
Optimal Sequential Selection of a Unimodal Subsequence of a Random Sequence

ALESSANDRO ARLOTTO¹ and J. MICHAEL STEELE²

¹ Department of Operations and Information Management, Wharton School, Huntsman Hall 527.2,
University of Pennsylvania, Philadelphia, PA 19104, USA
(e-mail: alear@wharton.upenn.edu)

² Department of Statistics, Wharton School, Huntsman Hall 447,
University of Pennsylvania, Philadelphia, PA 19104, USA
(e-mail: steele@wharton.upenn.edu)

Received 19 October 2010; revised 25 August 2011; first published online 5 October 2011

We consider the problem of selecting sequentially a *unimodal subsequence* from a sequence of independent identically distributed random variables, and we find that a person doing optimal sequential selection does so within a factor of the square root of two as well as a prophet who knows all of the random observations in advance of any selections. Our analysis applies in fact to selections of subsequences that have $d + 1$ monotone blocks, and, by including the case $d = 0$, our analysis also covers monotone subsequences.

1. Introduction

A classical result of Erdős and Szekeres [9] tells us that, in any sequence x_1, x_2, \dots, x_n of n real numbers, there is a subsequence of length $k = \lceil n^{1/2} \rceil$ that is either monotone increasing or monotone decreasing. More precisely, given x_1, x_2, \dots, x_n , one can always find a subsequence $1 \leq n_1 < n_2 < \dots < n_k \leq n$ for which we either have

$$x_{n_1} \leq x_{n_2} \leq \dots \leq x_{n_k}, \quad \text{or} \quad x_{n_1} \geq x_{n_2} \geq \dots \geq x_{n_k}.$$

Many years later, Fan Chung [8] considered the analogous problem for unimodal sequences. Specifically, she sought to determine the maximum value ℓ_n such that, in any sequence of n real values x_1, x_2, \dots, x_n , one can find a subsequence $x_{i_1}, x_{i_2}, \dots, x_{i_k}$ of length $k = \ell_n$ and a ‘turning place’ $1 \leq t \leq k$ for which one either has

$$\begin{aligned} x_{i_1} \leq x_{i_2} \leq \dots \leq x_{i_t} \geq x_{i_{t+1}} \geq \dots \geq x_{i_k}, \quad \text{or} \\ x_{i_1} \geq x_{i_2} \geq \dots \geq x_{i_t} \leq x_{i_{t+1}} \leq \dots \leq x_{i_k}. \end{aligned}$$

Through a sustained and instructive analysis, she surprisingly obtained an exact formula:

$$\ell_n = \lceil (3n - 3/4)^{1/2} - 1/2 \rceil.$$

Shortly afterwards, Steele [14] considered unimodal subsequences of permutations, or equivalently, unimodal subsequences of a sequence of n independent, uniformly distributed random variables X_1, X_2, \dots, X_n . For the random variables

$$U_n = \max\{k : X_{i_1} \leq X_{i_2} \leq \dots \leq X_{i_t} \geq X_{i_{t+1}} \geq \dots \geq X_{i_k}, \text{ where } 1 \leq i_1 < i_2 < \dots < i_k \leq n\},$$

and

$$D_n = \max\{k : X_{i_1} \geq X_{i_2} \geq \dots \geq X_{i_t} \leq X_{i_{t+1}} \leq \dots \leq X_{i_k}, \text{ where } 1 \leq i_1 < i_2 < \dots < i_k \leq n\},$$

it was established that

$$\mathbb{E}[\max\{U_n, D_n\}] \sim \mathbb{E}[U_n] \sim \mathbb{E}[D_n] \sim 2(2n)^{1/2} \text{ as } n \rightarrow \infty. \tag{1.1}$$

Here we consider analogues of the random variables U_n, D_n and $L_n = \max\{U_n, D_n\}$, but instead of seeing the whole sequence all at once, one observes the variables sequentially. Thus, for each $1 \leq i \leq n$, the chooser must decide at time i , when X_i is first presented, whether to accept or reject X_i as an element of the unimodal subsequence. The sequential (or on-line) selection for the much simpler problem of a monotone subsequence – the analogue of the original Erdős and Szekeres [9] problem – was considered long ago by Samuels and Steele [13].

