J. Appl. Prob. **42,** 108–120 (2005)
Printed in Israel

WHAT IS KNOWN ABOUT ROBBINS' PROBLEM?

F. THOMAS BRUSS,* *Université Libre de Bruxelles*

Abstract

Let $X_1, X_2, \ldots, X_n$ be independent, identically distributed random variables, uniform on $[0, 1]$. We observe the X_k sequentially and must stop on exactly one of them. No recollection of the preceding observations is permitted. What stopping rule τ minimizes the expected rank of the selected observation? This full-information expected-rank problem is known as *Robbins' problem*. The general solution is still unknown, and only some bounds are known for the limiting value as n tends to infinity. After a short discussion of the history and background of this problem, we summarize what is known. We then try to present, in an easily accessible form, what the author believes should be seen as the essence of the more difficult remaining questions. The aim of this article is to evoke interest in this problem and so, simply by viewing it from what are possibly new angles, to increase the probability that a reader may see what seems to evade probabilistic intuition.

Keywords: Full information; optimal selection, secretary problem; half-prophet; memoryless rule; history dependence; truncation; embedding; integral-differential equation

2000 Mathematics Subject Classification: Primary 60G40

1. Introduction

I should like to see this problem solved before I die.
Herbert Robbins, 26 June 1990

With exactly these words, Professor Herbert Robbins presented the above described problem at the International Conference on Search and Selection in Real Time, Amherst, MA, 21–27 June 1990. He also used it to finish his splendid invited final address, which attracted much attention and admiration. We here restate his problem precisely and introduce the relevant notation.

1.1. Robbins' problem

Let $X_1, X_2, \ldots, X_n$ be independent, identically distributed (i.i.d.) $U[0, 1]$ random variables, and let $\mathcal{F}_k$ be the σ-algebra generated by $X_1, X_2, \ldots, X_k$. Let R_k, $k = 1, \ldots, n$, be the absolute rank of X_k, where ranks are defined in increasing order, i.e.

$$R_k = \sum_{j=1}^{n} \mathbf{1}_{\{X_j \le X_k\}}. \tag{1}$$

Furthermore, let T_n be the set of all finite stopping rules adapted to $X_1, X_2, \ldots, X_n$, i.e.

$$T_n = \{\tau : \{\tau = k\} \in \mathcal{F}_k, 1 \le k \le n\} \quad \text{and} \quad \sum_{k=1}^{n} \mathrm{P}(\tau = k) = 1.$$

Received 8 December 2003; revision received 28 September 2004.
* Postal address: Département de Mathématique, Université Libre de Bruxelles, CP 210, B-1050 Brussels, Belgium. Email address: tbruss@ulb.ac.be

https://doi.org/10.1239/jap/1110381374 Published online by Cambridge University Press

The problem is to find the value

$$V(n) = \inf_{\tau \in T_n} \mathrm{E}[R_\tau], \tag{2}$$

including its asymptotic behavior as n tends to infinity, and also the stopping rule τ^* that achieves this value, namely

$$\tau^* := \tau_n^* = \arg\min_{\tau \in T_n} \mathrm{E}[R_\tau]. \tag{3}$$

Remarks. We usually define the largest observations to have rank 1 and allocate ranks in reverse order of the order statistics of observations. Definition (1) does the contrary. This makes no difference, of course, because the X_k are uniform i.i.d. variables on [0, 1] if and only if the variables $1 - X_k$ are too. We use definition (1) because it assigns small ranks to small observations. Also, the $U[0, 1]$ distribution for the X_k can, of course, be replaced by any other continuous distribution function F, since the transformed variables $Y_k := F(X_k)$ bring us back to the $U[0, 1]$ distribution. Finally, we note that $\mathrm{E}[R_\tau] = \mathrm{E}[\mathrm{E}[R_\tau \mid \mathcal{F}_\tau]] = \mathrm{E}[r_\tau + (n - \tau)X_\tau]$, where r_k is the relative rank of the kth observation. Hence, although R_k is not $\mathcal{F}_k$-measurable for $k < n$, the problem of finding $V(n)$ and τ^* is well defined (see (2) and (3)).

This is Robbins' problem in detail.

1.2. Objective of this paper

Our aim is stimulate interest in this problem in order to increase the probability that it will be solved. Gardner (1960) made the secretary problem known, and this had a great impact on optimal stopping. Robbins' problem belongs to the same class of problems. However, it is much harder to solve, and a complete solution may break new ground.

2. What did Robbins say about his problem?

Robbins said almost nothing about the problem. He dropped the remark that $V(n)$ approaches a value not far from 2 as n grows, but, almost as if he regretted having said this, this was all we got from him. As we shall see, this turns out to be rather modest information, and it is here that the real problems begin.

No further indication was given of why he was interested in this problem, of how he had tried to tackle it, or where in it he saw the true difficulties. Nevertheless, his severe demeanour and his suddenly apprehensive way of speaking while presenting this problem – and then his final words cited above – left a permanent impression on me. It must have been similar for many colleagues. We soon spoke of Robbins' problem as if this name had been around forever.

3. How long has the problem been around?

I (the author) could not find out whether Robbins talked about this problem before the Amherst conference in 1990. Among those I could ask, who collaborated with him or knew Robbins well otherwise, nobody could recall him discussing the problem before.

