

QUICKEST ONLINE SELECTION OF AN INCREASING SUBSEQUENCE OF SPECIFIED SIZE

ALESSANDRO ARLOTTO, ELCHANAN MOSSEL,
AND J. MICHAEL STEELE

ABSTRACT. Given a sequence of independent random variables with a common continuous distribution, we consider the online decision problem where one seeks to *minimize the expected value of the time* that is needed to complete the selection of a monotone increasing subsequence of a prespecified length n . This problem is dual to some online decision problems that have been considered earlier, and this dual problem has some notable advantages. In particular, the recursions and equations of optimality lead with relative ease to asymptotic formulas for mean and variance of the minimal selection time.

MATHEMATICS SUBJECT CLASSIFICATION (2010): Primary: 60C05, 90C40; Secondary: 60G40, 90C27, 90C39

KEY WORDS: Increasing subsequence problem, online selection, sequential selection, time-focused decision problem, dynamic programming, Markov decision problem.

1. INCREASING SUBSEQUENCES AND TIME-FOCUSED SELECTION

If X_1, X_2, \dots is a sequence of independent random variables with a common continuous distribution F , then

$$L_n = \max\{k : X_{i_1} < X_{i_2} < \dots < X_{i_k}, \text{ where } 1 \leq i_1 < i_2 < \dots < i_k \leq n\}$$

represents the length of the longest monotone increasing subsequence in the sample $\{X_1, X_2, \dots, X_n\}$. This random variable has been subject to a remarkable sequence of investigations beginning with Ulam (1961) and culminating with Baik et al. (1999) where it was found that $n^{-1/6}(L_n - 2\sqrt{n})$ converges in distribution to the Tracy-Widom law, a new universal law that was introduced just a few years earlier in Tracy and Widom (1994). The review of Aldous and Diaconis (1999) and the monograph of Romik (2015) draw rich connections between this increasing subsequence problem and topics as diverse as card sorting, triangulation of Riemann surfaces, and — especially — the theory of integer partitions.

Here we consider a new kind of decision problem where one seeks to select *as quickly as possible* an increasing subsequence of a prespecified length n . More precisely, at time i , when the decision maker is first presented with X_i , a decision

Date: first version: December 26, 2014; this version: August 9, 2015.

A. Arlotto: The Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC, 27708. Email address: alessandro.arlotto@duke.edu.

E. Mossel: Wharton School, Department of Statistics, Huntsman Hall 459, University of Pennsylvania, Philadelphia, PA 19104 and Department of Statistics, University of California, Berkeley, 401 Evans Hall, Berkeley, CA 94720. Email address: mossel@wharton.upenn.edu.

J. M. Steele: Wharton School, Department of Statistics, Huntsman Hall 447, University of Pennsylvania, Philadelphia, PA 19104. Email address: steele@wharton.upenn.edu.

must be made either to accept X_i a member of the selected subsequence or else to reject X_i forever. The decision at time i is assumed to be a deterministic function of the observations $\{X_1, X_2, \dots, X_i\}$, so the times $1 \leq \tau_1 < \tau_2 < \dots < \tau_n$ of affirmative selections give us a strictly increasing sequence stopping times that are adapted to the sequence of σ -fields $\mathcal{F}_i = \sigma\{X_1, X_2, \dots, X_i\}$, $1 \leq i < \infty$. The quantity of most interest is

$$(1) \quad \beta(n) := \min_{\pi} \mathbb{E}[\tau_n]$$

where the minimum is over all sequences $\pi = (\tau_1, \tau_2, \dots, \tau_n)$ of stopping times such that

$$1 \leq \tau_1 < \tau_2 < \dots < \tau_n \quad \text{and} \quad X_{\tau_1} < X_{\tau_2} < \dots < X_{\tau_n}.$$

Such a sequence π will be called a *selection policy*, and the set of all such selection policies with $\mathbb{E}[\tau_n] < \infty$ will be denoted by $\Pi(n)$.

It is useful to note that the value of $\beta(n)$ is not changed if we replace each X_i with $F^{-1}(X_i)$, so we may as well assume from the beginning that the X_i 's are all uniformly distributed on $[0, 1]$. Our main results concern the behavior of $n \mapsto \beta(n)$ and the structure of the policy that attains the minimum (1).

Theorem 1. *The function*

$$n \mapsto \beta(n) = \min_{\pi \in \Pi(n)} \mathbb{E}[\tau_n]$$

is convex, $\beta(1) = 1$, and for all $n \geq 2$ one has the bounds

$$(2) \quad \frac{1}{2}n^2 \leq \beta(n) \leq \frac{1}{2}n^2 + n \log n.$$

We can also add some precision to this result at no additional cost. We just need to establish that there is always an optimal policy within the subclass of *threshold policies* $\pi = (\tau_1, \tau_2, \dots, \tau_n) \in \Pi(n)$ that are determined by real valued sequences $\{t_i \in [0, 1] : 1 \leq i \leq n\}$ and the corresponding recursions

$$(3) \quad \tau_{k+1} = \min\{i > \tau_k : X_i \in [X_{\tau_k}, X_{\tau_k} + t_{n-k}(1 - X_{\tau_k})]\}, \quad 0 \leq k < n.$$

Here, the recursion begins with $\tau_0 = 0$ and $X_0 = 0$, and one can think of t_{n-k} as the “threshold parameter” that specifies the maximum fraction that one would be willing to spend from a “residual budget” $(1 - X_{\tau_k})$ in order to accept a value that arrives after the time τ_k when the k 'th selection was made.

Theorem 2. *There is a unique threshold policy $\pi^* = (\tau_1^*, \tau_2^*, \dots, \tau_n^*) \in \Pi(n)$ for which one has*

$$(4) \quad \beta(n) = \min_{\pi \in \Pi(n)} \mathbb{E}[\tau_n] = \mathbb{E}[\tau_n^*],$$

and for this optimal policy π^* one has for all $\alpha > 2$ that

$$(5) \quad \text{Var}[\tau_n^*] = \frac{1}{3}n^3 + O(n^2 \log^\alpha n) \quad \text{as } n \rightarrow \infty.$$

In the next section, we prove the existence and uniqueness of an optimal threshold policy, and in Section 3 we complete the proof of Theorem 1 after deriving some recursions that permit the exact computation of the optimal threshold values. Section 4 deals with the asymptotics of the variance and completes the proof of Theorem 2.

In Sections 5 and 6, we develop the relationship between Theorems 1 and 2 and the more traditional *size-focused* online selection problem which was first studied in Samuels and Steele (1981) and then studied much more extensively by Bruss and Robertson (1991), Gnedin (1999), Bruss and Delbaen (2001, 2004), and Arlotto et al. (2015). On an intuitive level, the time-focused selection problem and the size-focused selection problems are dual to each other, and it is curious to consider the extent to which rigorous relationships that can be developed between the two. Finally, in Section 7 we underscore some open issues, and, in particular, we note that there are several other selection problems where it may be beneficial to explore the possibility of time versus size duality.