1.1. Main results

We denote by $\Pi(n)$ the set of all feasible policies for the unimodal sequential selection problem for $\{X_1, X_2, \dots, X_n\}$ where these random variables are independent with a common continuous distribution function F . Given any feasible sequential selection policy $\pi_n \in \Pi(n)$, if we let τ_k denote the index of the k th selected element, then for each k the value τ_k is a stopping time with respect to the increasing sequence of σ -fields $\mathcal{F}_i = \sigma\{X_1, X_2, \dots, X_i\}$, $1 \leq i \leq n$. In terms of these stopping times, the random variable

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

is the length of the unimodal subsequence that is selected by the policy π_n . For the moment, we just consider unimodal subsequences that begin with an increasing piece and end with a decreasing piece; either of these pieces is permitted to have size one.

For each n there is a policy $\pi_n^* \in \Pi(n)$ that maximizes the expected length of the selected subsequence, and the main issue is to determine the asymptotic behaviour of this expected value. The answer turns out to have an informative relationship to the off-line selection problem. A prophet with knowledge of the whole sequence before making his choices will do better than an optimal on-line chooser, but he will only do better by a factor of $\sqrt{2}$.

Theorem 1.1 (Expected length of optimal unimodal subsequences). *For each $n \geq 1$, there is a $\pi_n^* \in \Pi(n)$, such that*

$$\mathbb{E}[U_n^o(\pi_n^*)] = \sup_{\pi_n \in \Pi(n)} \mathbb{E}[U_n^o(\pi_n)],$$

and for such an optimal policy one has the upper bound

$$\mathbb{E}[U_n^o(\pi_n^*)] < 2n^{1/2}$$

and the lower bound

$$2n^{1/2} - 4(\pi/6)^{1/2}n^{1/4} - O(1) < \mathbb{E}[U_n^o(\pi_n^*)],$$

which combine to give the asymptotic formula

$$\mathbb{E}[U_n^o(\pi_n^*)] \sim 2n^{1/2} \text{ as } n \rightarrow \infty.$$

In a natural sense that we will shortly make precise, the optimal policy π_n^* is unique. Consequently, one can ask about the *distribution* of the length $U_n^o(\pi_n^*)$ of the subsequence that is selected by the optimal policy, and there is a pleasingly general argument that gives an upper bound for the variance. Moreover, that bound is good enough to provide a weak law for $U_n^o(\pi_n^*)$.

Theorem 1.2 (Variance bound). *For the unique optimal policy $\pi_n^* \in \Pi(n)$, we have the bounds*

$$\text{Var}[U_n^o(\pi_n^*)] \leq \mathbb{E}[U_n^o(\pi_n^*)] < 2n^{1/2}. \tag{1.2}$$

Corollary 1.3 (Weak law for unimodal sequential selections). *For the sequence of optimal policies $\pi_n^* \in \Pi(n)$, we have the limit*

$$U_n^o(\pi_n^*)/\sqrt{n} \xrightarrow{P} 2 \text{ as } n \rightarrow \infty.$$

Organization of the proofs. The proof of Theorem 1.1 comes in two halves. First, we show by an elaboration of an argument of Gnedin [10] that there is an *a priori* upper bound for $\mathbb{E}[U_n^o(\pi_n)]$ for all n and all $\pi_n \in \Pi(n)$. This argument uses almost nothing about the structure of the selection policy beyond the fact from Section 4 that it suffices to consider policies that are specified by acceptance *intervals*. For the lower bound we simply construct a good (but suboptimal) policy. Here there is an obvious candidate, but the proof of its efficacy seems to be more delicate than one might have expected.

The proof of Theorem 1.2 in Section 3 exploits a martingale that comes naturally from the Bellman equation. The summands of the quadratic variation of this martingale are then found to have a fortunate relationship to the probability that an observation is selected. It is this ‘self-bounding’ feature that leads one to the bound (1.2) of the variance by the mean.

In Section 5 we outline analogues of Theorems 1.1 and 1.2 for subsequences that can be decomposed into $d + 1$ alternating monotone blocks (rather than just two). If one takes $d = 0$, this reduces to the monotone subsequence problem, and in this case only the variance bound is new. Finally, in Section 6 we comment briefly on two conjectures. These deal with a more refined understanding of $\text{Var}[U_n^o(\pi_n^*)]$ and with the naturally associated central limit theorem.

2. Mean bounds: Proof of Theorem 1.1

Since the distribution F is assumed to be continuous and since the problem is unchanged by replacing X_i by its monotone transformation $F^{-1}(X_i)$, we can assume without loss of generality that the X_i are uniformly distributed on $[0, 1]$. Next, we introduce two tracking variables. First, we let S_i denote the value of the last element that has been selected up to and including time i . We then let R_i denote an indicator variable that tracks the monotonicity of the selected subsequence: specifically, we set $R_i = 0$ if the selections made up to and including time i are increasing; otherwise we set $R_i = 1$.