A colleague conjectured later that Robbins had seen that this problem is hard and had wanted to present it in such a way as to have it named after him, without having worked much on it or having a strong interest in it. However, though I hardly knew Professor Robbins personally, I think that this is not true. First, this does not fit the image of the truly creative scientist we know Robbins to have been. Second, it fits even less well the picture of a scientifically generous person presented so clearly by two of his most significant collaborators (see Siegmund (2003a), (2003b) and Lai and Siegmund (1986)). Finally, we will show below that Robbins had a long interest in expected-rank problems, and we will argue that he is likely to have looked into this problem as early as in the mid-sixties.

4. Robbins' problem as the fourth secretary problem

We should look at the context of Robbins' problem. Preliminary bounds for $V(n)$ will follow in a straightforward manner just by looking at related problems. These are rank-based selection problems without recall of preceding observations, often called 'secretary problems' (see, e.g. Ferguson (1989a) and Samuels (1991)). The following four problems stand out.

(I) The classical secretary problem (CSP). A decision-maker sequentially observes the relative ranks $r_1, r_2, \dots, r_n$ of the X_k and wants to maximize $\mathrm{P}(R_\tau = 1)$. Here, τ is adapted to the sub-σ-fields generated by the relative ranks. This is the best-known secretary problem. Lindley (1961) published the solution first, showing that the limiting optimal probability of choosing rank 1 is $1/\mathrm{e}$. See Gardner (1960) for a related problem. Several other independent solutions were published in later years; see, e.g. Dynkin and Juschkewitsch (1969). (Nowadays I would recommend direct use of the general odds theorem (see Bruss (2000)), of which the solution is just a corollary.) The CSP was later generalized and modified in many interesting ways (see the review by Samuels (1991)), including to the case of an unknown n (see, e.g. Presman and Sonin (1972)).

According to Ferguson (private communication), Blackwell, Dubins, and others, including himself, already knew the solution of the CSP in the late fifties. Whittle (private communication) also knew the solution before 1961. Unlike Lindley, who should get credit for this, most of the scientists involved did not expect the problem to attract as much interest as it did, which is why they did not publish their solutions. It is probable that Robbins knew the problem early, and it is certain that secretary and related problems attracted his interest (see, e.g. Vanderbei (1980) and, in particular, problems (III) and (IV), below.)

(II) The full-information secretary problem. The situation is as in problem (I), where the decision-maker wants to maximize the probability of selecting rank 1, except that he now knows the distribution of the X_k and can use the observations $X_1, X_2, \dots, X_n$ (not only their ranks). The optimal win probability is now $0.5801\cdots$ (Gilbert and Mosteller (1966)). The term 'full information' refers to the full knowledge of the distribution of the X_k. If the relative ranks are all the information that the decision-maker has, we speak of the 'no-information' case. The CSP is, hence, the corresponding no-information case. We note that the improvement of the optimal win probability from $1/\mathrm{e} = 0.3678\cdots$ (no information) to $0.5801\cdots$ (full information) is about 58%.

(III) The no-information expected-rank problem. In this problem, the decision-maker must again base his decision only on the sequential observation of relative ranks, and the problem is to minimize $\mathrm{E}[R_\tau]$. This problem is harder than (I) and (II). A heuristic approach was given by Lindley (1961), but the resulting equations were too crude to derive the existence and value of the limit. This was achieved by Chow *et al.* (1964): the limiting optimal value (minimal expected loss) is

$$\prod_{j=1}^{\infty}\left(\frac{j+2}{j}\right)^{1/(j+1)} \approx 3.8695.$$

https://doi.org/10.1239/jap/1110381374 Published online by Cambridge University Press

This paper was a little miracle, because it was finished in August 1964 but published in June 1964 (Steve Samuels (private communication)). However, another miracle cannot be hoped for: dividing 3.8695 by 1.58 (1 plus the 58% improvement observed in passing from (I) to (II)) yields about 2.44, which is greater than V, as we will see.

(IV) The full-information expected-rank problem. Here, the decision-maker can base his decision on the sequential observation of i.i.d. random variables with a fully known distribution, and wants to minimize the expected rank. Hence, this is Robbins' problem.

We see that Robbins' problem fits perfectly in the two-by-two pattern of these four related problems, generated by two typical payoff functions and the two opposite extremes of information. As Problem (III) is directly related to (IV), it is natural to consider it further. Robbins was very fond of the paper Chow *et al.* (1964) on (III) – he said so in his speech in Amherst. Among other remarks, he cited the first ten or twelve digits of the value $3.8695\cdots$. He was fond of it because this problem is relatively hard. Robbins did not only have a liking for rank-related problems, but singled out rank-payoff functions as being particularly interesting. This is evident from his vigorous, and somewhat surprising, 'Down with googol, up with such problems!' (see his comment on Ferguson (Robbins (1989)); see also Ferguson's (1989b) answer). His interest in rank-payoff functions has been confirmed by several colleagues; it is also apparent in Robbins (1991).

By the mid-sixties, had Robbins already looked at the problem naturally related to (III), that is, at 'his' problem? If he followed Pólya's golden rule, he certainly had, but we will never know for sure. However, we are sure that Robbins did not take long to discover that problem (IV) is in quite a different class – in the Hobbs class, as G. H. Hardy would have put it: easy to state and seemingly feasible, but subsequently revealed to be truly in the top class of related problems.

