OPTIMAL STOPPING WITH RANK-DEPENDENT LOSS

ALEXANDER V. GNEDIN,* *Utrecht University*

Abstract

For τ, a stopping rule adapted to a sequence of n independent and identically distributed observations, we define the loss to be $\mathrm{E}[q(R_\tau)]$, where R_j is the rank of the jth observation and q is a nondecreasing function of the rank. This setting covers both the best-choice problem, with $q(r) = \mathbf{1}(r > 1)$, and Robbins' problem, with $q(r) = r$. As n tends to ∞, the stopping problem acquires a limiting form which is associated with the planar Poisson process. Inspecting the limit we establish bounds on the stopping value and reveal qualitative features of the optimal rule. In particular, we show that the complete history dependence persists in the limit; thus answering a question asked by Bruss (2005) in the context of Robbins' problem.

Keywords: Optimal stopping; Robbins' problem; best-choice problem; planar Poisson process

2000 Mathematics Subject Classification: Primary 60G40
Secondary 60G70

1. Introduction

Let $X_1, \ldots, X_n$ be a sequence of independent and identically distributed (i.i.d.) observations, sampled from the uniform distribution on $[0, 1]$ (in the setup of this paper this assumption covers the general case of an arbitrary continuous distribution). For $j \in [n] := \{1, \ldots, n\}$, define the *final ranks* as

$$R_j = \sum_{k=1}^{n} \mathbf{1}(X_k \leq X_j),$$

so $(R_1, \ldots, R_n)$ is an equiprobable permutation of $[n]$. Let $q : \mathbb{N} \to \mathbb{R}_+$ be a nondecreasing loss function with $q(1) < q(\infty) := \sup q(r)$. In 'secretary problems' [20] we are typically interested in the large-n behaviour of the minimum risk,

$$V_n(\mathcal{T}_n) = \inf_{\tau \in \mathcal{T}_n} \mathrm{E}[q(R_\tau)], \tag{1}$$

where $\mathcal{T}_n$ is a given class of stopping rules with values in $[n]$. Two classical loss functions are

(i) $q(r) = \mathbf{1}(r > 1)$ for the best-choice problem of maximising the probability of stopping at the minimum observation, $X_{n,1} := \min(X_1, \ldots, X_n)$;

(ii) $q(r) = r$ for the problem of minimising the expected rank.

Many results are available for the case where $\mathcal{T}_n$ in (1) is the class $\mathcal{R}_n$ of *rank rules*, which are the stopping rules adapted to the sequence of *initial ranks*,

$$I_j = \sum_{k=1}^{j} \mathbf{1}(X_k \leq X_j) = \sum_{k=1}^{j} \mathbf{1}(R_k \leq R_j), \qquad j \in [n];$$

see [8], [9], and [10]. By independence of the initial ranks, the optimal decision to stop at the jth observation depends only on I_j. The limiting risk $V_\infty(\mathcal{R}) := \lim_{n \to \infty} V_n(\mathcal{R}_n)$ has an interpretation in terms of a continuous-time stopping problem [10]. Explicit formulae for $V_\infty(\mathcal{R})$ are known in some cases, for bounded and unbounded q, including the two classical loss functions and their generalisations [2], [7], [8], [17], [18].

Much less explored are the problems where $\mathcal{T}_n$ is the class $\mathcal{F}_n$ of all stopping rules adapted to the natural filtration $(\sigma(X_1, \ldots, X_j),\ j \in [n])$. The principal difficulty here is that, for general q, the decision to stop on X_j must depend not only on X_j but also on the full vector $(X_{j-1,1}, \ldots, X_{j-1,j-1})$ of order statistics of $X_1, \ldots, X_{j-1}$. In this sense the optimal rule is *fully history dependent*. Specifically, the $\mathcal{F}_n$-optimal rule has the form

$$\tau_n = \min\{j \colon X_j < h_{n,j}(X_{j-1,1}, \ldots, X_{j-1,j-1})\} \tag{2}$$

(with $h_{n,1} = \text{constant}$ and $h_{n,n} = 1$), where $(h_{n,j},\ j \in [n])$ is a collection of functions with certain monotonicity properties. The dependence on history is reducible to the first $m - 1$ order statistics if q is *truncated* at m, i.e. $q(r) = q(m)$ for $r \geq m$, but even then the analytical difficulties are severe. The asymptotic value $V_\infty(\mathcal{F}) := \lim_{n \to \infty} V_n(\mathcal{F}_n)$ is known explicitly only for the best-choice problem (hence, for any q truncated at $m = 2$); see [12] for the formula and history. *Robbins' problem* is problem (1) with $\mathcal{T}_n = \mathcal{F}_n$ and the linear loss function $q(r) = r$; see [1], [3], [4], and [5].

The full history dependence makes explicit analysis of the $\mathcal{F}_n$-optimal rule hardly possible; thus, it is natural to seek for tractable smaller classes of rules, with some kind of reduced dependence on the history. Of course, the rank rules is one of these such classes, and the optimal rule in $\mathcal{R}_n$ is also of the form of (2), with the special feature that $h_{n,j}(x_1, \ldots, x_{j-1}) = x_{\iota_n(j)}$ (for $x_0 := 0 \leq x_1 \leq \cdots \leq x_{j-1} \leq 1$ and $1 < j \leq n$), where $\iota_n(j) \in \{0, \ldots, j-1\}$ is some threshold value of I_j, and $h_{n,1} = 0$. Another interesting possibility is to consider the class $\mathcal{M}_n$ of *memoryless rules* of the form

$$\tau = \min\{j \colon X_j < f_{n,j}\},$$

where $(f_{n,j},\ j \in [n])$ is an increasing sequence of thresholds. These rules are again of the form of (2), where the functions $h_{n,j}$ are constant. By familiar monotonicity arguments (which we recall in Section 4), the limiting value $V_\infty(\mathcal{M}) := \lim_{n \to \infty} V_n(\mathcal{M})$ (finite or infinite) exists for arbitrary q. See [15] and [19] for other classes of stopping rules with restricted dependence on history.

Memoryless rules were intensively studied in the context of Robbins' problem, in which case they were found to outperform, asymptotically, the rank rules, meaning that $V_\infty(\mathcal{M}) < V_\infty(\mathcal{R}) = 3.869 \cdots$; see [1], [4], and [5]. In a recent survey of Robbins' problem Bruss [3] stressed that a further principal step would be to either prove or disprove that $V_\infty(\mathcal{F}) < V_\infty(\mathcal{M})$. Coincidence of the asymptotic values $V_\infty(\mathcal{F}) = V_\infty(\mathcal{M})$ would imply that the full history dependence of the overall optimal rule was negligible for large n, meaning that when deciding about some X_j we should essentially focus only on the current observation.

In this paper we extend the approach introduced in [11], [12], [13], and [14] by establishing that the stopping problem in $\mathcal{F}_n$ has a limiting '$n = \infty$' form based on the planar Poisson process. The interpretation of limit risks in terms of the infinite model makes obvious the inequality $V_\infty(\mathcal{F}) < V_\infty(\mathcal{M})$ for any q, provided the values are finite, which is true for both the best-choice problem and Robbins' problem. Thus, the complexity does not disappear in the limit, and the full history dependence persists. Finiteness of the limit risks is guaranteed if $q(r)$ does not grow too fast, e.g. $q(r) < c \exp(r^\beta)$, with $0 < \beta < 1$ being sufficient.

In connection with Robbins' problem, the limiting form was reported by the author at the INFORMS Conference on Applied Probability (Atlanta, 14–16 June 1995), although the Poisson embedding had been exploited earlier in the analysis of rank rules [6]. See [16] for a similar development in the problem of minimising $E[X_\tau]$.

2. A model based on the planar Poisson process

Throughout the paper, we shall use the notation $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$ and $\overline{\mathbb{R}}_+ = [0, \infty]$ for the compactified halfline.