The sequence of real values $\{S_i : R_i = 0, 1 \leq i \leq n\}$ is thus a monotone increasing sequence, though of course not in the strict sense because there will typically be long patches where the successive values of S_i do not change. Similarly, $\{S_i : R_i = 1, 1 \leq i \leq n\}$ is a monotone decreasing sequence, and the full sequence $\{S_i : 1 \leq i \leq n\}$ is a unimodal sequence – in the non-strict sense that permits ‘flat spots’. As a convenience for later formulas, we also set $S_0 = 0$ and $R_0 = 0$.

2.1. The class of feasible interval policies

Here we will consider feasible policies that have acceptance sets that are given by intervals. It is reasonably obvious that any optimal policy must have this structure, but for completeness we give a formal proof of this fact in Section 4.

Now, if the value X_i is under consideration for selection, two possible scenarios can occur. If $R_{i-1} = 0$ (so one is in the ‘increasing part’ of the selected subsequence), then a selectable X_i can be *above or below* S_{i-1} . On the other hand, if $R_{i-1} = 1$ (and one is in the ‘decreasing part’ of the selected subsequence), then any selectable X_i has to be smaller than S_{i-1} . Thus, to specify a feasible interval policy, we just need to specify for each i an interval $[a, b] \subset [0, 1]$ where we accept X_i if $X_i \in [a, b]$ and reject it otherwise. Here, the values of the end-points of the interval are functions of i , S_{i-1} , and R_{i-1} . In longhand, we write the acceptance interval as

$$\Delta_i(S_{i-1}, R_{i-1}) \equiv [a(i, S_{i-1}, R_{i-1}), b(i, S_{i-1}, R_{i-1})].$$

There are some restrictions on the functions $a(i, S_{i-1}, R_{i-1})$ and $b(i, S_{i-1}, R_{i-1})$. To make these explicit we consider two sets of functions, \mathcal{A} and \mathcal{B} . We say $a \in \mathcal{A}$ provided that $a : \{1, 2, \dots, n\} \times [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ and

$$0 \leq a(i, s, r) \leq s \quad \text{for all } s \in [0, 1], r \in \{0, 1\} \text{ and } 1 \leq i \leq n.$$

Similarly, we say $b \in \mathcal{B}$ provided that $b : \{1, 2, \dots, n\} \times [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ and

$$\begin{aligned} s &\leq b(i, s, 0) \leq 1 \quad \text{for all } s \in [0, 1] \text{ and } 1 \leq i \leq n, \\ 0 &\leq b(i, s, 1) = s \quad \text{for all } s \in [0, 1] \text{ and } 1 \leq i \leq n. \end{aligned}$$

Together a pair $(a, b) \in \mathcal{A} \times \mathcal{B}$ defines an *interval policy* $\pi_n \in \Pi(n)$, where we accept X_i at time i if and only if $X_i \in \Delta_i(S_{i-1}, R_{i-1})$. We let $\Pi'(n)$ denote the set of feasible interval policies.

2.2. Three representations

First we note that for S_i we have a simple update rule driven by whether X_i is rejected or accepted:

$$S_i = \begin{cases} S_{i-1} & \text{if } X_i \notin \Delta_i(S_{i-1}, R_{i-1}), \\ X_i & \text{if } X_i \in \Delta_i(S_{i-1}, R_{i-1}). \end{cases}$$

For the sequence $\{R_i\}$ the update rule is initialized by setting $R_0 = 0$; one should then note that only one change takes place in the values of the sequence $\{R_i\}$. Specifically, we change to $R_i = 1$ at the first i such that $S_i < S_{i-1}$, *i.e.*, the first instance where we have a decrease in our sequence of selected values. For specificity, we can rewrite this rule as

$$R_i = \begin{cases} 1 & \text{if } X_i \in \Delta_i(S_{i-1}, R_{i-1}) \\ & \text{and } S_{i-1} = \max\{S_k : 1 \leq k \leq i\}, \\ R_{i-1} & \text{otherwise.} \end{cases} \tag{2.1}$$

Finally, using $\mathbb{1}(E)$ to denote the indicator function of the event E , we see by counting the occurrences of the ‘selection events’ $X_i \in \Delta_i(S_{i-1}, R_{i-1})$, that for each $1 \leq k \leq n$ the number of selections made up to and including time k is given by the sum of the indicators

$$U_k^o(\pi_n) = \sum_{i=1}^k \mathbb{1}(X_i \in \Delta_i(S_{i-1}, R_{i-1})). \tag{2.2}$$

