an approximation to the 1-sla when ϵ is small,

$$N^* = \min\{t \geq 0 : (n - X(t))h(t) \leq c\}.$$

This is the infinitesimal look-ahead rule. This problem is treated in Starr, Wardrop and Woodroofe (1976), where the payoff is extended to be of the form $Y(t) = g(X(t)) - c(t)$, where $g(k+1) - g(k)$ is nonincreasing in k, and c is a convex function. The corresponding optimal rule is $N^* = \min\{t \geq 0 : (n - X(t))(g(X(t)+1) - g(X(t)))h(t) \leq c'(t)\}$.

As a simple numerical example, consider the Pareto distribution, $F(t) = 1 - \lambda/t$, with decreasing hazard function, $h(t) = \lambda/t$. It is optimal to stop as soon as the number of fish left, $n - X(t)$, is less than ct/λ. For exponential distributions, $F(x) = 1 - \exp\{-\lambda x\}$, the only distributions that have both DFR and IFR, $h(t)$ is the constant λ so that the optimal rule is to stop as soon as $n - X(t) \leq c/\lambda$. This is a fixed sample size rule; we stop as soon as we catch $n - c/\lambda$ fish. For the Rayleigh distribution, $F(t) = 1 - \exp\{-t^2/2\}$, with increasing hazard function, $h(t) = t$, the problem becomes harder because the 1-sla is not optimal. But since we will only stop at catch times, the problem has a finite horizon and can be solved in principle by the method of backward induction of Chapter 2.

A related model, allowing an unknown number of fish of differing sizes with the catch time of a fish, T, being dependent on its size, Z, is due to Cozzolino (1972). The number of fish, M, has a prior Poisson distribution, $\mathrm{P}(\lambda)$, for a known $\lambda > 0$. Given $M = m$, the catch times and sizes of the fish, $(T_1, Z_1), \ldots, (T_m, Z_m)$, are i.i.d. with common known distribution function, $F(t, z)$, independent of m. (Cozzolino takes Z to be have a gamma distribution and T given $Z = z$ to have an exponential distribution with rate γz.) At time $s > 0$, the known data are the catch times and sizes of those fish, if any, whose catch times are less than or equal to s. The payoff for stopping at time s is the sum of the sizes of the fish caught by time s minus a constant cost per unit time: $Y_s = \sum_j Z_j \mathrm{I}(T_j \leq s) - cs$. See Exercise 13.

The most interesting feature of this model is that if by time s, k fish of sizes $z_1, \ldots, z_k$ have been caught at times $t_1, \ldots, t_k$ respectively, then the posterior distribution of future catch times and sizes is independent of this information and may be described as follows. The number of fish remaining, $N - k$, at time s is $\mathcal{P}(\lambda S(s))$, where $S(s) = \mathrm{P}(T > s)$ is the survival function. Given the number of fish remaining, their catch times and sizes are i.i.d. with distribution function, $(F(t,z)-F(s,z))/S(s)$, on the half-plane $\{(t,z): t > s\}$. The conclusion to be drawn from this is that if there is an optimal stopping rule, then there is an optimal fixed time stopping rule. This conclusion is independent of the payoff function.

6. A BEST-CHOICE PROBLEM — SUM-THE-ODDS. Here we consider a far-reaching generalization of the secretary problem due to Hill and Krengel (1992) and Thomas Bruss (2000). In this generalization, the observations are independent random variables, $X_1, X_2, \ldots$, taking values 0 or 1, with 0 representing failure and 1 representing success. Our goal is to stop on the last success. Let the success probabilities be denoted by $p_n = P(X_n = 1) = 1 - P(X_n = 0)$ for $n = 1, 2, \ldots$. Since we would never stop at stage n if there is an $i > n$ such that $p_i = 1$, we assume that all p_i are strictly less than 1 for all $i > 1$, but allow an initial success with probability 1. If we stop at stage n, our payoff, the probability of stopping at the last success, is for $n = 1, 2, \ldots$,

$$Y_n = X_n \prod_{i=n+1}^{\infty} (1 - p_i) = \begin{cases} \prod_{i=n+1}^{\infty} (1 - p_i) & \text{if } X_n = 1 \\ 0 & \text{if } X_n = 0 \end{cases} \tag{15}$$

and we take $Y_0 = Y_\infty = 0$. Of course, if there are an infinite number of successes, then no stopping rule can achieve the goal of stopping on the last success. By the Borel-Cantelli Lemma, there will be a finite number of successes almost surely if and only if $\sum_1^\infty p_i < \infty$, or equivalently, $\prod_{i=2}^\infty (1 - p_i) > 0$.