Let $\mathcal{P}$ be the scatter of atoms of a homogeneous Poisson point process in the strip $[0, 1] \times \overline{\mathbb{R}}_+$, with the intensity measure being the Lebesgue measure $dt \, dx$. The infinite collection of atoms can be labelled $(T_1, X_{1,1}), (T_2, X_{1,2}), \ldots$ in the order of increase of the second component. Thus, $\boldsymbol{X}_1 := (X_{1,1}, X_{1,2}, \ldots)$ is the increasing sequence of points of a unit Poisson process on $\mathbb{R}_+$, the T_rs are i.i.d. uniform $[0, 1]$, and $\boldsymbol{X}_1$ and $(T_r, r = 1, 2, \ldots)$ are independent. An atom $(T_r, X_{1,r}) \in \mathcal{P}$ will be understood to be an observation with value $X_{1,r}$, arrival time T_r, and final rank r. We define the initial rank of $(T_r, X_{1,r})$ as one plus the number of atoms in the open rectangle $(0, T_r) \times (0, X_{1,r})$. Note that the coordinatewise ties among the atoms have only probability 0.

It is convenient to treat $\boldsymbol{X}_1$ as a random element of the space $\mathcal{X}$ of all nondecreasing, nonnegative sequences $\boldsymbol{x} = (x_1, x_2, \ldots)$, where $x_r \in \overline{\mathbb{R}}_+$ and $x_r \to \infty$ (as $r \to \infty$). Our convention is that a sequence with finitely many proper terms is always padded by infinitely many terms ∞. In particular, the sequence $\varnothing := (\infty, \infty, \ldots)$ is the sequence with no finite terms. The space $\mathcal{X}$ is endowed with the product topology of componentwise convergence, as inherited from $\overline{\mathbb{R}}_+^\infty$. We denote by $\boldsymbol{x} \cup x$ the nondecreasing sequence obtained by inserting $x \in \overline{\mathbb{R}}_+$ in $\boldsymbol{x}$, with the understanding that $\boldsymbol{x} \cup \infty = \boldsymbol{x}$. A strict partial order on $\mathcal{X}$ is defined by setting $\boldsymbol{x} \prec \boldsymbol{y}$ if $x_r \leq y_r$ for $r = 1, 2, \ldots$, with at least one of the inequalities strict. Clearly, $\boldsymbol{x} \cup x \prec \boldsymbol{x}$ for $x < \infty$, while $\boldsymbol{x} \cup \infty = \boldsymbol{x}$. The order $\prec$ has minimal and maximal elements $\boldsymbol{0} = (0, 0, \ldots)$ and $\varnothing$, respectively.

Let $\boldsymbol{X}_t$ be a random element of $\mathcal{X}$ derived from $\boldsymbol{X}_1$ by removing each entry $X_{1,r}$ if and only if $T_r > t$. Clearly, $\boldsymbol{X}_t$ is an increasing sequence of atoms of a Poisson process on $\mathbb{R}_+$ with mean measure $t \, dx$. The laws of the $\boldsymbol{X}_t$s are mutually singular because, by the law of large numbers,

$$X_{t,r} \sim \frac{r}{t} \quad \text{as } r \to \infty, \tag{3}$$

with probability 1. We regard $\boldsymbol{X}_1$ as the terminal state of the $\mathcal{X}$-valued process $(\boldsymbol{X}_t, t \in [0, 1])$, where the $\boldsymbol{X}_t$s are intensifying as t increases, that is, $\boldsymbol{X}_t \prec \boldsymbol{X}_s$ for $t > s$. For $t \in \{T_r\}$, let X_t, R_t, and I_t denote the value, the final rank, and the initial rank of the observation arrived at time t, respectively, and, for $t \notin \{T_r\}$, let $X_t = R_t = I_t = \infty$. We have $\boldsymbol{X}_t = \boldsymbol{X}_{t-} \cup X_t$, so $\boldsymbol{X}_t = \boldsymbol{X}_{t-}$ unless $t \in \{T_r\}$.

The process $(\boldsymbol{X}_t, t \in [0, 1])$ is Markovian, with right-continuous paths, initial state $\boldsymbol{X}_0 = \varnothing$, and jump times $\{T_r\}$ which comprise a dense subset of $[0, 1]$. Each component $(X_{t,i}, t \in [0, 1])$

is a nondecreasing jump process which satisfies $X_{0+,i} = \infty$ and changes its value at every *i-record* (i.e. the observation of initial rank i). The jump times of $(X_{t,i},\ t \in [0, 1])$ are the arrival times of i-records; these occur according to a Poisson process of intensity $t^{-1}\,\mathrm{d}t$ independently for distinct $i \in \mathbb{N}$, as is known from the extreme-value theory.

Define a *stopping rule* τ to be a variable which may assume only one of the random values $\{T_r\} \cup \{1\}$, and satisfies the measurability condition $\{\tau \le t\} \in \sigma(X_s,\ s \le t)$ for $t \in [0, 1]$. The condition states that the decision to stop not later than t is determined by the atoms of $\mathcal{P} \cap ([0, t] \times \mathbb{R}_+)$ that arrived within the time interval $[0, t]$. Such rules are called 'canonical stopping times' in [16, Definition 2.1].

We fix a nondecreasing, nonnegative loss function q satisfying $q(1) < q(\infty)$. The risk incurred by stopping rule τ is assumed to be

$$\mathrm{E}[q(R_\tau)] = \sum_{r=1}^{\infty} q(r)\,\mathrm{P}(\tau = T_r) + q(\infty)\,\mathrm{P}(\tau = 1), \tag{4}$$

where the terminal component is nonzero if and only if $\mathrm{P}(\tau = 1) > 0$. Let $\mathcal{F}$ be the set of all stopping rules, and let $V(\mathcal{F}) = \inf_{\tau \in \mathcal{F}} \mathrm{E}[q(R_\tau)]$ be the minimal risk.

The optimal stopping problem with criterion (4) would be exactly the same if we had chosen to work with the homogeneous planar Poisson process in a strip $[0, \gamma] \times \overline{\mathbb{R}}_+$, or in any other set of the form $A \times B$, with $A \subset \mathbb{R}$ being a set of finite Lebesgue measure and $B \subset \mathbb{R}$ being a set of infinite Lebesgue measure, but such that the measure of $B \cap [-\infty, 0]$ is finite. To see this, observe that it is possible to map $A \times B$ on $[0, 1] \times \overline{\mathbb{R}}_+$ in such a way that preserves both the coordinatewise order and the Lebesgue measure. In particular, for $[0, \gamma] \times \overline{\mathbb{R}}_+$, the mapping is just the linear transform

$$\begin{pmatrix} \dfrac{1}{\gamma} & 0 \\[2mm] 0 & \gamma \end{pmatrix}.$$

The invariance with respect to the choice of horizon γ is a feature which has no analogue in the discrete-time i.i.d. setting, where we face distinct problems for various values of the horizon n.

We shall exploit a related scaling property of $\mathcal{P}$, called here *self-similarity*, meaning that, for every fixed $t \in [0, 1]$, the affine mapping $(s, x) \mapsto ((s - t)/\bar{t}, x\bar{t})$ preserves both the coordinatewise order and the planar Lebesgue measure and, hence, transforms the restricted point process $\mathcal{P} \cap ([t, 1] \times \overline{\mathbb{R}}_+)$ into a distributional copy of $\mathcal{P}$ with the same ordering of atoms.