5. What did Robbins know about his problem?

This is a question that may be harder to answer than Robbins' problem itself, and it is an enigma to me as to why it had to turn out that way. Around 1995/6 I gave up working on Robbins' problem for good(!?), and I did not keep records. I think I had only one more occasion after the Amherst conference to speak briefly with Robbins (in around 1992), asking him what he knew. No doubt he liked to see serious interest, and I remember an encouraging 'Go on'. This left all interpretations open. For the remainder, it was a one-sided interview: I got no new information. I only heard recently from Larry Shepp that this may not have been atypical of Robbins' style. I quote (private communication): 'Robbins was not easy to talk to about problems. He liked to make jokes frequently, so one was never sure what he really thought about questions like this.'

After several other fruitless efforts to make progress on the problem, I wrote a letter to Robbins (probably in late 1995 or early 1996), saying in clear terms that we (Bruss, Ferguson, Assaf, and Samuel-Cahn) had put so much effort into his problem that I thought we deserved to be told what he knew. I got no answer, and I found this difficult to believe. Actually, I finally did get an answer, but it was not from Robbins himself but from Ester Samuel-Cahn, asking me not to be cross with Robbins. This was funny and touching at the same time. Robbins had told her about my letter when she was visiting him: this was the only reason I knew for sure that he had received it. She believed, without being sure, that Robbins knew no more than we did. Independently, Steve Samuels thinks it inconceivable that Robbins witheld part of his knowledge. Hence, this is the most probable case.

It was due to Ester's letter that I then saw Robbin's behavior in this matter as being very enigmatic, rather than unfair. Ester says Robbins liked a sarcastic sense of humor. Hence, he would pardon me that I resolved to say 'Now, Professor Robbins, I would like to know what *you* know about *your* problem before *I* die.' Sadly, both his wish and mine are now out of reach: Professor Robbins died on 12 February 2001.

6. What is known?

As far as I am aware, some ten people have worked on the problem for some time, and several others have shown interest in it. Also, the problem is the recurrent theme of small discussions (as, for example, in the conferences in Nagoya (1994), Athens (1995), Kyoto (1998), and Bedlewo, Poland (2002)). There are three papers specifically on Robbins' problem. These are Bruss and Ferguson (1993), Assaf and Samuel-Cahn (1996), and Bruss and Ferguson (1996). Also, Hubert and Pyke (1997) refer to Robbins' problem in an addendum, and add a related problem (also in the Hobbs class).

Throughout this section, we present the main results on Robbins' problem.

Proposition 1. *$V(n)$ is increasing in n.*

Proof. There is a neat nontechnical proof of this. Imagine a (modest) prophet who can, at time 0, foresee the worst observation, i.e. the largest order statistic $X_{(n)}$, but who has no other prophetical abilities. His value $V_p(n)$, say, satisfies $V_p(n) \leq V(n)$, because he can do at least as well as we can. However, $V_p(n) = V(n-1)$ because optimal behavior forces the prophet to refuse $X_{(n)}$ and to solve Robbins' problem for $n-1$ i.i.d. observations with a $U[0, X_{(n)}[$ distribution, with the same value $V(n-1)$. Hence, for $n \geq 2$, $V(n) \geq V(n-1)$.

Assaf and Samuel-Cahn (1996) had, independently, a similar argument.

Proposition 2. *$V = \lim_{n\to\infty} V_n$ exists.*

Proof. From Proposition 1, it suffices to show that $(V(n))_{n=1,2,\ldots}$ is bounded from above. Now, $V(n) \leq \tilde{V}(n)$, where $\tilde{V}(n)$ denotes the corresponding value in the rank-observation case (III), because full information includes rank information. Since $(\tilde{V}(n))_n$ converges (see Chow *et al.* (1964)), it is bounded and, hence, $(V(n))_n$ is bounded.

6.1. Upper bounds

We now look at upper bounds for V_n. From the solution of (III) in Section 4, we have the trivial upper bound $V \leq 3.8695$. The first idea to obtain better upper bounds is based on the following observation: there is a strong correlation between the values X_k and their ranks R_k. Indeed, it is easy to see (Bruss and Ferguson (1993)) that

$$\operatorname{corr}(X_k, R_k) = \sqrt{\frac{n-1}{n+1}} \to 1 \quad \text{as } n \to \infty,$$

from which we might hope that the optimal strategy to minimize $\mathrm{E}[X_{M_n}]$ over all stopping rules $M_n \in T_n$ should also be a reasonable approximation for Robbins' problem. The solution of the problem of minimizing $\mathrm{E}[X_{M_n}]$ is well known (Moser (1956)) and has asymptotic value $\inf_{M_n \in T_n} \mathrm{E}[X_{M_n}] \sim 2/n$ as n tends to infinity. Note that this is on the order of $\mathrm{E}[X_{(2,n)}]$, that is, the expected value of the second-smallest order statistic. Since, here, the loss for accepting X_k is just X_k, the optimal rule for Moser's problem is independent of the values of preceding observations and, therefore, can only depend, at each step, on the currently observed value and

https://doi.org/10.1239/jap/1110381374 Published online by Cambridge University Press

on the number of observations to come. We will call such rules *memoryless threshold rules*, and denote them by $\mathcal{M}_n$. Clearly, $\mathcal{M}_n \subset T_n$ for $n > 1$.

From the result that $\inf_{M \in \mathcal{M}_n} \mathrm{E}[X_M] \sim 2/n$, we are naturally encouraged to try the stopping rule defined by $\min\{k \geq 1\colon X_k \leq 2/(n-k+2)\}$, and this yields the following result.