2. THRESHOLD POLICIES: EXISTENCE AND OPTIMALITY

A beneficial feature of the time-focused monotone selection problem is that there is a natural similarity between the problems of size n and size $n - 1$. This similarity entails a “scaled regeneration” and leads to a useful recursion for $\beta(n)$.

Lemma 1 (Variational Beta Recursion). *For all $n = 1, 2, \dots$ we have*

$$(6) \quad \beta(n) = \inf_{\tau} \mathbb{E} \left[\tau + \frac{1}{1 - X_{\tau}} \beta(n - 1) \right],$$

where the minimum is over all stopping times τ and where we initialize the recursion by setting $\beta(0) = 0$.

Proof. Suppose that the first value X_{τ} has been selected. When we condition on the value of X_{τ} , we then face a problem that amounts to the selection of $n - 1$ values from a *thinned* sequence of uniformly distributed random variables on $[X_{\tau}, 1]$. The thinned sequence may be obtained from the original sequence of uniform random variables on $[0, 1]$ by eliminating those elements that fall in $[0, X_{\tau}]$. The remaining thinned observations are then separated by geometric time gaps that have mean $1/(1 - X_{\tau})$, so conditional on X_{τ} , the optimal expected time needed for the remaining selections is $\beta(n - 1)/(1 - X_{\tau})$. Optimization over τ then yields (6). \square

This quick proof of Lemma 1 is essentially complete, but there is a variation on this argument that adds further information and perspective. If we first note that $\beta(1) = 1$, then one can confirm (6) for $n = 1$ simply by taking $\tau = 1$. One can then argue by induction. In particular, we take $n \geq 2$ and consider a *arbitrary* selection policy $\pi = (\tau_1, \tau_2, \dots, \tau_n)$.

If we set $\pi' = (\tau_2 - \tau_1, \tau_3 - \tau_1, \dots, \tau_n - \tau_1)$, then one can view π' as a selection policy for the sequence $(X'_1, X'_2, \dots) = (X_{1+\tau_1}, X_{2+\tau_1}, \dots)$ where one can only make selections from those values that fall in the interval $[X_{\tau_1}, 1]$. As before, remaining thinned observations are then separated by geometric time gaps that have mean $1/(1 - X_{\tau_1})$. Thus, conditional on X_{τ_1} , the *optimal* expected time needed for the remaining selections is $\beta(n - 1)/(1 - X_{\tau_1})$. By the putative suboptimality of the selection times $\tau_1, \tau_2, \dots, \tau_n$ one has

$$\frac{\beta(n - 1)}{(1 - X_{\tau_1})} \leq \mathbb{E}[\tau_n - \tau_1 \mid \tau_1, X_{\tau_1}] = -\tau_1 + \mathbb{E}[\tau_n \mid \tau_1, X_{\tau_1}].$$

Moreover, we note that the sum

$$\tau_1 + \frac{\beta(n - 1)}{(1 - X_{\tau_1})}$$

is the conditional expectation given τ_1 of a strategy that first selects X_{τ_1} and then proceeds optimally with the selection of $n - 1$ further values. The suboptimality of this strategy and the further suboptimality of the policy $\pi = (\tau_1, \tau_2, \dots, \tau_n)$ give us

$$(7) \quad \beta(n) \leq \mathbb{E}\left[\tau_1 + \frac{\beta(n-1)}{(1-X_{\tau_1})}\right] \leq \mathbb{E}[\tau_n].$$

The bounds of (7) may now look obvious, but since they are valid for *any* strategy one gains at least a bit of new information. At a minimum, one gets (6) immediately just by taking the infimum in (7) over all π in $\Pi(n)$.

Thus, in a sense, the bound (7) generalizes the recursion (6) which has several uses. In particular, (6) helps one to show that there is a unique threshold policy that achieves the minimal expectation $\beta(n)$.

Lemma 2 (Existence and Uniqueness of an Optimal Threshold Policy). *For $1 \leq i \leq n$, there are constants $0 \leq t_i \leq 1$ such that the threshold policy $\pi^* \in \Pi(n)$ defined by (3) is the unique optimal policy. That is, for $\pi^* = (\tau_1^*, \tau_2^*, \dots, \tau_n^*)$ one has*

$$\beta(n) = \min_{\pi \in \Pi(n)} \mathbb{E}[\tau_n] = \mathbb{E}[\tau_n^*],$$

and π^* is the only policy in $\Pi(n)$ that achieves this minimum.

Proof. The proof again proceeds by induction. The case $n = 1$ is trivial since the only optimal policy is to take any element which presented. This corresponds to the threshold policy with $t_1 = 1$ and $\beta(1) = 1$.

For the moment, we consider an arbitrary policy $\pi = (\tau_1, \tau_2, \dots, \tau_n) \in \Pi(n)$. We have $1 \leq \mathbb{E}[\tau_1] < \infty$, and we introduce a parameter t by setting $t = (\mathbb{E}[\tau_1])^{-1}$. Next, we define a new, threshold stopping time τ_1^* by setting

$$\tau_1^* = \min\{i : X_i < t\},$$

and we note that this construction gives us $\mathbb{E}[\tau_1^*] = \mathbb{E}[\tau_1] = 1/t$. For $s \in [0, t]$, we also have the trivial inequality

$$\mathbb{1}(X_{\tau_1} < s) \leq \sum_{i=1}^{\tau_1} \mathbb{1}(X_i < s),$$

so by Wald's equation we also have

$$(8) \quad \mathbb{P}(X_{\tau_1} < s) \leq \mathbb{E}\left[\sum_{i=1}^{\tau_1} \mathbb{1}(X_i < s)\right] = s\mathbb{E}[\tau_1] = s/t.$$

The definition of τ_1^* implies that $X_{\tau_1^*}$ is uniformly distributed on $[0, t]$, so we further have $\mathbb{P}(X_{\tau_1^*} < s) = \min\{1, s/t\}$. Comparison with (8) then gives us the domination relation

$$(9) \quad \mathbb{P}(X_{\tau_1} < s) \leq \mathbb{P}(X_{\tau_1^*} < s) \quad \text{for all } 0 \leq s \leq 1.$$

From (9) and the monotonicity of $x \mapsto (1-x)^{-1}$, we have by integration that

$$(10) \quad \mathbb{E}\left[\frac{\beta(n-1)}{1-X_{\tau_1^*}}\right] \leq \mathbb{E}\left[\frac{\beta(n-1)}{1-X_{\tau_1}}\right],$$

and one has a strict inequality in (9) and (10) unless $\tau_1^* = \tau_1$ with probability one.

If we now add $\mathbb{E}[\tau_1^*] = \mathbb{E}[\tau_1]$ to the corresponding sides of (10) and take the infimum over all τ_1 , then the beta recursion (6) gives us

$$\mathbb{E}[\tau_1^*] + \mathbb{E}\left[\frac{\beta(n-1)}{1-X_{\tau_1^*}}\right] \leq \inf_{\tau_1} \left\{ \mathbb{E}[\tau_1] + \mathbb{E}\left[\frac{\beta(n-1)}{1-X_{\tau_1}}\right] \right\} = \beta(n).$$

In other words, the first selection of an optimal policy is given by uniquely by a threshold rule.