The class $\mathcal{R}$ of rank rules is defined by a more restrictive measurability condition $\{\tau \le t\} \in \sigma(I_s,\ s \le t)$ for $t \in [0, 1]$. That is to say, by a rank rule, the information of the observer at time t amounts to the collection of arrival times on $[0, t]$ of i-records for all $i \in \mathbb{N}$. The optimal stopping problem in $\mathcal{R}$ is equivalent to 'the infinite secretary problem' of [10]. By Theorem 4.1 of [10] there exists an optimal rank rule of the form $\tau = \inf\{t : I_t \le \iota(t)\}$, $\inf \varnothing = 1$, where $\iota : [0, 1) \to \mathbb{N} \cup \{0\}$ is a nondecreasing function. For instance, in the best-choice problem $\iota(t) = \mathbf{1}(t \ge \mathrm{e}^{-1})$, while in the minimum expected rank problem

$$\iota(t) = \sum_{i=1}^{\infty} \mathbf{1}(t \ge t_i),$$

where $t_i = \prod_{j=i}^{\infty}(1 + 2/j)^{-1/(j+1)}$ [18].

A memoryless rule is a stopping rule of the form

$$\tau = \inf\{t : X_t \le f(t)\}, \quad \inf \varnothing = 1, \tag{5}$$

where $f : [0, 1) \rightarrow \mathbb{R}$ is a nondecreasing function. Denote by $\mathcal{M}$ the class of memoryless rules and denote by $V(\mathcal{M}) = \inf_{\tau \in \mathcal{M}} \mathrm{E}[q(R_\tau)]$ its stopping value. We could consider a larger class of stopping rules by which the decision to stop depends only on the current observation. However, Lemma 1, below, analogous to [1, Lemma 2.1], shows that such an extension of $\mathcal{M}$ does not reduce the risk.

Lemma 1. *Let $A \subset [0, 1] \times \mathbb{R}_+$ be a Borel set. For the stopping rule $\tau = \inf\{t : (t, X_t) \in A\}$ there exists a memoryless rule whose expected loss is not larger than that of τ.*

Proof. It is sufficient to consider sets A such that the area of $A \cap ([0, t] \times \mathbb{R}_+)$ is finite for every $t < 1$. Indeed, if the area of $A \cap ([0, t] \times \mathbb{R}_+)$ is infinite for some $s < 1$ then $\tau < s$ almost surely (a.s.); hence, letting A' denote $A \cap ([0, s] \times \mathbb{R}_+)$ shifted by $1 - s$ to the right, we obtain a rule not worse than τ. Replace each vertical section of A by an interval adjacent to 0 of the same length; thus obtaining the subgraph of a function g. This preserves the distribution of the stopping rule and does not increase the risk, by the monotonicity of q. Break $[0, 1]$ into intervals of equal size δ and approximate g (in L^1) by a right-continuous function g_δ, which is constant on these intervals. Suppose that on some adjacent intervals $[t, t+\delta)$ and $[t+\delta, t+2\delta)$ we have $g_\delta(t) > g_\delta(t+\delta)$. Let g'_δ be another piecewise-constant function with values on these intervals of $g_\delta(t+\delta)$ and $g_\delta(t)$, but, outside $[t, t+2\delta]$, g'_δ coincides with g. Let $\mathcal{P}'$ be the scatter of atoms obtained by exchanging the strips $[t, t+\delta) \times \mathbb{R}_+$ and $[t+\delta, t+2\delta) \times \mathbb{R}_+$. Obviously, $\mathcal{P}' \stackrel{\mathrm{D}}{=} \mathcal{P}$, where '$\stackrel{\mathrm{D}}{=}$' denotes equality in distribution. To compare the two stopping rules τ and τ', defined as in (5), but with f replaced by g_δ and g'_δ, respectively, we consider the selected atom (τ, X_τ) as a function of $\mathcal{P}$ and the selected atom $(\tau', X_{\tau'})$ as a function of $\mathcal{P}'$. It is easy to see that $X_\tau = X_{\tau'}$ unless $([t+\delta, t+2\delta) \times [0, g(t+\delta)]) \cap \mathcal{P} \neq \varnothing$, whereas in the latter case $X_{\tau'}$ is stochastically smaller than X_τ. The advantage comes from the event that each of the strips contains an atom below the graph of g_δ. It follows that τ' performs better than τ. Iterating this exchange argument, we see that the rule defined by g_δ is improved by a memoryless rule with a piecewise-constant function. Letting δ tend to 0 shows that we can reduce A to a subgraph of a monotonic $f : [0, 1) \rightarrow \overline{\mathbb{R}}_+$.

Given the initial rank $I_t = i$ and the value $X_t = x$ of some observation at time t, the final rank of the atom (t, x) is i plus the number of atoms southeast of (t, x), the latter being a Poisson variable with parameter $\bar{t}x$, where and henceforth

$$\bar{t} := 1 - t.$$

By independence properties of $\mathcal{P}$, the *adapted loss* incurred by stopping at (t, x) is equal to $Q(\bar{t}x, i)$, where

$$Q(\xi, i) := \sum_{r=i}^{\infty} q(r) e^{-\xi} \frac{\xi^{r-i}}{(r-i)!}. \tag{6}$$

For instance, $Q(\bar{t}x, i) = 1 - e^{-\bar{t}x} \mathbf{1}(i = 1)$ in the best-choice problem, and $Q(\bar{t}x, i) = \bar{t}x + i$ in Robbins' problem. The formula for Q is extended for infinite values of the arguments as $Q(\cdot, \infty) = Q(\infty, \cdot) = q(\infty)$. It is seen, from the identity

$$\frac{\mathrm{d}[e^\xi Q(\xi, 1)]}{\mathrm{d}\xi^{i-1}} = e^\xi Q(\xi, i),$$

that the series $Q(\cdot, i)$ has the same convergence radius for all i.

Lemma 2. *If q is not truncated then $Q(\xi, i)$ is strictly increasing in both ξ and i. If q is truncated at m and $q(m-1) < q(m)$ then the strict monotonicity is valid only for $i \in [m]$, while $Q(\xi, i) = q(m) = q(\infty)$ for $i \geq m$.*

Proof. This follows from the monotonicity of q and (6).

In terms of the adapted loss, the risk functional becomes

$$\mathrm{E}[q(R_\tau)] = \mathrm{E}[Q((1-\tau)X_\tau, I_\tau)],$$

with the convention that the terminal component of (4) is included as

$$Q(0 \cdot \infty, \infty) = q(\infty)\, \mathrm{P}(\tau = 1).$$

3. Memoryless rules and finiteness of the risk

For τ, a memoryless rule (5) with monotone f, denote by $L(f) = \mathrm{E}[q(R_\tau)]$ the expected loss. Introduce the integrals

$$F(t) = \int_0^t f(s)\, \mathrm{d}s \quad \text{and} \quad S(x) = \int_0^x f^{-1}(y)\, \mathrm{d}y = x f^{-1}(x) - F(f^{-1}(x)),$$

where f^{-1} is the right-continuous inverse defined for $x < f(0)$ as $f^{-1}(x) = 0$. Note that $\mathrm{P}(\tau > t) = \mathrm{e}^{-F(t)}$ and that, given $\tau = t < 1$, the law of X_τ is uniform on $[0, f(t)]$. The formula for the risk follows by conditioning on the location of the leftmost atom below the graph of f and using the fact that the configurations of the atoms above and below the graph are independent:

$$L(f) = \int_0^1 \mathrm{e}^{-F(t)}\, \mathrm{d}t \int_0^{f(t)} Q(\bar{t}x + S(x), 1)\, \mathrm{d}x + \mathrm{e}^{-F(1)} q(\infty). \tag{7}$$

Computation of the first variation of $L(f)$ shows that an optimal f must satisfy the following rather complicated functional equation:

$$
\begin{aligned}
&Q(f(t) - F(t), 1) \\
&\quad = \int_t^1 \exp(F(t) - F(s))\, \mathrm{d}s \left(\int_0^{f(t)} Q(S(x) + x\bar{s}, 1)\, \mathrm{d}x + \int_{f(t)}^{f(s)} Q(S(x) + x\bar{s}, 2)\, \mathrm{d}x \right),
\end{aligned}
\tag{8}
$$

which should be complemented by a suitable boundary condition at $t = 1$. We did not succeed in solving (8), even in the simplest case of the best-choice problem; hence, we turn to simpler bounds.