Proposition 3. *The memoryless threshold rule* $\hat{M}_n$*, defined by*

$$\hat{M}_n = \min\left\{k \geq 1\colon X_k \leq \frac{2}{n-k+2}\right\},$$

yields

$$V(n) \leq \mathrm{E}[R_{\hat{M}_n}] \leq 1 + \frac{4(n-1)}{3(n+1)} \leq \tfrac{7}{3} \quad \textit{for all } n \in \mathbb{N}.$$

The proof can be found in Bruss and Ferguson (1993).

We give this rule here because it is convenient and already comes quite close to the optimum one can achieve in the class of memoryless threshold strategies. It can be improved within the same class of rules. A first improvement (Assaf and Samuel-Cahn (1996)) is obtained by replacing the thresholds $2/(n-k+2)$ by $c/(n-k+c)$ and then minimizing the corresponding expected rank with respect to c. This yields $c = 1.9469\cdots$, and $2.3318\cdots$ for the expected rank. By a more elaborate analysis, Assaf and Samuel-Cahn (1996) showed that the limiting optimal value among the class of memoryless threshold strategies, W say, satisfies $2.295 < W < 2.3267$. By extrapolation techniques, Bruss and Ferguson (1993) estimated that $W = 2.32659$. The upper bound $W < 2.3267$ of Assaf and Samuel-Cahn maintains its interest, however, because it is obtained through an explicit (complicated) threshold function that achieves this value (see Lemma 4.1 of Assaf and Samuel-Cahn (1996)). For n not too large, the precise solution is given by the next proposition.

Proposition 4. *For n fixed, let $p_k \equiv p_k(n)$, and let $N(\boldsymbol{p})$ be the memoryless stopping rule defined by $N(\boldsymbol{p}) = \inf\{k \geq 1\colon X_k \leq p_k\}$, where $\boldsymbol{p} = (p_1, p_2, \ldots, p_n)$ with $0 \leq p_k \leq 1$ for all k. For each n, the optimal memoryless rule, $\boldsymbol{p}^* = (p_1^*, p_2^*, \ldots, p_n^*)$ say, is unique and satisfies*

$$0 < p_1^* < p_2^* < \cdots < p_n^* = 1.$$

This rule can be explicitly computed from the system of equations

$$\frac{\partial \mathrm{E}[R_{N(\boldsymbol{p})}]}{\partial p_k} = 0, \qquad k = 1, 2, \ldots, n-1.$$

The proof can be found in Bruss and Ferguson (1993).

We should mention here that $\mathrm{E}[R_{N(\boldsymbol{p})}]$ can be seen to be quadratic in each variable when the others are kept fixed (see Equation (3.1) of Bruss and Ferguson (1993)), so that the $\partial \mathrm{E}[R_{N(\boldsymbol{p})}]/\partial p_k$ are easy to work with and the exact solution relatively easy to find. For large n, we would of course prefer the approximation given by Assaf and Samuel-Cahn (1996). (For general results on limiting threshold rules, also see Kennedy and Kertz (1990).)

However, an important part of Robbins' problem is, as we shall see, to prove or disprove the statement $V = W$, where, as before, W denotes the limiting optimal value within the class of memoryless threshold rules. Engaging in a full effort to compute W up to ten decimals, say, is hard to justify if $V < W$. For this reason, we state the results in a modest and safe form, as follows.

Proposition 5. $V \leq W \leq 2.3267$.

To my knowledge, nobody has found an upper bound for V better than W. At first, the failure to do so came as a real surprise. Intuitively, there should be one, and we should find one because memoryless threshold rules do not depend on the history of observations. It is this history that determines the relative rank of the stopped observation, and this is a substantial part of the loss function in Robbins' problem. Hence, it seems reasonable to expect that some modifications of memoryless threshold rules, which take relative ranks into account, will perform better, and one would think that memorizing a few smallest relative ranks would make a difference. However, this seems not to be the case. Clearly, strategies based on relative ranks alone cannot be optimal. It seems complicated to evaluate the limiting influence of modifications on possible improvements through a balanced mixture of rank information and a set of 'critical' values, and in this we had no success: trial and error can seemingly not give enough insight.

6.2. Lower bounds

Naturally, we now turn our interest to lower bounds for V. By definition, $V(n) \geq 1$ for all n and it is easy to show that, in fact, $V \geq 1.4198\cdots$. Indeed, suppose that we change from R_τ to the new loss function

$$R_{\tau,2} = \mathbf{1}_{\{R_\tau=1\}} + 2\,\mathbf{1}_{\{R_\tau\geq 2\}} = 1 + \mathbf{1}_{\{R_\tau \geq 2\}}\,.$$

Clearly, $R_{\tau,2} \leq R_\tau$ and, so, $\mathrm{E}[R_{\tau,2}] \leq \mathrm{E}[R_\tau]$. Minimizing $\mathrm{E}[R_{\tau,2}]$ over T_n means minimizing $2 - \mathrm{P}(R_\tau = 1)$ or, equivalently, maximizing $\mathrm{P}(R_\tau = 1)$ over T_n. However, this being the full-information secretary problem (II) with asymptotic value 0.580164 as n tends to infinity, we have $V \geq 2 - 0.5802 = 1.4198$.