To see that all subsequent selections must be made by threshold rules, we just need to note that given the time τ_1 and value $X_{\tau_1} = x$ of the first selection, one is left with a selection problem of size $n-1$ from the smaller set $\{X_i : i > \tau_1 \text{ and } X_i > x\}$. The induction hypothesis applies to this problem of size $n-1$, so we conclude that there is a unique threshold policy $(\tau_2^*, \tau_3^*, \dots, \tau_n^*)$ that is optimal for these selections. Taken as a whole, we have a unique threshold policy $(\tau_1^*, \tau_2^*, \dots, \tau_n^*) \in \Pi(n)$ for the problem of selecting an increasing subsequence of size n in minimal time. \square

Lemma 2 completes the proof of the first assertion (4) of Theorem 2. After we develop a little more information on the behavior of the mean, we will return to the proof of the second assertion (5) of Theorem 2.

3. LOWER AND UPPER BOUNDS FOR THE MEAN

The recursion (6) for $\beta(n)$ is informative, but to determine its asymptotic behavior we need more concrete and more structured recursions. The key relations are summarized in the next lemma.

Lemma 3 (Recursions for $\beta(n)$ and the Optimal Thresholds). *For each $x \geq 1$ and $t \in (0, 1)$ we let*

$$(11) \quad g(x, t) = \frac{1}{t} + \frac{x}{t} \log\left(\frac{1}{1-t}\right), \quad G(x) = \min_{0 < t < 1} g(x, t), \quad \text{and}$$

$$(12) \quad H(x) = \arg \min_{0 < t < 1} g(x, t).$$

We then have $\beta(1) = 1$, and we have the recursion

$$(13) \quad \beta(n+1) = G(\beta(n)) \quad \text{for all } n \geq 1.$$

Moreover, if the deterministic sequence t_1, t_2, \dots is defined by the recursion

$$(14) \quad t_1 = 1 \quad \text{and} \quad t_{n+1} = H(\beta(n)) \quad \text{for all } n \geq 1,$$

then the minimum in the defining equation (1) for $\beta(n)$ is uniquely achieved by the sequence of stopping times given by the threshold recursion (3).

Proof. An optimal first selection time has the form $\tau_1 = \min\{i : X_i < t\}$, so we can rewrite the recursion (6) as

$$(15) \quad \begin{aligned} \beta(n) &= \min_{0 < t < 1} \left\{ \frac{1}{t} + \mathbb{E}\left[\frac{\beta(n-1)}{1-X_{\tau_1}}\right] \right\} = \min_{0 < t < 1} \left\{ \frac{1}{t} + \frac{\beta(n-1)}{t} \int_0^t \frac{1}{1-s} ds \right\} \\ &= \min_{0 < t < 1} g(t, \beta(n-1)) \equiv G(\beta(n-1)). \end{aligned}$$

The selection rule for the first element is given by $\tau_1 = \min\{i : X_i < t_n\}$ so by (15) and the definitions of g and H we have $t_n = H(\beta(n-1))$. \square

Lemma 3 already gives us enough to prove the first assertion of Theorem 1 which states that the map $n \mapsto \beta(n)$ is convex.

Lemma 4. *The map $n \mapsto \Delta(n) := \beta(n+1) - \beta(n)$ is an increasing function.*

Proof. One can give a variational characterization of Δ that makes this evident. First, by the defining relations (11) and the recursion (13) we have

$$\begin{aligned} \beta(n+1) - \beta(n) &= G(\beta(n)) - \beta(n) \\ &= \min_{0 < t < 1} \left\{ \frac{1}{t} + \beta(n) \left[\frac{1}{t} \log \left(\frac{1}{1-t} \right) - 1 \right] \right\}, \end{aligned}$$

so if we set

$$\widehat{g}(x, t) = \frac{1}{t} + x \left[\frac{1}{t} \log \left(\frac{1}{1-t} \right) - 1 \right],$$

then we have

$$(16) \quad \Delta(n) = \beta(n+1) - \beta(n) = \min_{0 < t < 1} \widehat{g}(\beta(n), t).$$

Now, for $0 \leq x \leq y$ and $t \in (0, 1)$ we then have

$$\widehat{g}(x, t) - \widehat{g}(y, t) = (x - y) \left[\frac{1}{t} \log \left(\frac{1}{1-t} \right) - 1 \right] = (x - y) \sum_{k=2}^{\infty} \frac{1}{k} t^{k-1} \leq 0;$$

so from the monotonicity $\beta(n) \leq \beta(n+1)$, we get

$$\widehat{g}(\beta(n), t) \leq \widehat{g}(\beta(n+1), t) \quad \text{for all } 0 < t < 1.$$

When we minimize over $t \in (0, 1)$, we see that (16) gives us $\Delta(n) \leq \Delta(n+1)$. \square

We next show that the two definitions in (11) can be used to give an *a priori* lower bound on G . An induction argument using the recursion (13) can then be used to obtain the lower half of (2).

Lemma 5 (Lower Bounding G Recursion). *For the function $x \mapsto G(x)$ defined by (11), we have*

$$(17) \quad \frac{1}{2}(x+1)^2 \leq G\left(\frac{1}{2}x^2\right) \quad \text{for all } x \geq 1.$$

Proof. To prove (17), we first note that by (11) it suffices to show that one has

$$(18) \quad \delta(x, t) = (x+1)^2 t - 2 - x^2 \log \left(\frac{1}{1-t} \right) \leq 0$$

for all $x \geq 1$ and $t \in (0, 1)$. For $x \geq 1$ the map $t \mapsto \delta(x, t)$ is twice-continuous differentiable and concave in t . Hence there is a unique value $t^* \in (0, 1)$ such that $t^* = \operatorname{argmax}_{0 < t < 1} \delta(x, t)$, and such that $t^* = t^*(x)$ satisfies the first order condition

$$(x+1)^2 - (1-t^*)^{-1} x^2 = 0.$$

Solving this equation gives us

$$t^* = \frac{2x+1}{(x+1)^2}, \quad \text{and} \quad \delta(x, t^*) = -1 + 2x - 2x^2 \log \left(1 + \frac{1}{x} \right),$$

so the familiar bound

$$\frac{1}{x} - \frac{1}{2x^2} \leq \log \left(1 + \frac{1}{x} \right) \quad \text{for } x \geq 1,$$

gives us

$$\delta(x, t) \leq \delta(x, t^*) \leq -1 + 2x - 2x^2 \left(\frac{1}{x} - \frac{1}{2x^2} \right) = 0,$$

and this is just what we needed to complete the proof of (18). \square

Now, to argue by induction, we consider the hypothesis that one has

$$(19) \quad \frac{1}{2}n^2 \leq \beta(n).$$

This holds for $n = 1$ since $\beta(1) = 1$, and, if it holds for some $n \geq 1$, then by the monotonicity of G we have $G(n^2/2) \leq G(\beta(n))$. Now, by (17) and (13) we have

$$\frac{1}{2}(n+1)^2 \leq G\left(\frac{1}{2}n^2\right) \leq G(\beta(n)) = \beta(n+1),$$

and this completes our induction step from (19).