A rough upper bound

$$L(f) \leq \int_0^1 \mathrm{e}^{-F(t)}\, \mathrm{d}t \int_0^{f(t)} Q(x, 1)\, \mathrm{d}x + \mathrm{e}^{-F(1)} q(\infty), \tag{9}$$

follows from $\bar{t}x + S(x) \leq x$. Minimising the right-hand side of (9) amounts to solving an optimal stopping problem $\mathrm{E}[Q(X_\tau, 1)] \to \inf$ of the familiar type [16], with loss dependent on the value of the stopped observation only. For $\ell = 1$, we face a variant of 'Moser's problem' associated with $\mathcal{P}$; see [1], [3], [16], and the references therein.

The auxiliary problem is easy for the falling factorial loss functions

$$q(r) = (r - 1)(r - 2) \cdots (r - \ell), \qquad \ell \in \mathbb{N}, \tag{10}$$

in which case we have a very simple formula $Q(\xi, 1) = \xi^\ell$; hence, (9) becomes

$$L(f) \leq (\ell + 1)^{-1} \int_0^1 e^{-F(t)} f(t)^{\ell+1} \, dt.$$

Substituting $f(t) := (\ell + 1)/(1 - t)$, the right-hand side is evaluated as $(\ell + 1)^\ell$. We can arrive at the function $f_{\ell+1}$ by solving the variational problem with Lagrangian $e^{-F}(F')^{\ell+1}$ and boundary conditions $F(0) = 0$ and $F(1) = \infty$.

Notably, for the stopping problem $E[(X_\tau)^\ell] \to \inf$ (with any $\ell > 0$), the optimality of some memoryless rule with hyperbolic threshold

$$f_b(t) := \frac{b}{1 - t}, \qquad b > 0,$$

can be concluded solely from the self-similarity of $\mathcal{P}$ and the homogeneity of the power function x^ℓ. For such rules, (9) becomes

$$L(f_b) \leq \int_0^1 (1 - t)^b \int_0^{b/(1-t)} x^\ell \, dx = \frac{b^{\ell+1}}{(\ell + 1)(b - \ell)},$$

which indeed attains a minimum at $b = \ell + 1$.

For small ℓ, it is not hard to compute $L(f_b)$ for the memoryless rules with threshold f_b and q as given in (10). We calculate $e^{-F(t)} = (1 - t)^b$ and $S(x) = (x - b - b \log(x/b))$ (for $x > f(0) = b$). For $\ell = 1$, integrating (7) by parts we obtain

$$L(f_b) = \frac{b}{2} + \frac{1}{b^2 - 1}, \tag{11}$$

which is finite for all $b > 1$, with the minimum $1.3318 \cdots$ attained at $b = 1.9469 \cdots$ (which agrees with Example 4.2 of [1], where the minimum is $2.3318 \cdots$ for the linear loss $q(r) = r$). For $\ell = 2$,

$$L(f_b) = \frac{b^3}{3} + \frac{2(b^4 - 2b^3 + 2b^2 + 6b - 4)}{(b - 2)(b - 1)^2(b + 1)(b + 2)}, \tag{12}$$

which is finite for all $b > 2$, with minimum $4.4716 \cdots$ at $b = 2.96439 \cdots$. The formulae become more involved for larger ℓ, a common feature being that $L(f_b) < \infty$ for $b > \ell$. For $\ell = 3$, the minimum is 24.8061 at $b = 3.9734 \cdots$. For $\ell = 4$, the minimum is $194.756 \cdots$ at $b = 4.979 \cdots$.

We see that (see also [1] and [3]), for the linear loss $q(r) = r$, we have $V(\mathcal{M}) = \inf L(f) < V(\mathcal{R}) = 3.8695 \cdots$, where the value of the rank rules is taken from [7] (it is the $\ell = 1$ instance of (14), below). The minimiser of $L(f)$ (that is, the optimal $\mathcal{M}$-rule) is not known explicitly, but some approximations to it can be found in [1], where they appear in the course of the asymptotic analysis of the finite-n problem.

In the best-choice problem the $\mathcal{F}$-optimal rule is

$$\tau^* = \min\{t : I_t = 1, \ X_t \leq f_b\}, \qquad b = 0.804 \cdots, \tag{13}$$

and $V(\mathcal{F}) = 1 - 0.580 \cdots$ (see, e.g. [11, p. 682]). Note that τ^* is not memoryless since only stopping on 1-records is allowed! However, the most simple memoryless rule with constant threshold $f(t) = 1.503 \cdots$ achieves $L(f) = 1 - 0.517 \cdots$, and hence beats the rank rules $V(\mathcal{M}) < V(\mathcal{R}) = 1 - 0.368 \cdots$. Evaluation of (7), which was carried out computationally, shows that the best values of $L(f)$ that we can achieve with some parametric families of fs are

(i) $1 - 0.547 \cdots$ for $f(t) = a\,\mathbf{1}(t \le s) + b\,\mathbf{1}(t > s)$, where $0 < a < b$ and $0 < s < 1$;

(ii) $1 - 0.537 \cdots$ for $f(t) = b/(1 - t)$, where $b > 0$;

(iii) $1 - 0.556 \cdots$ for $f(t) = at + b$, where $a, b > 0$.

The best-choice and minimum-rank instances suggest that it is more advantageous to learn the value of the observation rather than its initial rank. It would be interesting to explore whether this was true in general, and if not then it would be interesting to determine for which values of q the memoryless rules outperform the rank rules.

For unbounded q, we also face the problem of distinguishing between the situations with finite or infinite risk. We sketch some known and some new results in this direction. From the above elementary estimates, $V(\mathcal{M}) < \infty$ provided that $q(r) < cr^\ell$ for some constants $c > 0$ and $\ell > 0$. For such q, the risk of the rank rules is also finite. Moreover, Mucci [18, p. 426] showed that, for the rising factorial loss function $q(r) = r(r + 1) \cdots (r + \ell - 1)$, $\ell \in \mathbb{N}$, the minimum risk of the rank rules is

$$V(\mathcal{R}) = \ell! \prod_{j=1}^{\infty} \left(1 + \frac{\ell + 1}{j}\right)^{\ell/(\ell+j)} \tag{14}$$

(which extends the $\ell = 1$ result from [7]). For $\ell = 2$, the formula yields $33.260 \cdots$ (while the f_b-rules do worse, with $\inf_b L(f_b) = 38.068 \cdots$, as computed from (11) and (12) using the linearity of $L(f)$ in q). In fact, $V(\mathcal{M}) < \infty$ for many loss functions growing much faster than polynomials.

Proposition 1. *If $q(r) < c\exp(x^\beta)$ for some $c > 0$ and $0 < \beta < 1$ then $V(\mathcal{M}) < \infty$.*

Proof. The risk is finite for the memoryless rule with $f(t) = (1-t)^{-\alpha}$ for any $\alpha > (1-\beta)^{-1}$. To see this, use the bound in (9) and the formulae

$$Q(x, 1) = O(\exp(2x^\beta)), \qquad x \to \infty,$$
$$e^{-F(t)} = \exp\left(-\frac{1}{(\alpha - 1)(1 - t)^{\alpha-1}}\right),$$

which also imply that, for this rule, $P(\tau = 1) = 0$. Then, $E[\exp((X_\tau)^\beta)]$ is estimated from asymptotics of the incomplete gamma function.

It is not known whether the risk of the rank rules is finite for $q(r) = \exp(x^\beta)$ with $0 < \beta < 1$. We shall now look for the growth conditions on q forcing infinite risks.