Therefore, to avoid the trivial case where every rule is optimal and gives payoff 0, we assume $\sum_1^\infty p_i < \infty$.

The secretary problem is a finite horizon problem. In the formulation above, a problem is said to have finite horizon n if $p_i = 0$ for all $i > n$. The secretary problem is obtained from this by taking X_i as the indicator function of the event that the ith object is relatively best, for $i = 1, \ldots, n$. A result of Rényi (1962) states that these events are independent and the probability that the ith object is best out of the first i is $1/i$. (For a generalization of this result, see Exercise 17.) Thus the classical secretary problem is obtained in this formulation if we put

$$p_i = \begin{cases} 1/i & \text{for } i = 1, \ldots, n \\ 0 & \text{for } i > n. \end{cases} \tag{16}$$

The secretary problem is not monotone and the optimal rule is not the one-stage look-ahead rule. But this is only because if one continues from a relatively best option, the next option may not be relatively best, in which case stopping at the next stage is obviously bad. In one of the early treatments of the secretary problem, Dynkin (1963) shows that the problem can be interpreted so that it is monotone simply by not allowing stopping on an observation that is not relatively best. With this restriction, the problem becomes monotone and the optimal rule is the one-stage look-ahead rule.

In the structure for the theory of optimal stopping as treated in Chapter 3, stopping is allowed after each observation. So to change the above problem so that stopping is allowed only after a success is observed, we change the definition of an observation. We pretend that the observations are the times at which successes occur, say $T_1, T_2, \ldots$, where T_k is the time at which the kth success occurs. Let K denote the time at which the last success occurs, with $K = \infty$ if no successes occur. We put $T_j = \infty$ if $T_n = K$ and $j > n$. With this change, the 1-sla may be computed as follows. The payoff if we stop at $T_n = t$ is 1 if $K = T_n$ and 0 otherwise. So the expected payoff at the time of the nth success is

$$Y_n = P(K = t | T_n = t) = \begin{cases} \prod_{i=t+1}^{\infty}(1 - p_i) & \text{if } t < \infty \\ 0 & \text{if } t = \infty, \end{cases} \tag{17}$$

and we take $Y_\infty = 0$. If we continue from $T_n = t < \infty$ and stop at T_{n+1}, we expect to receive

$$\begin{aligned} P(K = T_{n+1} | T_1, \ldots, T_n = t) &= p_{t+1} \prod_{t+2}^{\infty}(1 - p_i) + (1 - p_{t+1})p_{t+2} \prod_{t+3}^{\infty}(1 - p_i) + \cdots \\ &= [\prod_{i=t+1}^{\infty}(1 - p_i)] \sum_{i=t+1}^{\infty} \frac{p_i}{1 - p_i} \end{aligned} \tag{18}$$

Therefore, the one-stage look-ahead rule is

$$\begin{aligned} N_1 &= \min\{n \geq 0 : \sum_{i=T_n+1}^{\infty} \frac{p_i}{1 - p_i} \leq 1\} \\ &= \min\{t \geq 1 : X_t = 1 \quad \text{and} \quad \sum_{i=t+1}^{\infty} \frac{p_i}{1 - p_i} \leq 1\}. \end{aligned} \tag{19}$$

This is the rule that stops on a success at time t if the sum of the odds, $p_i/(1 - p_i)$, from $i = t + 1$ to infinity is is less than or equal to 1.

We now show that the rule N_1 is optimal. Let $r_i = p_i/(1-p_i)$ be odds of success on the ith trial. First, the problem is monotone because $\sum_{i=T_n+1}^{\infty} r_i \leq 1$ implies $\sum_{i=T_{n+1}+1}^{\infty} r_i \leq 1$. A_3 is satisfied because $Y_n \to 0 = Y_\infty$ a.s. as $n \to \infty$ and we are assuming that $\sum_1^{\infty} p_i < \infty$. Finally, the conditions of Corollary 2 hold with $W_n = 0$ since $|Y_n| \leq 1$.