AN ALTERNATIVE OPTIMIZATION ARGUMENT¹.

From the definition of the threshold policy (3) we also know that for $0 \leq k < n$ the difference $\tau_{k+1} - \tau_k$ of the selection times is a geometrically distributed random variable with parameter $\lambda_k = t_{n-k}(1 - X_{\tau_k})$, so one has the representation

$$(20) \quad \beta(n) = \sum_{k=0}^{n-1} \mathbb{E}\left[\frac{1}{\lambda_k}\right].$$

Moreover, we also have

$$\mathbb{E}[X_{\tau_{k+1}} - X_{\tau_k} | X_{\tau_k}] = \frac{1}{2}\lambda_k$$

so, when we take expectations and sum, we see that telescoping gives us

$$(21) \quad \frac{1}{2} \sum_{k=0}^{n-1} \mathbb{E}[\lambda_k] = \mathbb{E}[X_{\tau_n}] \leq 1,$$

where the last inequality is a trivial consequence of the fact that $X_{\tau_n} \in [0, 1]$.

The relations (20) and (21) now suggest a natural optimization problem. Specifically, for $n \geq 2$ we let $\{l_k : 0 \leq k \leq n-1\}$ denote a sequence of random variables, and we consider the following convex program

$$(22) \quad \begin{aligned} & \text{minimize} && \sum_{k=0}^{n-1} \frac{1}{\mathbb{E}[l_k]} \\ & \text{subject to} && \sum_{k=0}^{n-1} \mathbb{E}[l_k] \leq 2 \\ & && \mathbb{E}[l_k] \geq 0, \quad k = 0, 2, \dots, n-1. \end{aligned}$$

This is a convex program in the n real variables $\mathbb{E}[l_k]$, $k = 0, 1, \dots, n-1$, and, by routine methods, one can confirm that its unique optimal solution is given by taking $\mathbb{E}[l_k] = 2/n$ for all $0 \leq k < n$. The corresponding optimal value is obviously equal to $n^2/2$.

By the definition of the random variables $\{\lambda_k : 0 \leq k < n\}$, we have the bounds $0 \leq \lambda_k \leq 1$, and by (21) we see that their expectations satisfy the first constraint of the optimization problem (22), so $\{\mathbb{E}[\lambda_k] : 0 \leq k < n\}$ is a feasible solution of (22). Feasibility of the values $\{\mathbb{E}[\lambda_k] : 0 \leq k < n\}$ then implies that

$$(23) \quad \frac{1}{2}n^2 \leq \sum_{k=0}^{n-1} \frac{1}{\mathbb{E}[\lambda_k]},$$

¹This subsection is based on the kind suggestions of an anonymous referee.

and, by convexity of $x \mapsto 1/x$ on $[0, 1]$, we obtain from Jensen's inequality that

$$(24) \quad \sum_{k=0}^{n-1} \frac{1}{\mathbb{E}[\lambda_k]} \leq \sum_{k=0}^{n-1} \mathbb{E}\left[\frac{1}{\lambda_k}\right] = \beta(n).$$

Taken together (23) and (24) tell us that $n^2/2 \leq \beta(n)$, so we have a second proof of (19), or the lower bound (2) of Theorem 1.

Here we should note that Gnedin (1999) used a similar convex programming argument to prove a tight *upper bound* in the corresponding *size-focused* selection problem. Thus, we have here our first instance of the kind of duality that is discussed more fully in Section 6.

COMPLETION OF THE PROOF OF THEOREM 1.

We now complete the proof of Theorem 1 by proving the upper half of (2). The argument again depends on an *a priori* bound on G . The proof is brief but delicate.

Lemma 6 (Upper Bounding G Recursion). *For the function $x \mapsto G(x)$ defined by (11) one has*

$$G\left(\frac{1}{2}x^2 + x \log(x)\right) \leq \frac{1}{2}(x+1)^2 + (x+1) \log(x+1) \quad \text{for all } x \geq 1.$$

Proof. If we set $f(x) := x^2/2 + x \log(x)$, then we need to show that

$$G(f(x)) \leq f(x+1).$$

If we take $t' = 2/(x+2)$ then the defining relation (11) for G tells us that

$$(25) \quad G(f(x)) \leq g(f(x), t') = \frac{x+2}{2} + \frac{x+2}{2} \log\left(1 + \frac{2}{x}\right) f(x).$$

Next, for any $y \geq 0$ integration over $(0, y)$ of the inequality

$$\frac{1}{u+1} \leq \frac{u^2 + 2u + 2}{2(u+1)^2} \quad \text{implies the bound} \quad \log(1+y) \leq \frac{y(y+2)}{2(y+1)}.$$

If we now set $y = 2/x$ and substitute this last bound in (25), we obtain

$$\begin{aligned} G(f(x)) &\leq \frac{x}{2} + 1 + \left(1 + \frac{1}{x}\right) f(x) \\ &= f(x+1) + \frac{1}{2} + (x+1) \{\log(x) - \log(x+1)\} \leq f(x+1), \end{aligned}$$

just as needed to complete the proof of the lemma. \square

One can now use Lemma 6 and induction to prove that for all $n \geq 2$ one has

$$(26) \quad \beta(n) \leq \frac{1}{2}n^2 + n \log(n).$$

Since $\beta(2) = G(1) = \min_{0 < t < 1} g(1, t) < 3.15$ and $2(1 + \log(2)) \approx 3.39$, one has (26) for $n = 2$. Now, for $n \geq 2$, the monotonicity of G gives us that

$$\beta(n) \leq \frac{1}{2}n^2 + n \log(n) \quad \text{implies} \quad G(\beta(n)) \leq G\left(\frac{1}{2}n^2 + n \log(n)\right).$$

Finally, by the recursion (13) and Lemma 6 we have

$$\begin{aligned}\beta(n+1) &= G(\beta(n)) \leq G\left(\frac{1}{2}n^2 + n \log(n)\right) \\ &\leq \frac{1}{2}(n+1)^2 + (n+1) \log(n+1).\end{aligned}$$

This completes the induction step and establishes (26) for all $n \geq 2$. This also completes last part of the proof of Theorem 1.

4. ASYMPTOTICS FOR THE VARIANCE

To complete the proof of Theorem 2, we only need to prove that one has the asymptotic formula (5) for $\text{Var}[\tau_n^*]$. This will first require an understanding of the size of the threshold t_n , and we can get this from our bounds on $\beta(n)$ once we have an asymptotic formula for H . The next lemma gives us what we need.