This is the starting point of a recursive truncation procedure. We will only indicate the idea here; for the proofs, see Bruss and Ferguson (1993, pp. 623–625). For $m = 1, 2, \ldots$ and $n = m+1, m+2, \ldots$, we define a new loss function by

$$Y_k := Y_k(m,n) = 1 + \min\left\{m, \sum_{j=1}^{k-1} \mathbf{1}_{\{X_j < X_k\}}\right\} + (n-k)X_k.$$

That is, if the relative rank exceeds m, it counts for $m+1$, and the expected number of observations filling the current relative rank up to the expected rank, is, as before, $(n-k)X_k$. We note that the latter does not depend on m, whereas the truncated relative rank is clearly nondecreasing in m. Hence, the $Y_k(m,n)$ are nondecreasing in m. We note that $Y_k(m,n)$ must yield a better lower bound than that which we find by just truncating the rank. (This explains why the value $1.4198\cdots$ obtained above will not be the same as the value for $m = 1$, below.) With the weak-prophet comparison we used in the proof of Proposition 1, we can again show that the new value is also nondecreasing in n. Solving the respective minimization problems for $Y_k(m,n)$, $m = 1, 2, \ldots$, then yields a sequence of lower bounds $V^{(1)}(n), V^{(2)}(n), \ldots$, all of which are themselves lower bounds for $V(n)$ and, hence, for V. Bruss and Ferguson (1993) have done this for $m = 1, \ldots, 5$, yielding

$$V^{(1)} = 1.462, \qquad V^{(2)} = 1.658, \qquad V^{(3)} = 1.782, \qquad V^{(4)} = 1.860, \qquad V^{(5)} = 1.908.$$

Although the described truncation presents a considerable simplification of Robbins' problem, the computational problem imposes a severe constraint. The computational work increases twice exponentially, in the sense that both time and storage requirements increase exponentially in m. (Actually, in 1991, we had to out-source part of the work and were grateful to Hardwick and Schork (1992, private communication to T. S. Ferguson) for having done the computations for $m = 4$ and $m = 5$ at the University of Michigan.) Nowadays, with highly increased computation speed, we could do the same for some larger m, but we certainly could not go much further. Computers can provide valuable help in calculating approximations to the value but, here, they cannot give the answer, which is why no further effort was invested into performing the calculations for larger values of m. Moreover, no insight is gained about the 'limiting form' of strategies that might achieve V. (An independent and purely analytical approach by Assaf and Samuel-Cahn (1996) yields the lower bound $V > 1.8502$.) In summary, we have the next proposition.

Proposition 6. $V \geq 1.908$.

As we have said, this lower bound can be improved. (However, so far, we cannot obtain $V - \varepsilon$ for arbitrary $\varepsilon > 0$.) The first-order differences of the $V^{(m)}$ displayed above, i.e. 0.196, 0.124, 0.078, and 0.048, and their decreasing ratios 0.632, 0.629, and 0.615, indicate a concave shape pointing towards a value somewhere below 2. If none of the subsequent ratios were to exceed 0.615, we would obtain $V \leq 1.908 + 0.048(0.615/(1 - 0.615)) < 1.985$. This gives observational support to the conjecture that $V < 2$ (and, in particular, that $V < W$). However, we do not know if this is true. There are instances of nonmonotonic behavior of functions in Robbins' problem (Assaf and Samuel-Cahn (1996), Bruss and Ferguson (1996)) and it seems hard to say something general about these differences.

6.3. The optimal rule

The goal is now to get a better understanding of the structure of the overall optimal rule. We will summarize the known results about its form and its properties. As before, let $\tau^* := \tau_n^*$ denote the optimal stopping time for Robbins' problem with n observations. Unlike the memoryless threshold rules considered above, we must now deal with a larger class of rules based on threshold functions, which, as we shall see, depend stepwise on (essentially) the whole preceding history.

Proposition 7. *For all n, the optimal rule is a stepwise-monotone-increasing threshold function rule defined by the stopping time*

$$\tau_n^* = \inf\{1 \leq k \leq n \colon X_k \leq p_k^{(n)}(X_1, X_2, \dots, X_k)\},$$

where the functions $p_k^{(n)}(\cdot)$ satisfy

$$0 \leq p_k^{(n)}(X_1, X_2, \dots, X_k) \leq p_{k+1}^{(n)}(X_1, X_2, \dots, X_k, X_{k+1}) \quad \textit{almost surely,}$$

and

$$p_{n-1}^{(n)}(X_1, X_2, \dots, X_{n-1}) < 1 = p_n^{(n)}(\cdot) \quad \textit{almost surely.}$$

Note that 'stepwise' is added here to avoid confusion with any stronger form of monotonicity: whenever it is optimal to stop on X_k, it is optimal to stop at k for a smaller value $X_k' < X_k$ (this justifies the term 'threshold') and, for each history $X_1, X_2, \dots, X_k$ for which it is optimal to stop on X_k, it is then optimal to stop on an X_{k+1} with $X_{k+1} \leq X_k$. Although this result is intuitive, the proof is somewhat technical and, thus, the interested reader is referred to Bruss and Ferguson (1996, pp. 9–11).

6.4. Full history dependence.

Full history dependence is a strong result, actually proving that the problem is complex. It says that, for any n, the optimal rule achieving $V(n)$ depends essentially on the whole history of values, and that the preceding history allows for no simplification.