As an example, consider the secretary problem with the p_i given by (16). The stopping rule (19) reduces to

$$N_1 = \min\{t \geq 1 : X_t = 1 \quad \text{and} \quad \sum_{i=t+1}^{n} \frac{(1/i)}{1 - (1/i)} \leq 1\}.$$

This is the rule that stops at time t if the tth object is relatively best, and if $\sum_{i=t+1}^{n} 1/(i - 1) \leq 1$. This agrees with the rule found in equation (4) of Chapter 2.

§5.5 Exercises.

1. *Bayesian Estimation of the mean of an exponential distribution.* Consider the problem of sequential statistical estimation of a real parameter θ with squared error loss and constant cost $c > 0$ per observation, so that the loss for stopping at n is given by (3). For given $\theta > 0$, let $X_1, X_2, \ldots$ be i.i.d. according to an exponential distribution with mean θ and density $f(x|\theta) = (1/\theta)\exp\{-x/\theta\}$ for $x > 0$. Let the prior density of θ be the reciprocal gamma distribution with density

$$g(\theta) = \frac{\lambda^\alpha}{\Gamma(\alpha)} \exp\{-\lambda/\theta\}\theta^{-(\alpha+1)}, \quad \theta > 0,$$

denoted by $\mathcal{G}^{-1}(\alpha, \lambda)$. Suppose that $\alpha > 2$.
(a) Show that the posterior distribution of θ given $X_1, \ldots, X_n$, is reciprocal gamma, $\mathcal{G}^{-1}(\alpha + n, \lambda + S_n)$, where $S_n = \sum_1^n X_i$.
(b) Show $Y_0 = \mathrm{Var}\,(\theta) = \lambda^2/((\alpha-1)^2(\alpha-2))$.
(c) Show that the 1-sla calls for stopping without taking any observations if and only if

$$\lambda^2/(\alpha(\alpha-1)^2(\alpha-2)) \le c.$$

(d) Show that the 2-sla calls for stopping without taking any observations if and only if

$$\lambda^2/(\alpha(\alpha-1)^2(\alpha-2)) \le c - (2c/(\alpha-2))(\lambda^2/(c(\alpha+1)\alpha^2(\alpha-1)))^{\alpha/2}.$$

2. *The adaptive burglar.* A burglar moves to a new city where he does not know precisely the distribution of the rewards he might expect from his burglaries, $X_1, X_2, \ldots$, but he is willing to assume that they are i.i.d. with an exponential density $f(x|\theta) = \theta e^{-\theta x} I(x > 0)$ for some unknown θ, whose distribution in turn he is willing to approximate by the gamma, $\mathcal{G}(1, 1/\lambda)$, with density $g(\theta) = \lambda e^{-\lambda\theta} I(\theta > 0)$ for some known λ. Let β denote the probability he is caught, $0 < \beta < 1$, and write his reward for stopping at n in the Dubins/Teicher form, $Y_n = \beta^n S_n$ for $n = 0, 1, 2, \ldots$, where $S_n = X_1 + \ldots + X_n$, and $Y_k = 0$.
(a) Show that the problem is monotone.
(b) Note that $\mathrm{E}X_1 = \infty$ so that A1 is not satisfied. But since the burglar should be especially eager to take the first observation anyway, let us suppose that he has done so and that $X_1 = x_1$ is known. Show that the 1-sla is optimal in this conditional problem.

3. *Vingt-et-un.* The burglar problem may also be generalized to allow capture times to be dependent on the rewards. (See Taylor (1975).) We also generalize to a concave utility of the accumulated payoff. Let the returns for the burglaries, $X_1, X_2, \ldots$, be nonnegative i.i.d. random variables with known distribution function, F, and finite mean, $\mu > 0$, and let $S_n = X_1 + \cdots + X_n$. Let T denote the random integer-valued time of capture and assume that for $n = 0, 1, \ldots$,

$$\mathrm{P}(T = n+1 | X_1, X_2, \ldots, T > n) = 1 - r(S_n),$$

where $0 \le r(z) \le 1$ is a known nonincreasing function on $[0, \infty)$ not identically 1. (In the casino games vingt-et-un and blackjack, one loses when one's total card count exceeds 21. Such games may be modeled as this burglar problem with $r(z) = \mathrm{P}(z + X \le 21)$.)
(a) *Maximizing utility of reward.* Let $u(z)$ denote the burglar's utility of retiring with an accumulated gain of z, where u is assumed to be a concave nondecreasing function on $[0, \infty)$, normalized so that $u(0) = 0$. Thus,

$$\begin{aligned} Y_n &= u(S_n) I(T > n) \quad \text{for} \quad n = 0, 1, \ldots \\ Y_\infty &= 0. \end{aligned}$$