Lemma 7. *For $x \mapsto G(x)$ and $x \mapsto H(x)$ defined by (11) and (12), we have for $x \rightarrow \infty$ that*

$$(27) \quad G(x) = (x^{1/2} + 2^{-1/2})^2 \{1 + O(1/x)\} \quad \text{and}$$

$$(28) \quad H(x) = \sqrt{\frac{2}{x}} \{1 + O(x^{-1/2})\}.$$

Proof. For any fixed $x \geq 1$ we have $g(t, x) \rightarrow \infty$ when $t \rightarrow 0$ or $t \rightarrow 1$, so the minimum of $g(t, x)$ is obtained at an interior point $0 < t < 1$. Computing the t -derivative $g_t(t, x)$ gives us

$$g_t(t, x) = -\frac{1}{t^2} - \frac{x}{t^2} \log\left(\frac{1}{1-t}\right) + \frac{x}{t(1-t)},$$

so the first-order condition $g_t(t, x) = 0$ implies that at the minimum we have the equality

$$\frac{1}{t^2} = -\frac{x}{t^2} \log\left(\frac{1}{1-t}\right) + \frac{x}{t(1-t)}.$$

Writing this more informatively as

$$(29) \quad \frac{1}{x} = \log(1-t) + \frac{t}{1-t} = \frac{t^2}{2} + \sum_{i=3}^{\infty} \frac{i-1}{i} t^i,$$

we see the right-hand side is monotone in t , so there is a unique value $t_* = t_*(x)$ that solves (29) for t . The last sum on the right-hand side of (29) tells us that

$$\frac{1}{2}t_*^2 \leq \frac{1}{x} \quad \text{or, equivalently,} \quad t_* \leq \sqrt{\frac{2}{x}},$$

and when we use these bounds in (29) we have

$$\frac{t_*^2}{2} \leq \frac{1}{x} \leq \frac{t_*^2}{2} + \sum_{i=3}^{\infty} \left(\frac{2}{x}\right)^{i/2} \leq \frac{t_*^2}{2} + O(x^{-3/2}).$$

Solving these inequalities for t_* , we then have by the definition (12) of H that

$$(30) \quad H(x) = t_* = \sqrt{\frac{2}{x}} \{1 + O(x^{-1/2})\}.$$

Finally, to confirm the approximation (27), we substitute $H(x) = t_*$ into the definition (11) of G and use the asymptotic formula (30) for $H(x)$ to compute

$$\begin{aligned} G(x) = g(x, t_*) &= \frac{1}{t_*} + \frac{x}{t_*} \log \left(\frac{1}{1-t_*} \right) = \frac{1}{t_*} \left(1 + x \sum_{i=1}^{\infty} \frac{t_*^i}{i} \right) \\ &= \frac{1}{t_*} \left\{ 1 + xt_* + \frac{xt_*^2}{2} + O(xt_*^3) \right\} = \frac{1}{t_*} \left(xt_* + 1 + \frac{xt_*^2}{2} \right) + O(1) \\ &= x + 2\sqrt{\frac{x}{2}} + O(1) = \left(x^{1/2} + 2^{-1/2} \right)^2 \{ 1 + O(1/x) \}, \end{aligned}$$

and this completes the proof of the lemma. \square

The recursion (14) tells us that $t_n = H(\beta(n))$ and the upper and lower bounds (2) of Theorem 1 tell us that $\beta(n) = n^2/2 + O(n \log n)$, so by the asymptotic formula (28) for H we have

$$(31) \quad t_n = \frac{2}{n} + O(n^{-2} \log n).$$

To make good use of this formula we only need two more tools. First, we need to note that the stopping time τ_n^* satisfies a natural distributional identity. This will lead in turn to a recursion from which we can extract the required asymptotic formula for $v(n) = \text{Var}(\tau_n^*)$.

If t_n is the threshold value defined by the recursion (14), we let $\gamma(t_n)$ denote a geometric random variable of parameter $p = t_n$ and we let $U(t_n)$ denote a random variable with the uniform distribution on the interval $[0, t_n]$. Now, if we take the random variables $\gamma(t_n)$, $U(t_n)$, and τ_{n-1} to be independent, then we have the distributional identity,

$$(32) \quad \tau_n^* \stackrel{d}{=} \gamma(t_n) + \frac{\tau_{n-1}^*}{(1 - U(t_n))},$$

and this leads to a useful recursion for the variance of τ_n^* . To set this up, we first put

$$R(t) = (1 - U(t))^{-1}$$

where $U(t)$ is uniformly distributed on $[0, t]$, and we note

$$\mathbb{E}[R(t)] = -t^{-1} \log(1 - t) = 1 + t/2 + t^2/3 + O(t^3);$$

moreover, since $\mathbb{E}[R^2(t)] = (1 - t)^{-1} = 1 + t + t^2 + O(t^3)$, we also have

$$(33) \quad \text{Var}[R(t)] = (1 - t)^{-1} - t^{-2} \log^2(1 - t) = \frac{t^2}{12} + O(t^3).$$

Lemma 8 (Approximate Variance Recursion). *For the variance $v(n) := \text{Var}(\tau_n^*)$ one has the approximate recursion*

$$(34) \quad v(n) = \left\{ 1 + \frac{2}{n} + O(n^{-2} \log n) \right\} v(n-1) + \frac{1}{3} n^2 + O(n \log n).$$

Proof. By independence of the random variables on the right side of (32), we have

$$(35) \quad v(n) = \text{Var}(\gamma(t_n)) + \text{Var}[R(t_n)\tau_{n-1}^*].$$

From (31) we have $t_n = 2/n + O(n^{-2} \log n)$, so for the first summand we have

$$(36) \quad \text{Var}(\gamma(t_n)) = \frac{1}{t_n^2} - \frac{1}{t_n} = \frac{1}{4} n^2 + O(n \log n).$$

To estimate the second summand, we first use the complete variance formula and independence to get

$$\begin{aligned} \text{Var}(R(t_n)\tau_{n-1}^*) &= \mathbb{E}[R^2(t_n)]\mathbb{E}[(\tau_{n-1}^*)^2] - \mathbb{E}[R(t_n)]^2(\mathbb{E}[\tau_{n-1}^*])^2 \\ (37) \qquad \qquad \qquad &= \mathbb{E}[R^2(t_n)]v(n-1) + \text{Var}[R(t_n)](\mathbb{E}[\tau_{n-1}^*])^2. \end{aligned}$$

Now from (31) and (33) we have

$$\text{Var}[R(t_n)] = \frac{1}{3}n^{-2} + O(n^{-3} \log n),$$

and from (2) we have

$$(\mathbb{E}[\tau_{n-1}^*])^2 = \left\{ \frac{1}{2}(n-1)^2 + O(n \log n) \right\}^2 = \frac{1}{4}n^4 + O(n^3 \log n),$$

so from (35), (36) and (37) we get

$$v(n) = \left\{ 1 + \frac{2}{n} + O(n^{-2} \log n) \right\} v(n-1) + \frac{1}{3}n^2 + O(n \log n),$$

and this completes the proof of (34). \square

To conclude the proof of Theorem 2, it only remains to show that the approximate recursion (34) implies the asymptotic formula

$$(38) \qquad v(n) = \frac{1}{3}n^2(n+1) + O(n^2 \log^\alpha n) \quad \text{for } \alpha > 2.$$

If we define $r(n)$ by setting $v(n) = 3^{-1}n^2(n+1) + r(n)$, then substitution of $v(n)$ into (34) gives us a recursion for $r(n)$,

$$r(n) = \left\{ 1 + \frac{2}{n} + O(n^{-2} \log n) \right\} r(n-1) + O(n \log n).$$

We then consider the normalized values $\hat{r}(n) = r(n)/(n^2 \log^\alpha(n))$, and we note they satisfy the recursion

$$\hat{r}(n) = \{1 + O(n^{-2})\} \hat{r}(n-1) + O(n^{-1} \log^{\alpha-1} n).$$

This is a recursion of the form $\hat{r}(n+1) = \rho_n \hat{r}(n) + \epsilon_n$, and one finds by induction its solution has the representation

$$\hat{r}(n) = \hat{r}(0)\rho_0\rho_1 \cdots \rho_{n-1} + \sum_{k=0}^{n-1} \epsilon_k \rho_{k+1} \cdots \rho_{n-1}.$$

Here, the product of the ‘‘evolution factors’’ ρ_n is convergent and the sum of the ‘‘impulse terms’’ ϵ_n is finite, so the sequence $\hat{r}(n)$ is bounded, and, consequently, (34) gives us (38). This completes the proof of the last part of Theorem 2.