2.3. Proof of the upper bound: an *a priori* prophet inequality

The immediate task is to show that for all $n \geq 1$ and all $\pi_n \in \Pi'(n)$, we have the inequality

$$\mathbb{E}[U_n^o(\pi_n)] < 2n^{1/2}. \tag{2.3}$$

It will then follow from Proposition 4.1 that the bound (2.3) holds for all $\pi_n \in \Pi(n)$. We start with the representation (2.2), and then after two applications of the Cauchy–Schwarz inequality we have

$$\begin{aligned} \mathbb{E}[U_n^o(\pi_n)] &= \sum_{i=1}^n \mathbb{E}[b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1})] \\ &\leq n^{1/2} \left\{ \sum_{i=1}^n (\mathbb{E}[b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1})])^2 \right\}^{1/2} \\ &\leq n^{1/2} \left\{ \sum_{i=1}^n \mathbb{E}[(b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1}))^2] \right\}^{1/2}. \end{aligned}$$

The target bound (2.3) is therefore an immediate consequence of the following – curiously general – lemma.

Lemma 2.1 (Telescoping bound). *For each $n \geq 1$ and for any strategy $\pi_n \in \Pi'(n)$, we have the inequality*

$$\sum_{i=1}^n \mathbb{E}[(b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1}))^2] < 4. \tag{2.4}$$

Proof. We first introduce a bookkeeping function $g : [0, 1] \times \{0, 1\} \rightarrow [0, 2]$ by setting

$$g(s, r) = \begin{cases} s & \text{if } r = 0, \\ 2 - s & \text{if } r = 1. \end{cases}$$

Trivially g is bounded by 2, and we will argue by conditioning and telescoping that the left-hand side of inequality (2.4) is bounded above by $2 \mathbb{E}[g(S_n, R_n)] < 4$. Specifically, if we condition on \mathcal{F}_{i-1} , then the independence and uniform distribution of X_i gives us, after a few lines of straightforward calculation, that

$$\begin{aligned} &\mathbb{E}[g(S_i, R_i) - g(S_{i-1}, 0) \mid \mathcal{F}_{i-1}] \\ &= \int_{a(i, S_{i-1}, 0)}^{S_{i-1}} (g(x, 1) - S_{i-1}) dx + \int_{S_{i-1}}^{b(i, S_{i-1}, 0)} (g(x, 0) - S_{i-1}) dx \\ &= \frac{1}{2} (b(i, S_{i-1}, 0) - a(i, S_{i-1}, 0))^2 \\ &\quad + (S_{i-1} - a(i, S_{i-1}, 0))(2 - S_{i-1} - b(i, S_{i-1}, 0)). \end{aligned}$$

Since the last summand is non-negative we have the tidier bound

$$(b(i, S_{i-1}, 0) - a(i, S_{i-1}, 0))^2 \leq 2 \mathbb{E}[g(S_i, R_i) - g(S_{i-1}, 0) \mid \mathcal{F}_{i-1}]. \tag{2.5}$$

By an analogous direct calculation we also have the identity

$$\begin{aligned} \mathbb{E}[g(S_i, 1) - g(S_{i-1}, 1) \mid \mathcal{F}_{i-1}] &= \int_{a(i, S_{i-1}, 1)}^{S_{i-1}} (g(x, 1) - g(S_{i-1}, 1)) dx \\ &= \frac{1}{2} (b(i, S_{i-1}, 1) - a(i, S_{i-1}, 1))^2. \end{aligned} \tag{2.6}$$

Since $R_{i-1} = 1$ implies $R_i = 1$, we can write $g(S_i, R_i) - g(S_{i-1}, R_{i-1})$ as the sum

$$\{g(S_i, R_i) - g(S_{i-1}, 0)\} \mathbb{1}(R_{i-1} = 0) + \{g(S_i, 1) - g(S_{i-1}, 1)\} \mathbb{1}(R_{i-1} = 1),$$

so the two bounds (2.5) and (2.6) give us the key estimate

$$(b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1}))^2 \leq 2 \mathbb{E}[g(S_i, R_i) - g(S_{i-1}, R_{i-1}) \mid \mathcal{F}_{i-1}].$$