Proposition 2. *If $Q(b, 1) = \infty$ for some $b \in \mathbb{R}_+$ then $V(\mathcal{F}) = \infty$, i.e. there is no stopping rule $\tau \in \mathcal{F}$ with finite risk.*

Proof. For such b, the conditional loss by stopping above f_b is infinite; thus, we can only consider stopping rules τ which never stop above f_b and which satisfy $P(\tau = 1) = 0$. Choose any x with

$$S(x) = x - b - b \log\left(\frac{x}{b}\right) > b.$$

On the nonzero event

$$A = \{\mathcal{P} \cap \{(t, y): y < \min(x, f(t))\} = \varnothing\},$$

stopping occurs at some atom (s, z) with $s > 1 - b/x$, $z > x$, and by averaging we see that the expected loss is infinite.

In the sequel we shall assume that the convergence radius of $Q(\cdot, 1)$ is infinite (hence, $Q(x, i) < \infty$ for all finite x, i). Without this assumption, the stopping problem is trivial.

See Proposition 5.3 of [10] for an analogue of Proposition 2 for the rank rules. By Section 5 of [10] we have $V(\mathcal{R}) = \infty$ if $\sum_r (\log q(r))/r^2 = \infty$; hence, the risk of the rank rules is infinite for the exponential loss function $q(r) = e^{cr}$, $c > 0$. A sufficient condition for $V(\mathcal{M}) = \infty$ follows by comparison with an auxiliary problem.

Proposition 3. *Given* $0 < c < 1$, *suppose that* $E[Q(cX_\tau, 1)] = \infty$ *for every stopping rule* $\tau < 1$. *Then* $V(\mathcal{M}) = \infty$ *also.*

Proof. For any increasing f, we have $S(x)/x \to 1$ as $x \to \infty$; thus, we can choose x_0 large enough to satisfy $S(x) > cx$ for all $x > x_0$. For A as given in the proof of Proposition 2, we have

$$
\begin{aligned}
L(f) &= E[Q(S(X_\tau) + (1 - \tau)X_\tau, 1)] \\
&\geq P(A)\, E[Q(cX_\tau, 1) \mid A] \\
&\geq P(A)\, E[Q(cX_\tau, 1) \mid \tau > f^{-1}(x_0)],
\end{aligned}
$$

which is finite exactly when $E[Q(cX_\tau, 1)] < \infty$, as follows from the independence properties of $\mathcal{P}$ and the definition of the memoryless rule.

As a corollary, we obtain $V(\mathcal{M}) = \infty$ for $q(r) = e^{cr}$, $c > 0$. Indeed, for this loss function, $Q(x, 1) > c_1 e^{cx}$; hence,

$$E[Q(2^{-1}X_\tau, 1)] \geq c_1 \int_0^1 e^{-F(t)}\, dt \int_0^{F'(t)} e^{cx/2}\, dx.$$

To show that the right-hand side is infinite for every F with $f = F' > 0$ and $F(0) = 0$, $F(1) = \infty$, it is enough to check that, for every $c_2 > 0$,

$$\int_0^1 \exp(F(t) - c_2 F'(t))\, dt = \infty,$$

which is achieved by solving the variational problem with fixed boundary value $F(1)$ and then letting $F(1)$ tend to ∞.

In view of (3) we can expect that $V(\mathcal{F}) = \infty$ exactly when the risk in the problem with loss $Q(X_\tau, 1)$ is infinite. This seems difficult to prove, even in the exponential case $q(r) = e^{cr}$. In contrast, however, the problem of minimising $Q((1 - \tau)X_\tau, 1)$ is always trivial, i.e. the rules with threshold f_b approach the infimum $q(1)$ as b tends to 0.

4. Properties of the optimal rule

The optimal stopping problem in $\mathcal{F}$ is of Markovian type, associated with the time-homogeneous Markov process $((X_t, I_t), t \in [0, 1])$, with state-space $\mathcal{X} \times \overline{\mathbb{N}}$ and time-dependent loss $Q(\bar{t}X_t, I_t)$ for stopping at time t. If I_t assumes some finite value i then $t \in \{T_r\}$ and $X_{t,i} = X_t$, which combined with the fact that ranking of the arrivals after t depends on $\mathcal{P} \cap ([0, t] \times \mathbb{R}_+)$ through X_t shows that (X_t, I_t) indeed summarises all the relevant information up to time t. By virtue of $X_{I_t,t} = X_t$ (where $X_{\infty,t} := \infty$), the data (X_t, I_t) contains the same information as (X_{t-}, X_t).

Following a well-known approach, we will consider a family of conditional stopping problems parametrised by (t, x). This corresponds to the class of stopping rules $\tau > t$ measurable with respect to $\mathcal{P} \cap ([t, 1] \times \mathbb{R}_+)$ that operate under the condition $X_t = x$. The effect of conditioning is that each $x_r < X_\tau$ contributes one unit to R_τ in the event $\tau < 1$. The dependence of conditional risk on t can be eliminated using the self-similarity of $\mathcal{P}$.

Lemma 3. *The optimal stopping problem from time t with history x is equivalent to the optimal stopping problem starting with $X_0 = \bar{t}x$ at time 0.*

We denote by E_x the expected value given $X_0 = x$, and we denote by

$$v(x) = \inf_{\tau \in \mathcal{F}} \mathrm{E}_x[q(R_\tau)]$$

the minimum risk given $X_0 = x$. In the setting with varying initial configuration x, all stopping rules and variables X_t, I_t, and R_t are considered to be functions of both x and $\mathcal{P}$. In the search for the optimality within $\mathcal{F}$ it is sufficient to consider stopping rules τ with the following invariance property. Given $\tau \geq t$ and that at time t the observed value is $X_t = x$ and the history is $X_{t-} = x$, the decision to stop at (t, x) depends on these data through $(\bar{t}x, \bar{t}x)$. Note that a memoryless rule with constant threshold does not fulfill this invariance condition, while a rule with threshold f_b does.

The function v, defined on the whole of $\mathcal{X}$, satisfies a lower bound

$$v(x) \geq \sum_{r=1}^{\infty} q(r)(\exp(-x_{r-1}) - \exp(-x_r)), \qquad x_0 = 0 \tag{15}$$

(strict if the series converges), which follows since $X_{1,1}$ is exponentially distributed and $X_\tau \geq X_{1,1}$. The bound is a continuous-time analogue of the finite-n 'half-prophet' bounds in Lemma 3.2 of [4]. From (15), $v(x) \to q(\infty)$ as $x \to \mathbf{0}$.

If q is truncated at m then clearly v depends only on the first $m-1$ components of x and satisfies $v(x) < q(m)$. Obviously, from Lemma 2, if q is not truncated then $x \prec y$ implies that $v(x) < v(y)$, provided that these values are finite; if q is truncated at m with $q(m-1) < q(m)$ then the analogous implication holds for the strict order defined on the first $m-1$ components.

The monotonicity of v allows us to further restrict our consideration to the stopping rules that satisfy

$$\mathrm{E}_x[q(R_\tau)] \geq \mathrm{E}_y[q(R_\tau)] \quad \text{for } x \prec y,$$

and, by Lemma 4, below, we can further restrict our consideration to the stopping rules that satisfy ess sup $\tau = 1$.

Lemma 4. *Every stopping rule τ with $\tau < 1 - \varepsilon$ a.s. can be improved by a stopping rule τ' with ess sup $\tau' = 1$.*

Proof. Note that τ depends on the observations that arrive within the time interval $[0, 1-\varepsilon]$. Since $\mathcal{P} \cap ([\varepsilon, 1] \times \mathbb{R})$ is a copy of $\mathcal{P} \cap ([0, 1-\varepsilon] \times \mathbb{R})$, defining τ'' to be a copy of τ acting on $\mathcal{P} \cap ([\varepsilon, 1] \times \mathbb{R})$ we have $\mathrm{E}[q(R_\tau)] = \mathrm{E}[q(R_{\tau''})]$. However, τ'' can be improved by the rule which picks the first arrival (if any) on $[0, \varepsilon]$ below some sufficiently small x, and coincides with τ'' otherwise.