5. SUBOPTIMAL POLICIES AND A BLOCKING INEQUALITY

Several inequalities for $\beta(n)$ can be obtained through the construction of suboptimal policies. The next lemma illustrates this method with an inequality that leads to an alternative proof of (19), the uniform lower bound for $\beta(n)$.

Lemma 9 (Blocking Inequality). *For nonnegative integers n and m one has the inequality*

$$(39) \qquad \beta(mn) \leq \min\{m^2\beta(n), n^2\beta(m)\}.$$

Proof. First, we fix n and we consider a policy π^* that achieves the minimal expectation $\beta(n)$. The idea is to use π^* to build a suboptimal π' policy for the selection of an increasing subsequence of length mn . We take X_i , $i = 1, 2, \dots$ to be a sequence of independent random variables with the uniform distribution on $[0, 1]$, and we partition $[0, 1]$ into the subintervals $I_1 = [0, 1/m)$, $I_2 = [1/m, 2/m)$, \dots , $I_m = [(m-1)/m, 1]$. We then define π' by three rules:

- (i) Beginning with $i = 1$ we say X_i is feasible value if $X_i \in I_1$. If X_i is feasible, we accept X_i if X_i would be accepted by the policy π^* applied to the sequence of feasible values after we rescale those values to be uniform in $[0, 1]$. We continue this way until the time τ'_1 when we have selected n values.
- (ii) Next, beginning with $i = \tau'_1$, we follow the previous rule except that now we say X_i is feasible value if $X_i \in I_2$. We continue in this way until time $\tau'_1 + \tau'_2$ when n additional increasing values have been selected.
- (iii) We repeat this process $m - 2$ more times for the successive intervals I_3, I_4, \dots, I_m .

At time $\tau'_1 + \tau'_2 + \dots + \tau'_m$, the policy π' will have selected nm increasing values. For each $1 \leq j \leq m$ we have $\mathbb{E}[\tau'_j] = m\beta(n)$, so by suboptimality of π' we have

$$\beta(mn) \leq \mathbb{E}[\tau'_1 + \tau'_2 + \dots + \tau'_m] = m^2\beta(n).$$

We can now interchange the roles of m and n , so the proof of the lemma is complete. \square

The blocking inequality (39) implies that even the crude asymptotic relation $\beta(n) = n^2/2 + o(n^2)$ is strong enough to imply the uniform lower bound $n^2/2 \leq \beta(n)$. Specifically, one simply notes from (39) and $\beta(n) = n^2/2 + o(n^2)$ that

$$\frac{\beta(mn)}{(mn)^2} \leq \frac{\beta(n)}{n^2} \quad \text{and} \quad \lim_{m \rightarrow \infty} \frac{\beta(mn)}{(mn)^2} = \frac{1}{2}.$$

This third derivation of the uniform bound $n^2/2 \leq \beta(n)$ seems to have almost nothing in common with the proof by induction that was used in the proof of Lemma 6. Still, it does require the bootstrap bound $\beta(n) = n^2/2 + o(n^2)$, and this does require at least some of the machinery of Lemma 3.

6. DUALITY AND THE SIZE-FOCUSED SELECTION PROBLEM

In the online *size-focused* selection problem one considers a set of policies $\Pi_s(n)$ that depend on the size n of a sample $\{X_1, X_2, \dots, X_n\}$, and the goal is to make sequential selections in order to maximize the expected size of the selected increasing subsequence. More precisely, a policy $\pi_n \in \Pi_s(n)$ is determined by stopping times τ_i , $i = 1, 2, \dots$ such that $1 \leq \tau_1 < \tau_2 < \dots < \tau_k \leq n$ and $X_{\tau_1} \leq X_{\tau_2} \leq \dots \leq X_{\tau_k}$. The random variable of interest is

$$L_n^o(\pi_n) = \max\{k : X_{\tau_1} < X_{\tau_2} < \dots < X_{\tau_k} \text{ where } 1 \leq \tau_1 < \tau_2 < \dots < \tau_k \leq n\},$$

and most previous analyses have focused on the asymptotic behavior of

$$\ell(n) := \max_{\pi_n \in \Pi_s(n)} \mathbb{E}[L_n^o(\pi_n)].$$

For example, Samuels and Steele (1981) found that $\ell(n) \sim \sqrt{2n}$, but now a number of refinements of this are known. Our goal here is to show how some of these refinements follow from the preceding theory.

UNIFORM UPPER BOUND FOR $\ell(n)$ VIA DUALITY.

Perhaps the most elegant refinement of $\ell(n) \sim \sqrt{2n}$ is the following uniform upper bound that follows independently from related analyses of Bruss and Robertson (1991) and Gnedin (1999).

Bruss and Robertson (1991) obtained the upper bound (40) by exploiting the equivalence of the size-focused monotone subsequence problem with a special knapsack problem that we will briefly revisit in Section 7. On the other hand, Gnedin (1999) showed that (40) can be proved by an optimization argument like the one used to prove (22). Here, we give a third argument with the curious feature that one gets a sharp bound that holds for all $n \geq 1$ from an apparently looser relation that only holds approximately for large n .

Proposition 1 (Uniform Upper Bound). *For all $n \geq 1$, one has*

$$(40) \quad \ell(n) \leq \sqrt{2n}.$$

As we noted above, this proposition is now well understood, but it still seems instructive to see how it can be derived from $\beta(n) = (1/2)n^2 + O(n \log n)$. The basic idea is that one exploits duality with a suboptimality argument like the one used in Section 5, although, in this case a bit more work is required.

We fix n , and, for a much larger integer k , we set

$$(41) \quad N_k = \lfloor (k - 2k^{2/3})\ell(n) \rfloor \quad \text{and} \quad r_k = \lfloor k - k^{2/3} \rfloor.$$

The idea of the proof is to give an algorithm that is guaranteed to select from $\{X_1, X_2, \dots\}$ an increasing subsequence of length N_k . If T_k is the number of the elements that the algorithm inspects before returning the increasing subsequence, then by the definition of $\beta(\cdot)$ we have $\beta(N_k) \leq \mathbb{E}[T_k]$. One then argues that (40) follows from this relation.