Find the 1-sla and show it is optimal.
(b) *Maximizing the duration of operation.* Suppose instead that

$$\begin{aligned} Y_n &= u(n) I(T > n) \quad \text{for} \quad n = 0, 1, \ldots \\ Y_\infty &= 0, \end{aligned}$$

where $u(n)$ is such that $u(0) = 0$ and $u(n)/u(n+1)$ is nondecreasing. We also need a condition to make $\mathrm{E}(u(T)) < \infty$, for example, $u(n)/u(n+1) \to 1$. (The function $u(n) = n$, where we are trying to maximize the duration of operation, satisfies these conditions.) Find the 1-sla and show it is optimal.

4. *Adaptively selling an asset with recall.* (Ferguson (1974) and Rothschild (1974)) Offers $X_1, X_2, \ldots$ for an asset you own come in, independently drawn from a distribution F that you do not know exactly, but for which you have a Dirichlet process prior $\mathcal{D}(\alpha)$ with some parameter $\alpha = MF_0$. (All you need to know about the Dirichlet process to, solve this problem is that the expectation of $F(x)$ is $F_0(x)$ and that the posterior distribution of F given $X_1, \ldots, X_n$ is also Dirichlet but with parameter $\alpha_n = (M+n)F_n$ where $F_n = p_n F_0 + (1 - p_n) F_n^*$ and $p_n = M/(M+n)$ and F_n^* is the sample distribution function.) Assume that F_0 has a finite second moment.
(a) Assume a cost model $Y_n = M_n - nc$, find the 1-sla and show it is optimal.
(b) Assume a discount model $Y_n = \beta^n M_n$, find the 1-sla and show it is optimal.

5. *Selling several assets.* (MacQueen and Miller (1960).) A company desires to hire r persons. Interviews at a cost of $c > 0$ each yield applicants with expected worths to the company of $X_1, X_2, \ldots$, an i.i.d. sequence with known distribution function F with finite variance. Hiring is done with recall so that $Y_n = X_1^{(n)} + X_2^{(n)} + \ldots + X_r^{(n)} - nc$ for $n = r, r+1, \ldots$ and $Y_0 = \ldots = Y_{r-1} = Y_\infty = -\infty$, where $X_1^{(n)} \ge X_2^{(n)} \ge \ldots \ge X_n^{(n)}$ are the order statistics of $X_1, \ldots, X_n$. Find the 1-sla and show it is optimal.

6. *The parking problem.* Suppose the target destination $t > 0$ is known, and that the points $X_1, X_2, \ldots$ at which parking is available form a random walk; more precisely, we assume that the differences $X_1, X_2 - X_1, X_3 - X_2, \ldots$ are independent identically distributed positive random variables. The cost of stopping at stage n is taken to be

$$Y_n = g(X_n) \quad \text{for} \quad n = 1, 2, \ldots \quad (Y_0 = g(0), Y_\infty = \infty),$$

where g is a nonnegative convex function with minimum 0 at t: $g(t) = 0$. Assume that $g(x) \to \infty$ as $x \to \infty$, so that A3 is satisfied.
(a) Show that the problem is monotone.
(b) Show that the 1-sla is optimal.

7. *Finding the parking place closest to the destination.* (Boneh (1989)) Consider the parking problem with parking places occurring at sites $S_1, S_2, \ldots$ chosen according to a Poisson process with intensity λ. This time you want to find the parking place closest to the target t, assumed to be known. That is, you win if and only if you stop at the S_n closest to t. Thus,

$$Y_n = \begin{cases} 1 & \text{if } S_n > t \text{ and } S_n - t < |S_{n-1} - t| \\ 0 & \text{if } S_n > t \text{ and } S_n - t > |S_{n-1} - t| \\ \exp\{-2\lambda(t - S_n)\} & \text{if } S_n < t. \end{cases}$$

(a) Find the 1-sla.
(b) Show it is optimal.