We will call the stopping rule τ *admissible* if it is invariant and satisfies the above two conditions.

Lemma 5. *If* $\mathrm{E}_{\boldsymbol{x}}[q(R_\tau)] < \infty$ *for an admissible stopping rule* τ *then also* $\mathrm{E}_{\boldsymbol{x}'}[q(R_\tau)] < \infty$ *for* $\boldsymbol{x}' = \boldsymbol{x} \cup x$, $x > 0$.

Proof. On the positive event $\tau > 1 - \varepsilon$, the loss of τ is

$$\mathrm{P}(\tau > 1 - \varepsilon) \, \mathrm{E}_{\boldsymbol{y}}[q(R_\tau)] < \infty,$$

where $\boldsymbol{y} \prec \varepsilon \boldsymbol{x}$. Thus, $\mathrm{E}_{\varepsilon \boldsymbol{x}}[q(R_\tau)] < \infty$. Conversely, $\varepsilon \boldsymbol{x} \prec \boldsymbol{x}'$ for small enough ε, so $\mathrm{E}_{\boldsymbol{x}'}[q(R_\tau)] \leq \mathrm{E}_{\varepsilon \boldsymbol{x}}[q(R_\tau)]$ is finite.

With respect to the product topology in $\mathcal{X}$, the continuity of v is a nontrivial issue. If $V(\mathcal{F}) = \infty$ then, of course, $v(\boldsymbol{x}) = \infty$ everywhere, but, for arbitrary unbounded q, there exists a dense in $\mathcal{X}$ set of sequences $\boldsymbol{x} = (x_r)$ for which $x_r \uparrow \infty$ so slowly that $v(\boldsymbol{x}) = \infty$. Thus, if $q(\infty) = \infty$, the function v is discontinuous at every point where it is finite.

Proposition 4. *If* $v(\boldsymbol{x}) < \infty$ *then* $v(\boldsymbol{x} \cup x)$ *is finite and continuous in* x.

Proof. From Lemma 5 we see that $v(\boldsymbol{x} \cup x) < \infty$. Fix $\boldsymbol{x}$ and x, and suppose that a rule τ is ε-optimal under $X_0 = \boldsymbol{x} \cup x$, and is such that X_τ is the same for all $\boldsymbol{x} \cup x'$ as x' varies. Then, for all x' sufficiently close to x, the stopped rank R_τ assumes the same value. It follows, by dominated convergence, that $\mathrm{E}_{\boldsymbol{x} \cup x'}[q(R_\tau)] \to \mathrm{E}_{\boldsymbol{x} \cup x}[q(R_\tau)]$ as $x' \to x$. Since ε is arbitrary and the rule τ is suboptimal under $\boldsymbol{x} \cup x'$, we conclude that $v(\boldsymbol{x} \cup x) + \varepsilon \geq v(\boldsymbol{x} \cup x')$. The analogous inequality with x and x' swapped is derived by considering a rule τ' ε-optimal under $\boldsymbol{x} \cup x'$, so the continuity follows.

Let $i(\boldsymbol{x}, x) := \#\{r : x_r \leq x\}$, and suppose that $\boldsymbol{x}$ has $x_1 > 0$. Applying Lemma 2, we see that if q is not truncated then the function $Q(x, i(\boldsymbol{x}, x))$ is strictly increasing in x from $q(1)$ to $q(\infty)$. If q is truncated at m and $q(m-1) < q(m)$ then $Q(x, i(\boldsymbol{x}, x))$ is strictly increasing as x varies from 0 to x_{m-1}, with $Q(x, i(\boldsymbol{x}, x)) = q(m)$ for $x \geq x_{m-1}$. Conversely, $(\boldsymbol{x} \cup x) \prec (\boldsymbol{x} \cup y)$ for $x < y$; hence, $v(\boldsymbol{x} \cup x)$ is nonincreasing in x. Thus, introducing

$$h(\boldsymbol{x}) := \sup\{x : Q(x, i(\boldsymbol{x}, x)) < v(\boldsymbol{x} \cup x)\},$$

we have $Q(x, i(\boldsymbol{x}, x)) < v(\boldsymbol{x} \cup x)$ for $x < h(\boldsymbol{x})$ and $Q(x, i(\boldsymbol{x}, x)) \geq v(\boldsymbol{x} \cup x)$ for $x \geq h(\boldsymbol{x})$. Subject to obvious adjustments, the definition of $h(\boldsymbol{x})$ makes sense for every $\boldsymbol{x} \neq \boldsymbol{0}$ in the untruncated case and for $x_{m-1} > 0$ in the truncated case.

We are ready to show that memoryless rules are not optimal.

Proposition 5. *If* $V(\mathcal{F}) < \infty$ *then* $V(\mathcal{F}) < V(\mathcal{M})$.

Proof. For a memoryless rule with threshold function f to be optimal, we must have $v(\bar{t}X_t) < Q(\bar{t}X_t, i(X_{t-}, X_t))$ for $X_t > f(t)$ and $v(\bar{t}X_t) > Q(\bar{t}X_t, i(X_{t-}, X_t))$ for $X_t < f(t)$, because otherwise the rule can be improved. This forces $f(t) = h(\bar{t}\boldsymbol{x})$, which cannot hold since h is not constant.

To concretely demonstrate how a memoryless rule with threshold f can be improved, let us apply the same approach as used in Section 5 of [4]. Assume that $q(\infty) = \infty$. Suppose that (t, x) is above the graph of f and, hence, should be skipped by the memoryless rule, irrespective of the initial rank or history. Let $i(x, x) = 1$ under some history x. Keeping x fixed while letting sufficiently many of the components x_r decrease, but ensuring that they remain larger than x, we can always show that, for the modified history x', the bound in (15) is arbitrarily close to $q(\infty)$ while the expected loss of stopping remains unaltered at $Q(\bar{t}x, 1)$. For such x', we have $v(\bar{t}(x' \cup x)) > Q(\bar{x}, 1)$, from which it is easy to conclude that stopping at (t, x) will strictly reduce the adapted loss and that this is an event of positive probability.

Based on the function $h \colon \mathcal{X} \to \overline{\mathbb{R}}_+$, we construct a predictable process

$$H_t := h(X_{t-} \setminus \{X_{1,r} \colon T_r < t, \ X_{1,r} < h(X_{T_r-})\}), \qquad t \in [0, 1].$$

Let Y_t be a thinned sequence obtained by removing the terms in $\{\ldots\}$ from X_{t-}, so $H_t = h(Y_t)$. Intuitively, H_t is a history-dependent threshold which depends on the configuration of atoms X_{t-} that arrived within the time interval $[0, t)$ and are above the curve $(H_s, \ s \in [0, t))$. As t starts increasing from 0, the process H_t coincides with $h(X_{t-})$ as long as there are no atoms below the threshold, while at the first moment this occurs the atom is discarded and does not affect the future path of the process.

The operation of thinning $\mathcal{P}$ allows us to introduce (H_t) as a random function defined on $[0, 1)$, as opposed to considering $h(X_{t-})$ killed as soon as the threshold is undershot by some X_t. Alternatively, (H_t) could be defined as $h(X_{t-})$ conditioned on $X_t \geq H_t$, $t \in [0, 1)$.

We list some properties of (H_t) which follow directly from the definition and Lemmas 4 and 2 (under $X_0 = \varnothing$).