We now consider $[0, 1]$ and for $1 \leq i \leq r_k$, we consider the disjoint intervals $I_i = [(i-1)/k, i/k]$ and a final ‘‘reserve’’ interval $I^* = [r_k/k, 1]$ that is added to complete the partition of $[0, 1]$. Next, we let $\nu(1)$ be the first integer such that

$$\mathcal{S}_1 := \{X_1, X_2, \dots, X_{\nu(1)}\} \cap I_1$$

has cardinality n , and for each $i > 1$ we define $\nu(i)$ to be least integer greater $\nu(i-1)$ for which the set $\mathcal{S}_i := \{X_{\nu(i-1)+1}, X_{\nu(i-1)+2}, \dots, X_{\nu(i)}\} \cap I_i$ has cardinality n . By Wald’s lemma and (41) we have

$$(42) \quad \mathbb{E}[\nu(r_k)] = nkr_k \quad \text{where} \quad r_k = \lfloor k - k^{2/3} \rfloor.$$

Now, for each $1 \leq i \leq n$, we run the *optimal size-focused* sequential selection algorithm on \mathcal{S}_i , and we let $L(n, i)$ be the length of the subsequence that one obtains. The random variables $L(n, i)$, $1 \leq i \leq r_k$, are independent, identically distributed, and each has mean equal to $\ell(n)$. We then set

$$\mathcal{L}(n, r_k) = L(n, 1) + L(n, 2) + \dots + L(n, r_k),$$

and we note that if $\mathcal{L}(n, r_k) \geq N_k$, for N_k as defined in (41), then we have extracted an increasing subsequence of length at least N_k ; in this case, we halt the procedure.

On the other hand if $\mathcal{L}(n, r_k) < N_k$, we need to send in the reserves. Specifically, we recall that $I^* = [r_k/k, 1]$ and we consider the post- $\nu(r_k)$ reserve subsequence

$$\mathcal{S}^* := \{X_i : i > \nu(r_k) \text{ and } X_i \in I^*\}.$$

We now rescale the elements of \mathcal{S}^* to the unit interval, and we run the optimal *time-focused* algorithm on \mathcal{S}^* until we get an increasing sequence of length N_k . If we let $R(n, k)$ denote the number of observations from \mathcal{S}^* that are examined in this case, then we have $\mathbb{E}[R(n, k)] = \beta(N_k)$ by the definition of β . Finally, since I^* has length at least $k^{-1/3}$, the expected number of elements of $\{X_i : i > \nu(r_k)\}$ that need to be inspected before we have selected our increasing subsequence of length N_k is bounded above by $k^{1/3}\beta(N_k)$.

The second phase of our procedure may seem wasteful, but one rarely needs to use the reserve subsequence. In any event, our procedure does guarantee that we find an increasing subsequence of length N_k in a finite amount of time T_k . By (42) and the upper bound $k^{1/3}\beta(N_k)$ on the incremental cost when one needs to use the reserve subsequence, we have

$$(43) \quad \beta(N_k) \leq \mathbb{E}[T_k] \leq knr_k + \{knr_k + k^{1/3}\beta(N_k)\}\mathbb{P}(\mathcal{L}(n, r_k) < N_k),$$

where, as noted earlier, the first inequality comes from the definition of β .

The summands of $\mathcal{L}(n, r_k)$ are uniformly bounded by n and $\mathbb{E}[\mathcal{L}(n, r_k)] = r_k\ell(n)$, so by the definition (41) of N_k and r_k we see from Hoeffding's inequality that

$$(44) \quad \mathbb{P}(\mathcal{L}(n, r_k) < N_k) \leq \mathbb{P}\left(\mathcal{L}(n, r_k) - \mathbb{E}[\mathcal{L}(n, r_k)] < -(k^{2/3} - 1)\ell(n)\right) \\ \leq \exp\{-A_n k^{1/3}\},$$

for constants A_n, K_n , and all $k \geq K_n$. The exponential bound (44) tells us that for each n there is a constant C_n such that the last summand in (43) is bounded by C_n for all k . By the bounds (2) of Theorem 1 we have $\beta(N_k) = (1/2)N_k^2 + O(N_k \log N_k)$, and by (41) we have

$$N_k = (k - 2k^{2/3})\ell(n) + O(1), \quad r_k = k - k^{1/3} + O(1),$$

so, in the end, our estimate (43) tell us

$$\frac{1}{2}\ell^2(n)\{k^2 - 2k^{5/3} + 4k^{4/3}\} \leq k^2n + o_n(k^2).$$

When we divide by k^2 and let $k \rightarrow \infty$, we find $\ell(n) \leq \sqrt{2n}$, just as we hoped.

LOWER BOUNDS FOR $\ell(n)$ AND THE DUALITY GAP.

One can use the time-focused tools to get a lower bound for $\ell(n)$, but in this case the slippage, or duality gap, is substantial. To sketch the argument, we first let T_r denote the time required by the optimal time-focused selection policy to select r values. We then follow the r -target time-focused policy. Naturally, we stop if we have selected r values, but if we have not selected r values by time n , then we quit, no matter how many values we have selected. This suboptimal strategy gives us the bound $r\mathbb{P}(T_r \leq n) \leq \ell(n)$, and from this bound and Chebyshev's inequality, we then have

$$(45) \quad r\{1 - \text{Var}(T_r)/(n - \mathbb{E}[T_r])^2\} \leq \ell(n).$$

If we then use the estimates (2) and (5) for $\mathbb{E}[T_r]$ and $\text{Var}[T_r]$ and optimize over r , then (45) gives us the lower bound $(2n)^{1/2} - O(n^{1/3})$. However in this case the time-focused bounds and the duality argument leave a big gap.

Earlier, by different methods — and for different reasons — Rhee and Talagrand (1991) and Gnedin (1999) obtained the lower bound $(2n)^{1/2} - O(n^{1/4}) \leq \ell(n)$. Subsequently, Bruss and Delbaen (2001) studied a continuous time interpretation

of the online increasing subsequence problem where the observations are presented to the decision maker at the arrival times of a unit-rate Poisson process on the time interval $[0, t)$, and, in this new formulation, they found the stunning lower bound $\sqrt{2t} - O(\log t)$. Much later, Arlotto et al. (2015) showed by a de-Poissonization argument that the lower bound of Bruss and Delbaen (2001) can be used to obtain

$$\sqrt{2n} - O(\log n) \leq \ell(n) \quad \text{for all } n \geq 1$$

under the traditional discrete-time model for sequential selection. Duality estimates such as (45) seem unlikely to recapture this bound.

7. OBSERVATIONS, CONNECTIONS, AND PROBLEMS

A notable challenge that remains is that of understanding the asymptotic distribution of τ_n^* , the time at which one completes the selection of n increasing values by following the unique optimal policy π^* that minimizes the expected time $\mathbb{E}[\tau_n^*] = \beta(n)$. By Theorems 1 and 2 we know the behavior of the mean $\mathbb{E}[\tau_n^*]$ and of the variance $\text{Var}[\tau_n^*]$ for large n , but the asymptotic behavior of higher moments or characteristic functions does not seem amenable to the methods used here.