8. *Proofreading.* (Ferguson and Hardwick (1989)) (a) Consider the proofreading problem of the text with $M \in \mathcal{P}(\lambda)$, but with the change that the success probability on the nth proofreading depends on n and is denoted by p_n. Find conditions on the p_n, assumed known, under which the problem is still monotone. Note that if the p_n are decreasing, as they well might be in practice, the 1-sla is still optimal, but that even if the 1-sla is not optimal, it is still easy to find the optimal rule because it is a fixed sample size rule.
(b) Consider a manuscript with a known number of words, W, each of which has a known probability π of being a misprint, so that M, the number of misprints, has a binomial distribution with sample size W and success probability π. Assume that the conditional distribution of the number of misprints found is as in part (a). Show that the posterior distribution of M_n given $X_1, \ldots, X_n$ is binomial. Find conditions on the p_n under which the problem is monotone. Note that the 1-sla is optimal if the p_n are nonincreasing.

9. *A search problem.* (Chew (1967)) An object is placed in a box with known probability π, $0 < \pi < 1$. You may search for it there as many times as you like. If it is there, the probability that you find it in any given search is p, $0 < p < 1$, independent of how many times you have already searched. Each search costs $c > 0$ and you win 1 if and only if you find the object.
(a) Find the 1-sla and show it is optimal.
(b) Generalize to the case where π and p are unknown with an arbitrary prior distribution.
(c) Extend to k objects placed in the box independently with probabilities π_j, $j = 1, \ldots, k$. Searches may find many objects, the objects being found independently with probabilities p_j, $j = 1, \ldots, k$, independent of the number of times you have searched and the number of objects found so far. Object j is worth x_j, $j = 1, \ldots, k$.

10. *Search for new species.* (Rasmussen and Starr (1979)) Independent observations may be drawn from a population Π consisting of finitely many species, $\Pi_1, \ldots, \Pi_s$, with probability p_j of observing a member of Π_j for $j = 1, \ldots, s$ where $p_j \geq 0$ and $\sum p_j = 1$.

There is a cost $c > 0$ for each observation, and a reward for each new species discovered. If we stop after n observations for some $n = 0, 1, \ldots$, our payoff is $Y_n = h(K_n) - nc$, where K_n is the number of distinct species observed by time n, and $h(k) \geq 0$ is the reward for having observed k distinct species. ($Y_\infty = -\infty$.) Assume that $h(k+1) - h(k)$ is nonincreasing in k.
(a) Suppose the p_j are known numbers. Find the 1-sla and show it is optimal.
(b) Suppose the p_j are unknown, but have a Dirichlet prior distribution, $\mathcal{D}(\alpha_1, \ldots, \alpha_s)$, where $\alpha_j > 0$ for all j. Find the 1-sla and show it is optimal.
(c) Show by an example that for a general prior distribution for the p_j, the 1-sla may not be optimal.

11. *Dispatching.* (See Ross (1969)) Passengers arrive randomly in the time interval (0,1) at a bus depot and wait for the next bus to leave. One bus is scheduled to leave at time 1; the problem is to choose the time t, $0 < t < 1$, at which an unscheduled bus should be dispatched in order to minimize the total expected waiting time of all the passengers. Discretize the interval (0,1) into T intervals of length $\delta = 1/T$, and let $X_j \geq 0$ denote the number of passengers who arrive in the time interval $(j-1)\delta$ to $j\delta$, for $j = 1, \ldots, T$. The total waiting time if the unscheduled bus leaves at time $t = n\delta$ is $W_n = \delta[\sum_{j=1}^{n}(n-j)X_j + \sum_{j=n+1}^{T}(T-j)X_j]$. Assume that the X_j are independent with means λ_j (e.g. $X_j \in \mathcal{P}(\lambda_j)$) for $j = 1, \ldots, T$.
(a) Find the 1-sla stopping rule N_1 for minimizing $\mathrm{E}(W_N)$ and show it is optimal provided that $\lambda_j(T-j)$ is nonincreasing in j.
(b) Suppose the λ_j are constant, equal to λ, unknown, but that λ has a known gamma prior distribution. Find the 1-sla and show it is optimal.