Lemma 6. (i) *The function (H_t) is nondecreasing on $[0, 1)$.*

(ii) *If $V(\mathcal{F}) < \infty$ then H_0 is the unique root of $Q(x, 1) = v(x \cup \infty)$.*

(iii) *We have $H_{1-} = Y_{1, m-1}$ if q is truncated at m and $q(m-1) < q(m)$.*

(iv) *We have $H_{1-} = \infty$ unless q is truncated.*

To reveal qualitative features of (H_t) and to gain some intuition, we shall gradually increase the complexity of the loss function. In the simplest instance of the best-choice problem, v depends only on x_1 and there is an explicit formula for the threshold:

$$H_t = \min(f_b(t), Y_{t,1}), \qquad b = 0.804 \cdots;$$

cf. (13) and [12, Equations (8) and (13)]. That is to say, as t starts increasing from 0, H_t is a deterministic drift process until it hits the level of the lowest atom above the graph. The drift is hyperbolic due to the self-similarity of $\mathcal{P}$ (Lemma 3). After this random hitting time, H_t has a *flat*, which appears because it is never optimal to stop at the observation with initial rank 2 or larger. On the first part of the path, H_t satisfies $Q(H_t, 1) = v(\bar{t}(Y_t \cup H_t))$, and on the second part of the path, H_t satisfies $Q(H_t, 1) < v(\bar{t}(Y_t \cup H_t))$.

If q is strictly truncated at $m = 3$, meaning that $q(2) < q(3) = q(\infty)$, a new effect appears. For sufficiently small t, as long as $H_t < Y_{t,1}$, each 1-record above the threshold causes a *jump* because $v(\bar{t}Y_t)$ jumps and the threshold must go up to compensate. Thus, (H_t) has both drift and jump components. The jump locations are the 1-record times accumulating near 0 at rate $t^{-1} \, dt$. As H_t hits $Y_{t,1}$, there is a possible flat, then a period of deterministic drift, where $Q(H_t, 2) = v(\bar{t}(Y_t \cup H_t))$, and finally there is a flat at some level $Y_{t,2}$ (then $Y_{t,2} = Y_{1,2}$).

For q strictly truncated at $m > 3$, the jump locations are included in $m - 2$ record-time processes of atoms with initial rank at most $m - 2$, there are $m - 1$ potential flats, and a drift component between the flats. We do not assert that the number of flats is always exactly $m - 1$, because it is not at all clear if (H_t) can break a level $Y_{t,r}$ for $r < m - 1$ by jumping through it; hence sparing a flat.

Now suppose that q is not truncated and that $H_t < \infty$ everywhere on $[0, 1)$ with probability 1. Then, outside the union of flat intervals, *every* arrival above H_t causes a jump; thus, the set of jump locations is dense there. The number of flats may be infinite, and outside the flats $Q(H_t, i(\boldsymbol{Y}_t, H_t)) = v(\bar{t}(\boldsymbol{Y}_t \cup H_t))$. The durations of the flats are random variables, depending on both the history before the flat starts and the observations as the flat unfolds.

In the case of Robbins' problem we have, by linearity of the loss,

$$Q(x, i + 1) - Q(x, i) = 1 \quad \text{and} \quad v(\boldsymbol{x} \cup x) - v(\boldsymbol{x}) < 1$$

(if $v(\boldsymbol{x} \cup x) < \infty$). Thus, $Q(x, i(\boldsymbol{x}, x)) = v(\boldsymbol{x} \cup x)$ implies that

$$Q(x, i(\boldsymbol{x}, x) + 1) > v(\boldsymbol{x} \cup x \cup x')$$

for arbitrary x'. But this means that (H_t) cannot cross any $Y_{t,i}$ by a jump. It follows that (H_t) has infinitely many flats at all levels $Y_{1,r}$, $r \in \mathbb{N}$. The presence of all three effects (drift, jumps, and flats) and the lack of independence of increments property all leave little hope for a kind of analytic description of (H_t).

The optimality principle requires stopping at atom (t, x) when the history $\boldsymbol{X}_{t-} = \boldsymbol{x}$ satisfies $Q(\bar{t}x, i(\boldsymbol{x}, x)) < v(\bar{t}\boldsymbol{x})$, whence we derive the following analogue of (2).

Proposition 6. *If $V(\mathcal{F}) < \infty$ then $H_t < \infty$ for all $t < 1$ a.s. and the stopping rule*

$$\tau^* := \inf\{t : X_t < H_t\}, \qquad \inf \varnothing = 1,$$

is optimal in $\mathcal{F}$. (Note that $H_t(\boldsymbol{X}_{t-}) = H_t(\boldsymbol{Y}_{t-})$ for $t \leq \tau^*$.)

Proof. For bounded q, a general result [21, Theorem 3] is applicable since the function $Q(x, i(\boldsymbol{x}, x))$ is bounded and continuous on $\mathcal{X} \times \overline{\mathbb{N}}$.

Alternatively, for q truncated at some m, we can use results from the optimal stopping theory for discrete-time processes. To fit exactly in this framework, focus on the sequences of i-records (for $i \leq m - 1$) that arrive within the time interval $[\varepsilon, 1]$, and then let $\varepsilon \to 0$. The general bounded case follows in the limit $m \to \infty$.

For unbounded q, we use another kind of truncation (analogous to that used in [3, Section 4]). For m fixed, let

$$Q^{(m)}(x, i) = Q(x, \max(i, m)),$$

and consider the stopping problem with loss $Q^{(m)}(\bar{t}x, i(\boldsymbol{x}, x))$ for stopping at (t, x) with history $\boldsymbol{x}$. This corresponds to ranking x relative to at most m atoms before t, but fully accounting for all future observations below x. In this problem it is never optimal to stop at an atom with relative rank m or higher. Indeed, stopping at (t, x) with such rank can be improved by continuing and then exploiting any hyperbolic memoryless rule with $b < \bar{t}x$ (stopping is guaranteed before 1 since the subgraph of f_b has infinite area). By discrete-time methods, optimality of the rule $\tau^{(m)} = \inf\{t : X_t < H_t^{(m)}\}$ in the truncated problem is readily acquired, with a nondecreasing predictable process $(H_t^{(m)})$ defined through

$$h^{(m)}(\boldsymbol{x}) := \sup\{x : Q^{(m)}(x, i(\boldsymbol{x}, x)) < v^{(m)}(\boldsymbol{x} \cup x)\},$$

where $v^{(m)}$ is the minimum loss analogous to v. Obviously, $Q^{(m)}(x, i(\mathbf{x}, x))$ and $v^{(m)}(\mathbf{x})$ are nondecreasing in m.

A decisive property of this kind of truncation is that $Q^{(m)}(x, i) = Q(x, i)$ for $m \geq i$. This implies that $H_t^{(m)}$ is eventually nondecreasing in m; hence, there exists a pointwise limit $H_t' = \lim_{m \to \infty} H_t^{(m)}$, which defines a legitimate stopping rule τ' as the time of the first arrival under H'. For notational convenience, let

$$L(\tau) = \mathrm{E}[Q(X_\tau, I_\tau)] \quad \text{and} \quad L^{(m)}(\tau) = \mathrm{E}[Q^{(m)}(X_\tau, I_\tau)],$$

and let u and $u^{(m)}$ denote the minimum risks (so $u = V(\mathcal{F})$). Trivially, $\lim_{m \to \infty} u^{(m)} \leq u$. Conversely, by monotone convergence $L^{(m)}(\tau') \uparrow L(\tau) \geq u$. If follows that $u^{(m)} \leq u$ and τ' is optimal. The convergence $v^{(m)}(\mathbf{x}) \uparrow v(\mathbf{x})$ is shown in the same way, from which $H_t' = H_t$ and $\tau' = \tau^*$ is optimal.

As a by-product, we have shown that the risk in the truncated problem with loss function $q(\min(r, m))$ converges to $V(\mathcal{F})$. Indeed, the loss is squeezed between the loss in the modified truncated problem and the original untruncated loss.

From the formula for the distribution of the optimal rule,

$$\mathrm{P}(\tau^* > t) = \mathrm{E}\left[\exp\left(-\int_0^s H_s \, \mathrm{d}s\right)\right],$$

and arguing as in Lemma 1, we see that H_t cannot explode to ∞ at some $t < 1$ if $V(\mathcal{F}) < \infty$.