Some second order asymptotic distribution theory should be available for τ_n^* , but it is not likely to be easy. Still, some encouragement can be drawn from the central limit theorem for the size-focused selection problem that was obtained by Bruss and Delbaen (2004) under the Poisson model introduced in Bruss and Delbaen (2001). Also, by quite different means, Arlotto et al. (2015) obtained the central limit theorem for the size-focused selection problem under the traditional model discussed in Section 6. Each of these developments depends on martingale arguments that do not readily suggest any natural analog in the case of time-focused selection.

There are also several other size-focused selection problems where it may be informative to investigate the corresponding time-focused problem. The most natural of these is probably the unimodal subsequence selection problems where one considers the random variable

$$U_n^o(\pi_n) = \max\{k : X_{\tau_1} < X_{\tau_2} < \cdots < X_{\tau_t} > X_{\tau_{t+1}} > \cdots > X_{\tau_k}, \text{ where} \\ 1 \leq \tau_1 < \tau_2 < \cdots < \tau_k \leq n\},$$

and where, as usual, each τ_k is a stopping time. Arlotto and Steele (2011) found that

$$\max_{\pi_n} \mathbb{E}[U_n^o(\pi_n)] \sim 2\sqrt{n}, \quad \text{as } n \rightarrow \infty,$$

and an analogous time-focused selection problem is easy to pose. Unfortunately, it does not seem easy to frame a useful analog of Lemma 3, so the time-focused analysis of the unimodal selection problem remains open.

For another noteworthy possibility one can consider the time-focused problem in a multidimensional setting. Here one would take the random variables X_i , $i = 1, 2, \dots$, to have the uniform distribution on $[0, 1]^d$, and the decision maker's task is then to select as quickly as possible a subsequence that is monotone increasing in each coordinate. The dual, size-focused, version of this problem was studied by Baryshnikov and Gnedin (2000) who characterized the asymptotic behavior of the optimal mean.

Finally, a further possibility is the analysis of a time-focused knapsack problem. Here one considers a sequence X_1, X_2, \dots of independent item sizes with common

continuous distribution F , and the decision maker's task is to select as quickly as possible a set of n items for which the total size does not exceed a given knapsack capacity c . The analytical goal is then to estimate the size of the objective function

$$\phi_{F,c}(n) = \min_{\tau_1, \dots, \tau_n} \mathbb{E} \left[\tau_n : \sum_{k=1}^n X_{\tau_k} \leq c \right],$$

where, just as before, the decision variables $\tau_1 < \tau_2 < \dots < \tau_n$ are stopping times. The optimal mean $\phi_{F,c}(n)$ now depends on the model distribution F and on the capacity parameter c . Never the less, if F represents the uniform distribution on $[0, 1]$ and if one takes $c = 1$, then the time-focused knapsack problem is equivalent to the time-focused increasing subsequence problem; specifically, one has

$$\phi_{\text{Unif},1}(n) = \beta(n), \quad \text{for all } n \geq 1.$$

The dual, size-focused, version of this problem was first studied by Coffman et al. (1987), Rhee and Talagrand (1991), and Bruss and Robertson (1991) who proved asymptotic estimates for the mean. While the structure of the optimal policy is now well understood (see, e.g. Papastavrou et al., 1996), only little is known about the limiting distribution of the optimal number of size-focused knapsack selections. From Arlotto et al. (2014) one has a suggestive upper bound on the variance, but at this point one does not know for sure that this bound gives the principle term of the variance for large n .

ACKNOWLEDGEMENTS

Elchanan Mossel acknowledges the support of NSF grants DMS 1106999 and CCF 1320105, ONR grant number N00014-14-1-0823 and grant 328025 from the Simons Foundation.

The authors also thank the referees for their useful suggestions, including a clarification of the proof of Lemma 1 and the alternative optimization proof of (19) given in Section 3.

REFERENCES

- Aldous, D. and Diaconis, P. (1999), 'Longest increasing subsequences: from patience sorting to the Baik-Deift-Johansson theorem', *Bull. Amer. Math. Soc. (N.S.)* **36**(4), 413–432.
- Arlotto, A., Gans, N. and Steele, J. M. (2014), 'Markov decision problems where means bound variances', *Operations Research* **62**(4), 864–875.
- Arlotto, A., Nguyen, V. V. and Steele, J. M. (2015), 'Optimal online selection of a monotone subsequence: a central limit theorem', *Stochastic Processes and their Applications* **125**(9), 3596–3622.
- Arlotto, A. and Steele, J. M. (2011), 'Optimal sequential selection of a unimodal subsequence of a random sequence', *Combin. Probab. Comput.* **20**(6), 799–814.
- Baik, J., Deift, P. and Johansson, K. (1999), 'On the distribution of the length of the longest increasing subsequence of random permutations', *J. Amer. Math. Soc.* **12**(4), 1119–1178.
- Baryshnikov, Y. M. and Gnedin, A. V. (2000), 'Sequential selection of an increasing sequence from a multidimensional random sample', *Ann. Appl. Probab.* **10**(1), 258–267.

- Bruss, F. T. and Delbaen, F. (2001), ‘Optimal rules for the sequential selection of monotone subsequences of maximum expected length’, *Stochastic Process. Appl.* **96**(2), 313–342.
- Bruss, F. T. and Delbaen, F. (2004), ‘A central limit theorem for the optimal selection process for monotone subsequences of maximum expected length’, *Stochastic Process. Appl.* **114**(2), 287–311.
- Bruss, F. T. and Robertson, J. B. (1991), ‘“Wald’s lemma” for sums of order statistics of i.i.d. random variables’, *Adv. in Appl. Probab.* **23**(3), 612–623.
- Coffman, Jr., E. G., Flatto, L. and Weber, R. R. (1987), ‘Optimal selection of stochastic intervals under a sum constraint’, *Adv. in Appl. Probab.* **19**(2), 454–473.
- Gnedin, A. V. (1999), ‘Sequential selection of an increasing subsequence from a sample of random size’, *J. Appl. Probab.* **36**(4), 1074–1085.
- Papastavrou, J. D., Rajagopalan, S. and Kleywegt, A. J. (1996), ‘The dynamic and stochastic knapsack problem with deadlines’, *Management Science* **42**(12), 1706–1718.
- Rhee, W. and Talagrand, M. (1991), ‘A note on the selection of random variables under a sum constraint’, *J. Appl. Probab.* **28**(4), 919–923.
- Romik, D. (2015), *The Surprising Mathematics of Longest Increasing Subsequences*, Cambridge University Press, Cambridge.
- Samuels, S. M. and Steele, J. M. (1981), ‘Optimal sequential selection of a monotone sequence from a random sample’, *Ann. Probab.* **9**(6), 937–947.
- Tracy, C. A. and Widom, H. (1994), ‘Level-spacing distributions and the Airy kernel’, *Comm. Math. Phys.* **159**(1), 151–174.
- Ulam, S. M. (1961), Monte Carlo calculations in problems of mathematical physics, in ‘Modern mathematics for the engineer: Second series’, McGraw-Hill, New York, pp. 261–281.