12. *Detecting a change-point.* (Shiryaev (1963)) The change-point K is unobservable but has known distribution $\mathrm{P}(K = k) = \pi_k$, $k = 1, 2, \ldots$, with finite mean. The observable quantities, $X_1, X_2, \ldots$ are independent given $K = k$, with $X_1, \ldots, X_{k-1}$ having a distribution with density $f_0(x)$ and $X_k, X_{k+1}, \ldots$ having a distribution with density $f_1(x)$. If you stop at n your loss is an inspection cost $c > 0$ if $n < K$, and the length of time since the change, $n - K$, if $n \geq K$. Thus, $Y_n = cI(n < K) + (n-K)I(n \geq K)$ or, conditioned on $\mathcal{F}_n$,

$$Y_n = c\mathrm{P}\{K > n|\mathcal{F}_n\} + \mathrm{E}\{(n-K)^+|\mathcal{F}_n\} \quad \text{for} \quad n = 1, 2, \ldots,$$

and we take $Y_\infty = \infty$.
(a) Find the 1-sla in the form $N_1 = \min\{n \geq 0 : U_n \geq c\}$.
(b) Find a recurrence to compute the U_n.
(c) Take K to be geometric with parameter π, $\mathrm{P}(K = n) = (1-\pi)\pi^{n-1}$ for $n = 1, 2, \ldots$. Take $f_0(x)$ to be exponential with mean 1, and $f_1(x)$ to be exponential with mean $\mu > 1$. Show the 1-sla is optimal if $c\pi\mu \leq 1 + c$.

13. *A Poisson fishing model.* (Cozzolino (1972) and Ferguson (1997)) Let $M \in \mathcal{P}(\lambda)$ and suppose that $X_1, \ldots, X_m$ given $M = m$ are i.i.d. according to a probability distribution P on some space. Let A be any measurable set in the space, and let $N(A)$ be the random variable denoting the number of $X_i \in A$. Then, $N(A)$ and $N(A^c)$ are independent with $\mathcal{P}(\mathrm{P}(A)\lambda)$ and $\mathcal{P}(\mathrm{P}(A^c)\lambda)$ distributions respectively. Moreover, given $N(A)$

and $N(A^c)$, the $X_i \in A$ (resp. the $X_i \in A^c$) are i.i.d. with distribution $\mathrm{P}/\mathrm{P}(A)$ on A (resp. $\mathrm{P}/\mathrm{P}(A^c)$ on A^c).
(a) Using this result, establish the following. Suppose the number, M, of fish in the lake is Poisson, $\mathcal{P}(\lambda)$, and given $M = m$ let the catch times and sizes of the fish, $(T_1, Z_1), \ldots, (T_m, Z_m)$, be i.i.d. with known distribution $F(t,z)$. Then at time t, independent of the number of fish caught and their catch times and sizes, the number of fish remaining in the lake is $\mathcal{P}(\lambda S(t))$, where $S(t) = \mathrm{P}(T > t)$ and given the number of fish remaining, their catch times and sizes are i.i.d. with distribution function, $\mathrm{P}(T \le s, Z \le z | T > t) = (F(s,z) - F(t,z))/S(t)$, on the set $\{s > t\}$.
(b) Suppose the joint distribution of T and Z has density $f(t,z)$ with finite $\mathrm{E}|Z|$. Let $Y_t = \sum_{j=1}^{M} Z_j I(T_j \le t) - ct$. Find $\mathrm{E}Y_t$, the expected return of the fixed time stopping rule, $N \equiv t$.
(c) Suppose T has the inverse power distribution with density, $f(t) = \theta/(1+t)^{\theta+1}$ on $(0,\infty)$, and let the distribution of Z given $T = t$ be the gamma, $\mathcal{G}(\alpha, \gamma(t+1))$, where $\theta > 0$, $\alpha > 0$, and $\gamma > 0$ are given constants. (Smaller fish are easier to catch.) Find the optimal stopping rule.
(d) Suppose the distribution of the size Z is gamma, $\mathcal{G}(\alpha,\beta)$ and that given $Z = z$ the catch time is exponential $\mathcal{G}(1, 1/(z\gamma))$. (The larger fish are easier to catch.) Find the optimal rule.

14. *Additive Damage Model.* (See Derman and Sacks (1960)) Consider a machine that accumulates observable damage. Let X_n denote the damage accrued on day n, where $X_1, X_2, \ldots$ are i.i.d. with known distribution function $F(x)$ on $[0,\infty)$. The accumulated damage on day n is $S_n = X_1 + \cdots + X_n$. When the accumulated damage exceeds a known threshold L, the machine breaks down. If the machine has not broken down by day n, it will produce a random return with mean $r_n(S_n)$ on that day. If it breaks down on that day, a penalty of b_n is assessed. Let Z denote the time of breakdown, $Z = \min\{n \ge 1 : S_n > L\}$. Then, the return for stopping on day n is

$$Y_n = \begin{cases} \sum_1^n r_j(S_j) & \text{if } n < Z \\ \sum_1^Z r_j(S_j) - b_Z & \text{if } n \ge Z. \end{cases}$$

(a) Find the 1-sla.
(b) Find some reasonable conditions on the r_n, b_n and F so that the problem is monotone.
(c) What if L is random?