The risk can be bounded from below in the spirit of (7) as

$$\mathrm{E}[q(R_{\tau^*})] \geq \mathrm{E}\left[\int_0^1 \exp\left(-\int_0^s H_s \, \mathrm{d}s\right) \int_0^{H_t} Q(\bar{t}x, \phi_H(x)) \, \mathrm{d}x\right],$$

where $\phi_H(x)$ is the number of flats of (H_t) below x. If the loss function q has the property that the flats of (H_t) occur at all levels $X_{1,r}$, $r \in \mathbb{N}$ (like in Robbins' problem), the equality holds. The same kind of lower bound is valid for every stopping rule τ defined by means of an arbitrary nondecreasing predictable process like (H_t).

Remark. A small-time decomposition of the risk leads in the m-truncated case to the 'dynamic programming' equation

$$\sum_{i=1}^{m-1} x_i \frac{\partial v(\mathbf{x})}{\partial x_i} = \sum_{i=1}^{m-1} \int_{x_{i-1}}^{x_i} (\min\{Q(\mathbf{x} \cup x, i), v(\mathbf{x} \cup x)\} - v(\mathbf{x})) \, \mathrm{d}x, \qquad x_0 = 0,$$

in the $(m-1)$-dimensional 'Weyl chamber' $\{(x_1, \ldots, x_{m-1}): x_1 \leq \cdots \leq x_{m-1}\}$ with the initial condition $v(\mathbf{0}) = q(m)$. The equation is not completely trivial for the case in which $m = 2$; see [12, Section 2.2], where t stands for our x_1. Trying to justify this equation in the case of unbounded q we need to overcome many conceptual difficulties like infinitely many variables, the domain of definition of v, condition (3), and the fact that the initial condition at $\mathbf{0}$ evaporates. We remind the reader that in the simpler setting of rank rules, the question concerning the uniqueness of the solution to a similar equation (see [10, p. 432]) remains open.

5. The infinite Poisson model as a limit of finite-n problems

To connect the finite-n problem with its Poisson counterpart it is convenient to realise an i.i.d. sequence in the following way [9], [11], [14]. Divide the strip $[0, 1] \times \mathbb{R}_+$ into n vertical strips of the same width $1/n$. Let X_j be the atom of $\mathcal{P}$ with the lowest x-value. By properties of the Poisson process, $X_1, \ldots, X_n$ are i.i.d. with exponential distribution of rate $1/n$. Note that the optimal stopping of $X_1, \ldots, X_n$ is equivalent to the optimal stopping of $\mathcal{P}$ with the lookback option that allows the observer to return to any atom within a given $1/n$-strip (equivalently, at time $(j - 1)/n$ to foresee the configuration of atoms up to time j/n). This embedding in $\mathcal{P}$ immediately implies that $V_n(\mathcal{F}_n) < V(\mathcal{F})$. Moreover, as n tends to ∞, each i-record process derived from $X_1, \ldots, X_n$ converges almost surely to the i-record process derived from $\mathcal{P}$. From this it is easy to conclude, first for truncated then for any bounded q, that $V_\infty(\mathcal{F}) = V(\mathcal{F})$, where $V_\infty(\mathcal{F}) = \lim_{n \to \infty} V_n(\mathcal{F}_n)$ is as defined in the introduction.

For the general q, the relations

$$V_\infty(\mathcal{F}) = V(\mathcal{F}), \qquad V_\infty(\mathcal{R}) = V(\mathcal{R}), \quad \text{and} \quad V_\infty(\mathcal{M}) = V(\mathcal{M})$$

follow (as in [1], [2], [4], [7], [9], and [17]) from that in the truncated case, by combining monotonicity of risks in the truncation parameter m with the monotonicity in n stated in Lemma 7, below.

Lemma 7. $V_n(\mathcal{F}_n)$, $V_n(\mathcal{R}_n)$, and $V_n(\mathcal{M}_n)$ are increasing with n.

Proof. This all is standard; see the references above. We only add small details to Theorem 2.4 of [1] for the $\mathcal{M}$-case. Let τ be an optimal memoryless rule in the problem of size $n + 1$, and let τ' be a modified memoryless strategy which always skips the worst value $X_{n+1,n+1}$, but otherwise has the same thresholds as τ. (To apply τ', the observer must be able to recognise $X_{n+1,n+1}$ as it arrives.) Then τ' strictly improves τ in the event that τ stops at $X_{n+1,n+1}$. Conversely, strategy τ' performs as a mixture of memoryless rules in the problem of size n, because given $X_{n+1,n+1} = x$ the other X_js are i.i.d. uniform on $[0, x]$. Therefore, $V_n(\mathcal{M}_n) < V_{n+1}(\mathcal{M}_{n+1})$.

References

[1] 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.

[2] BEREZOVSKY, B. A. AND GNEDIN, A. V. (1984). *The Best Choice Problem*. Nauka, Moscow.

[3] BRUSS, F. T. (2005). What is known about Robbins' problem? *J. Appl. Prob.* **42,** 108–120.

[4] BRUSS, F. T. AND FERGUSON, T. S. (1993). Minimizing the expected rank with full information. *J. Appl. Prob.* **30,** 616–626.

[5] BRUSS, F. T. AND FERGUSON, T. S. (1996). Half-prophets and Robbins' problem of minimizing the expected rank. In *Athens Conf. Appl. Prob. Time Series Anal.* (Lecture Notes Statist. **114**), Springer, New York, pp. 1–17.

[6] BRUSS, F. T. AND ROGERS, L. C. G. (1991). Embedding optimal selection problems in a Poisson process. *Stoch. Process. Appl.* **38,** 267–278.

[7] CHOW, Y. S., MORIGUTI, S., ROBBINS, H. AND SAMUELS, S. M. (1964). Optimal selection based on relative rank. (The "secretary problem".) *Israel J. Math.* **2,** 81–90.

[8] FRANK, A. AND SAMUELS, S. M. (1980). On an optimal stopping problem of Gusein-Zade. *Stoch. Process. Appl.* **10,** 299–311.

[9] GIANINI, J. (1977). The infinite secretary problem as the limit of the finite problem. *Ann. Prob.* **5,** 636–644.

[10] GIANINI, J. AND SAMUELS, S. M. (1976). The infinite secretary problem. *Ann. Prob.* **4,** 418–432.

[11] GNEDIN, A. V. (1996). On the full-information best-choice problem. *J. Appl. Prob.* **33,** 678–687.

[12] GNEDIN, A. V. (2004). Best choice from the planar Poisson process. *Stoch. Process. Appl.* **111,**, 317–354.

[13] GNEDIN, A. V. (2007). Recognising the last record of a sequence. *Stochastics* **79,** 199–210.

[14] GNEDIN, A. V. AND MIRETSKIY, D. I. (2007). Winning rate in the full-information best-choice problem. *J. Appl. Prob.* **44,** 560–565.

[15] HILL, T. AND KENNEDY, D. (1992). Sharp inequalities for optimal stopping with rewards based on ranks. *Ann. Appl. Prob.* **2,** 503–517.

[16] KÜHNE, R. AND RÜSCHENDORF, L. (2000). Approximation of optimal stopping problems. *Stoch. Process. Appl.* **90,** 301–325.

[17] MUCCI, A. (1973). Differential equations and optimal choice problems. *Ann. Statist.* **1,** 104–113.

[18] MUCCI, A. (1973). On a class of best-choice problems. *Ann. Prob.* **1,** 417–427.

[19] RUBIN, H. AND SAMUELS, S. M. (1977). The finite-memory secretary problem. *Ann. Prob.* **5,** 627–635.

[20] SAMUELS, S. M. (1991). Secretary problems. In *Handbook of Sequential Analysis*, eds B. K. Ghosh and P. K. Sen, Marcel Dekker, New York, pp. 381–405.

[21] SHIRYAEV, A. N. (1978). *Optimal Stopping Rules*, Springer, New York.