Proposition 8. *The stopping time τ_n^* is fully history dependent in the sense that, at each step k, the decision to accept or refuse X_k depends on all $X_1, X_2, \dots, X_k$, and no nontrivial statistic of $X_1, X_2, \dots, X_k$ is sufficient to achieve $V(n)$.*

Although technicalities are unavoidable in the proof (see Section 4.2 of Bruss and Ferguson (1996)), here we can summarize the idea in a simpler style. Imagine two decision-makers, one (male) having no prophetical ability (like us) and the other one (female) being a 'half-prophet' in the following sense: if she comes into the problem at stage k, with no decision to stop having been taken before, she can foresee the future observations $X_{k+1}, X_{k+2}, \dots, X_n$ *provided* that she refuses X_k and enters stage $k+1$. Thus, with respect to the expected loss by stopping, i.e. $L_k = r_k + (n-k)X_k$ (the expected rank of X_k), the two decision-makers are on the same footing at stage k. However, to compute the expected loss by going on, L_k^+ say, *he* must solve Robbins' problem at stage k. *She*, on the other hand, need only compute the expected rank generated by the history and the smallest order statistics of what follows, because if she decides to go on, she knows that she can stop at that smallest order statistic. Her future expected loss is easily found to be

$$L_k^{\mathrm{H}} = \int_0^1 \left(1 + \sum_{m=1}^{k} \mathbf{1}_{\{X_m \le s\}}\right) \mathrm{d}F_{i(k)}(s) = 1 + \sum_{m=1}^{k} (1 - X_m)^{n-k},$$

where $F_{i(k)}$ is the distribution function of the infimum of the $n-k$ i.i.d. $U[0,1]$ variables $X_{k+1}, \dots, X_n$ and, hence, $\mathrm{d}F_{i(k)}(s) = (n-k)(1-s)^{n-k-1}\,\mathrm{d}s$. She now simply compares L_k and L_k^{H} and stops if $L_k \le L_k^{\mathrm{H}}$. Note the influence of k and the X_j on the last sum. If she only came in at stage $k+1$, she would again need all preceding values; knowing L_k^{H} and X_{k-1}, for instance, is insufficient. Note also that, clearly, he must stop when she stops under optimal behavior, since he cannot possibly do better than her. The contrary is not necessarily the case, however; hence, even if she does not stop, he may still have to.

How can we make sure that, in this case, the dependence on the entire history does not partially cancel out? The trick is to look ahead from stage k to stage $n-1$. By the optimality principle, he must do this because, as we know from Proposition 7, this stage will be reached, under optimal behavior, with a strictly positive probability. At stage $n-1$, however, he has the same power as her if she only came in at step $n-1$. Hence, at stage $n-1$ their optimal behavior is the same! She needs all the preceding observations and, hence, so does he. However, he could not possibly have all the information unless he had remembered each step.

Unfortunately this is not sufficient to show that remembering all the observations makes a difference to the *asymptotic* value. That is, we can show that the probability of reaching stage $n-1$ under an optimal rule tends to 0 as $n \to \infty$. Thus, in the limit, the trick is no longer effective.

6.5. Half-prophets and finite-*n* improvements on memoryless threshold rules.

Proposition 9. *For each n, we can strictly improve on the optimal memoryless threshold rule. (However, to my knowledge, nobody has so far shown that one can find improvements that are still strictly positive in the limit as n tends to infinity.)*

https://doi.org/10.1239/jap/1110381374 Published online by Cambridge University Press

The idea is that we can improve by using a combined rule on memoryless threshold rules. We can do it, for example, by making the decree: Stop whenever a half-prophet would stop! That this indeed yields a strict improvement over the optimal memoryless threshold rule is proved in Section 5 of Bruss and Ferguson (1996). The proof is based on the half-prophet's value $\max\{L_k, L_k^{\mathrm{H}}\}$ and a result on the limiting behavior of the optimal memoryless threshold values $p_k^*(n)$. Perhaps this is a useful beginning for future work. However, so far, these improvements vanish as n tends to infinity and, hence, it is only a conjecture that a true improvement can be maintained in the asymptotic case. (Establishing the latter would prove that $V < W$, of course.)

6.6. Poisson embedding of Robbins' problem

Encouraged by the successes of embedding in the monotone subsequence problem (Bruss and Delbaen (2001)) and in the best-choice problem (see, e.g. Gnedin (2003)) we may take the following approach. We consider a planar Poisson process on $[0, t] \times [0, 1]$ with homogeneous rate 1, and interpret $[0, t]$ as the time interval and $[0, 1]$ as the value interval. We sequentially observe $(T_1, X_1), (T_2, X_2), \dots$ with $T_1 < T_2 < \cdots$, which constitutes a Poisson arrival-counting process $(N(s))_{0 \le s \le t}$ with associated values $(X_k)_{k=1,2,\dots}$. Let $v(t)$ be the value of the corresponding Robbins problem on $[0, t]$. Since we may miss the last arrival, we should enforce a loss of $\frac{1}{2}(t+1)$ for missing all the points, just to stay consistent with the discrete setting. Hence, we propose the initial condition $v(0) = \frac{1}{2}$.