15. *A Win-Lose-or-Draw Sum-the-Odds Problem.* Extend the Bruss sum-the-odds result of Example 6 Section 5.4, to the win-lose-or-draw problem of Sakaguchi. Suppose you win 1 if you stop on the last success, win nothing if you stop on a success that is not the last, and win an amount θ if you don't stop, where $0 < \theta < 1$. Assume that successes are independent events, that the probability of success on trial n is p_n and that $\sum_1^\infty p_i < \infty$. Find the 1-sla, note that it is a sum-the-odds rule, and show it is optimal.

16. *The Group Interview Secretary Problem.* (Hsiau and Yang (2000)) Consider the secretary problem in which groups of applicants are interviewed together. It is possible to select any applicant of the present group, but one may not recall applicants of previous

groups. Suppose there are n rankable applicants arranged in a completely random order, that are to be interviewed sequentially in m groups of sizes $k_1, k_2, \ldots, k_m$, where $\sum_1^m k_i = n$. You will stop only if the present group has a relatively best applicant, and if you stop you will select that applicant. Find the optimal rule.

17. *Another Sum-the-Odds Stopping Rule.* Let $Z_1, Z_2, \ldots,$ be independent random variables, and suppose Z_n has an exponential distribution, $\mathcal{E}(\theta_n)$, with density $f(z|\theta_n) = \theta_n e^{-\theta_n z}$ for $z > 0$. The θ_n are given numbers that satisfy the restriction, $\sum_1^\infty \theta_i < \infty$. The Z_n are observed sequentially, and it is desired to choose a stopping rule that maximizes the probability of stopping on the smallest Z_n. Let X_n denote the indicator of the event that $Z_n = \min_{1\le i\le n} Z_i$, that Z_n is a relatively best observation. Thus, we want to stop on the last X_n that is equal to 1.
(a) Show $P(X_n = 1) = \theta_n/(\theta_1 + \cdots + \theta_n)$.
(b) Show that the X_i are independent.
(c) Find the optimal stopping rule.

18. *A lower bound for the value of a sum-the-odds problem.* (Hill and Krengel (1992) and Bruss (2003)) When the secretary problem first appeared, it was found surprising that as the number of applicants tends to infinity, the probability of selecting the very best does not tend to zero, but instead is always greater than e^{-1}. It is even more surprising that this same bound, e^{-1}, holds for these more general problems, provided only that the sum of the odds is at least 1. In fact, there is a simple bound for finite horizon problems as well.
(a) Let $q_i = 1 - p_i$ and $r_i = p_i/(1-p_i)$. Consider the problem with finite horizon, n, and suppose that $\sum_1^n r_i \ge 1$. Let s denote an integer such that $\sum_{i=s}^n r_i \ge 1$ and $\sum_{i=s+1}^n r_i \le 1$. Show that the optimal probability of success is

$$V_n^* = \begin{cases} [\prod_{i=s}^n q_i] \sum_{i=s}^n r_i & \text{if } p_s < 1 \\ \prod_{i=s+1}^n q_i & \text{if } p_s = 1. \end{cases}$$

(b) Show that when $q_{s+1}, \ldots, q_n$ are fixed, V_n^* is increasing in p_s and so is smallest when p_s is as small as it can be made, which is just small enough to make $\sum_{i=s}^n r_i = 1$.
(c) Show that $\prod_{i=s}^n q_i$ is minimized subject to the constraint $\sum_{i=s}^n r_i = 1$ when all the q_i for $i \ge s$ are equal.
(d) Conclude that when $\sum_1^n r_i \ge 1$,

$$V_n^* \ge \left(1 + \frac{1}{n}\right)^{-n}$$

with equality if and only if $p_i = 1/(n+1)$ for $i = 1, \ldots, n$.
(e) Conclude that when $\sum_1^\infty r_i \ge 1$ and $\sum_1^\infty p_i < \infty$,

$$V_\infty^* \ge e^{-1}.$$