We can show the following. The value $v(t)$ is Lipschitz and continuously differentiable on $]0, \infty]$, and increasing in t. Its limit v exists. Let $v(t \mid x)$ denote the optimal value conditioned on a first (artificial) arrival at time 0 with value x, which *cannot* be selected. Then, $v(t \mid x)$ is monotone decreasing in x on $[0, 1]$, with

$$v(t) + 1 = v(t \mid 0) \ge v(t \mid x) \ge v(t \mid 1) = v(t), \quad 0 \le x \le 1.$$

Also, for all $\varepsilon > 0$, $\lim_{t\to\infty}(v(t \mid \varepsilon) - v) = 0$. Now, $v(t)$ and $v(t \mid \cdot)$ satisfy the integral-differential equation

$$v'(t) = -v(t) + \int_0^1 \min\{1 + xt, v(t \mid x)\}\,\mathrm{d}x$$

with initial condition $v(0) = \frac{1}{2}$. If, in this equation, $v(t \mid x)$ were the same function as $v(t)$, we could mimick the approach of Bruss and Delbaen (2001). Hence, the question is to estimate $v(t \mid x)$ in terms of $v(t)$, t, and x sufficiently precisely. Unfortunately, this *precision problem* seems to share some of the flavor of Robbins' problem itself.

7. What is special about Robbins' problem?

Whenever one gets hooked on a problem, one of course finds it somewhat special; this is one form of selection bias. Nevertheless, I believe it is worth saying a little about Robbins' problem by way of comparison with other problems, and singling out distinctions.

In asymptotic-type analytical problems, the difficult part is typically to show that the limit of a sequence (ℓ_k), say, does exists. Once this is proved, the limit ℓ is usually either evident or else is found by some fixpoint argument $\ell = f(\ell)$, say. In the worst case, i.e. if the latter can only be solved numerically, monotonicity of (ℓ_k) usually makes things feasible. Robbins' problem has all of these helpful properties: monotonicity of $V(n)$ and the existence of V are easy to show – we even have nice bounds for V. Yet where is the $\ell = f(\ell)$ that should arise from it? Where, at least, is the recursion?

Unable to answer, we should ask the next natural question: has Robbins' problem too complicated a structure? Would this explain the difficulty we face above?

No, I think it has a comparatively easy structure, and it may not be the structure that makes the difference. If we stop with X_k and history $X_1, X_2, \ldots, X_{k-1}$ in Robbins' problem, we have a simple formula for our expected loss at time n, namely $\mathbf{1}_{\{X_1 \le X_k\}} + \cdots + \mathbf{1}_{\{X_{k-1} \le X_k\}} + 1 + (n-k)X_k$, for the appropriate value of n. So, where does intuition fail us?

The following answer is perhaps easiest to defend. Strategies are, a priori, difficult functions. There are many interesting problems in optimal stopping where these are, fortunately, still tractable. Fully history-dependent functions, however, have an unbounded number of arguments and would therefore figure in the list of most undesirable mathematical objects. It is these, as we have shown, that we have to face in Robbins' problem. It is understandable, therefore, that the theory of optimal stopping, still a relatively young discipline by the standards of mathematics, need not be able to cope easily with such problems.

I have never looked into possible solutions by stochastic approximation, because Monro and Robbins (1951) can be regarded as the fathers of stochastic approximation and, so, it is probable that Robbins would have seen this first.

Another thought: if $V < W$ then the speed of convergence of empirical distribution functions $F_n(t)$ to the limit $F(t) = t$ is bound to play an essential role. Indeed, the sequence $(F_n(t))$ describes how new points fall between the earlier ones and, thus, how patterns of points appear sequentially. It is these that must then be responsible for the 'occasions' on which memoryless threshold rules are (possibly) insufficient. If occasions of improvement quickly become less frequent as n increases, then, again in casual terms, the law of large numbers will take over, implying that $V = W$. Thus, a sound knowledge of the domain of spacings (clusters) and empirical distribution functions might be the right prerequisite to recognizing or solving the essential part of the problem. (See, e.g. Brennan and Durrett (1987); see also the constructions of proofs in Pyke and van Zwet (2004), although these are focused on Kakutani's (1976) model and deal with the $U[0, 1]$ model only for comparison.)

I think the solution to Robbins' problem may give something new. Robbins may have expected this. As I discovered later, Ronald Pyke would share this point of view. He says (see Hubert and Pyke (1997, p. 123)): 'Throughout his work, Herb shows his special knack for capturing and describing clearly the essence of a new type of problem, often by means of a deceptively simple example . . . sometimes opening up completely new directions for research.'

8. What is the essence of the open questions?

Let us single out the remaining essential questions in the form of a guide. We must prove that either

(a) $V < W$ or, alternatively,

(b) $V = W$.

If (a) holds then there are many possible further questions, but the complete answers to the following questions should really be considered to be the *complete* solution:

(i) What is V?

(ii) What is the form of the optimal strategy?

(iii) For any $n \in \mathbb{N}$ and $\varepsilon > 0$, how can we determine an $n(\varepsilon)$-memory algorithm to compute an ε-optimal strategy $\sigma_n \in T_n$, i.e. a strategy that achieves a value of at least $V(n) - \varepsilon$ by remembering, at each stage, no more than a fixed number $n(\varepsilon)$ of observed points? (The answer to the following slightly weaker question may suffice: what are the dominant features in the sequentially observed patterns of points?)

If (b) holds, fewer interesting questions remain than in the case $V < W$. Question (i) is then trivial, of course, and in answer to question (ii) we would now know that the history dependence fades as n tends to infinity; that is, the advantage we have by memorizing the whole history at each step tends to 0, and V can be obtained by the optimal memoryless threshold rule. However, it remains interesting to compute or approximate the optimal strategy for the n-problem and the value $V(n)$, that is, to answer question (iii).

If neither (a) nor (b) can be proved, we suggest finding an analytical proof that $V < 2$ or $V \geq 2$. Indeed, 2 stands out as an interesting constant just by comparison with Moser's optimal rule for finding the smallest expected value (see Section 6). What is easier: minimizing the expected rank, or minimizing the expected value?

A final word. Most of us would agree that mathematics becomes particularly interesting when teaching us something our intuition fails to grasp. In Robbins' problem, the optimal strategy may surprise us. If so, it may remind us of the apocryphal man who experienced the nature of strategies: 'My enemy had two possibilities, but then he chose the third.'

Acknowledgements

I thank Professors F. Delbaen, A. Gnedin, T. Lai, G. Louchard, E. Samuel-Cahn, S. Samuels, L. Shepp, D. Siegmund, M. Tamaki, R. Vanderbei, and, in particular, T. Ferguson, for checking the problem's history, and for helpful comments.

References

Assaf, D. and Samuel-Cahn, E. (1996). The secretary problem: minimizing the expected rank with i.i.d. random variables. *Adv. Appl. Prob.* **28,** 828–852.

Brennan, M. D. and Durrett, R. (1987). Splitting intervals. II. Limit laws for lengths. *Prob. Theory and Relat. Fields* **75,** 109–127.

Bruss, F. T. (2000). Sum the odds to one and stop. *Ann. Prob.* **28,** 1384–1391.

Bruss, F. T. and Delbaen, F. (2001). Optimal rules for the sequential selection of monotone subsequences of maximum expected length. *Stoch. Process. Appl.* **96,** 313–342.

Bruss, F. T. and Ferguson, T. S. (1993). Minimizing the expected rank with full information. *J. Appl. Prob.* **30,** 616–626.

Bruss, F. T. and Ferguson, T. S. (1996). Half-prophets and Robbins' problem of minimizing the expected rank. In *Athens Conf. Appl. Prob. Time Ser. Anal.* (Lecture Notes Statist. **114**), Vol. 1, Springer, New York, pp. 1–17.

Chow, Y. S., Moriguti, S., Robbins, H. and Samuels, S. M. (1964). Optimal selection based on relative ranks. *Israel J. Math.* **2,** 81–90.

Dynkin, E. B. and Juschkewitsch, A. A. (1969). *Sätze und Aufgaben über Markoffsche Prozesse.* Springer, Berlin.

Ferguson, T. S. (1989a). Who solved the secretary problem? *Statist. Sci.* **4,** 282–289.

Ferguson, T. S. (1989b). [Who solved the secretary problem?]: rejoinder. *Statist. Sci.* **4,** 294–296.

Gardner, M. (1960). Mathematical games. *Scientific Amer.* **3,** 202–203.

Gilbert, J. P. and Mosteller, F. (1966). Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.* **61,** 35–73.

Gnedin, A. V. (2003). Best choice from the planar Poisson process. *Stoch. Process. Appl.* **111,** 317–354.

Hubert, S. L. and Pyke, R. (1997). A particular application of Brownian motion to sequential analysis. *Statistica Sinica* **7,** 109–127.

Kakutani, S. (1976). A problem of equidistribution on the unit interval [0, 1]. In *Measure Theory* (Proc. Conf. Oberwolfach, June 1975; Lecture Notes Math. **451**), eds A. Bellow and D. Kolzow, Springer, Berlin, pp. 369–376.

Kennedy, D. P. and Kertz, R. P. (1990). Limit theorems for threshold-stopped random variables with applications to optimal stopping. *Adv. Appl. Prob.* **22,** 396–411.

Lai, T. L. and Siegmund, D. (1986). The contributions of Herbert Robbins to mathematical statistics. *Statist. Sci.* **1,** 276–284.

LINDLEY, D. (1961). Dynamic programming and decision theory. *Appl. Statist.* **10,** 39–51.

MONRO, S. AND ROBBINS, H. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22,** 400–407.

MOSER, L. (1956). On a problem of Cayley. *Scripta Math.* **22,** 289–292.

PRESMAN, E. AND SONIN, I. (1972). The best choice problem for a random number of objects. *Theory Prob. Appl.* **17,** 657–668.

PYKE, R. AND VAN ZWET, W. R. (2004). Weak convergence results for the Kakutani interval splitting procedure. *Ann. Prob.* **32,** 380–423.

ROBBINS, H. (1989). [Who solved the secretary problem?]: comment. *Statist. Sci.* **4,** 291.

ROBBINS, H. (1991). Remarks on the secretary problem. *Amer. J. Math. Manag. Sci.* **11,** 25–37.

SAMUELS, S. M. (1991). Secretary problems. In *Handbook of Sequential Analysis* (Statist. Textbooks Monogr. **118**), eds B. K. Gosh and P. K. Sen, Marcel Dekker, New York, pp. 381–405.

SIEGMUND, D. (2003a). Herbert Robbins and sequential analysis. *Ann. Statist.* **31,** 349–365.

SIEGMUND, D. (2003b). The publications and writings of Herbert Robbins. *Ann. Statist.* **31,** 407–413.

VANDERBEI, R. J. (1980). The optimal choice of a subset of a population. *Math. Operat. Res.* **5,** 481–486.