Finally, when we take the total expectation and sum, we see that telescoping gives

$$\sum_{i=1}^n \mathbb{E}[(b(i, S_{i-1}, R_{i-1}) - a(i, S_{i-1}, R_{i-1}))^2] \leq 2 \mathbb{E}[g(S_n, R_n)] < 4,$$

just as needed. □

2.4. Proof of the lower bound: exploitation of suboptimality

We construct an explicit policy $\tilde{\pi}_n \in \Pi(n)$ that is close enough to optimal to give us the bound

$$2n^{1/2} - 4(\pi/6)^{1/2}n^{1/4} - O(1) < \mathbb{E}[U_n^o(\pi_n^*)]. \tag{2.7}$$

The basic idea is to make an approximately optimal choice of an increasing subsequence from the sample $\{X_i : 1 \leq i \leq n/2\}$ and an approximately optimal choice of a decreasing subsequence from the sample $\{X_i : n/2 + 1 \leq i \leq n\}$. The cost of giving up a flexible choice of the ‘turn-around time’ is substantial, but this class of policies is still close enough to optimal to give required bound (2.7).

For the moment, we assume that n is even. We then select observations according to the following process.

- For $1 \leq i \leq n/2$ we select the observation X_i if and only if X_i falls in the interval between S_{i-1} and $\min\{1, S_{i-1} + 2n^{-1/2}\}$.
- We set $S_{n/2} = 1$, and for $n/2 + 1 \leq i \leq n$ we select the observation X_i if and only if X_i falls in the interval between $\max\{0, S_{i-1} - 2n^{-1/2}\}$ and S_{i-1} .

Here, of course, the selections for $1 \leq i \leq n/2$ are increasing and the selections for $n/2 + 1 \leq i \leq n$ are decreasing, so the selected subsequence is indeed unimodal.

We then consider the stopping time

$$v = \min\{i : S_i > 1 - 2n^{-1/2} \text{ or } i \geq n/2\},$$

and we note that the representation (2.2), the suboptimality of the policy $\tilde{\pi}_n$, and the symmetry between our policy on $1 \leq i \leq n/2$ and on $n/2 + 1 \leq i \leq n$ will give us the lower bound

$$2 \mathbb{E} \left[\sum_{i=1}^v \mathbb{1}(X_i \in [S_{i-1}, S_{i-1} + 2n^{-1/2}]) \right] \leq \mathbb{E}[U_n^o(\tilde{\pi}_n)] \leq \mathbb{E}[U_n^o(\pi_n^*)]. \tag{2.8}$$

Wald’s lemma now tells us that

$$\mathbb{E} \left[\sum_{i=1}^v \mathbb{1}(X_i \in [S_{i-1}, S_{i-1} + 2n^{-1/2}]) \right] = 2n^{-1/2} \mathbb{E}[v],$$

so we have

$$4n^{-1/2} \mathbb{E}[v] \leq \mathbb{E}[U_n^o(\pi_n^*)].$$

The main task is to estimate $\mathbb{E}[v]$. It is a small but bothersome point that the summands $\mathbb{1}(X_i \in [S_{i-1}, S_{i-1} + 2n^{-1/2}])$ are not i.i.d. over the entirety of the range $i \in [1, n/2]$; the distribution of the last terms differs from that of the predecessors. To deal with this nuisance, we take $Z_j, 1 \leq j < \infty$, to be a sequence of random variables defined by setting

$$Z_j = \begin{cases} 0 & \text{with probability } 1 - 2n^{-1/2}, \\ U_j & \text{with probability } 2n^{-1/2}, \end{cases}$$

where the U_j are independent and uniformly distributed on $[0, 2n^{-1/2}]$. Easy calculations now give us for all $1 \leq j < \infty$ that

$$\mathbb{E}Z_j = \frac{2}{n}, \quad \text{Var}[Z_j] = \frac{8n^{1/2} - 12}{3n^2} < \frac{8}{3n^{3/2}}, \quad \text{and} \quad |Z_j - \mathbb{E}Z_j| < \frac{2}{n^{1/2}}. \tag{2.9}$$

Next, if we set $\tilde{S}_0 \equiv 0$ and put

$$\tilde{S}_i = \sum_{j=1}^i Z_j, \quad \text{for } 1 \leq i \leq n,$$

for $1 \leq i \leq v$, we have $S_i \stackrel{d}{=} \tilde{S}_i$. Setting $\tilde{v} = \min\{i : \tilde{S}_i > 1 - 2n^{-1/2} \text{ or } i \geq n/2\}$ we also have $v \stackrel{d}{=} \tilde{v}$, so to estimate $\mathbb{E}[v]$ it then suffices to estimate

$$\mathbb{E}[\tilde{v}] = \sum_{i=0}^{n/2-1} \mathbb{P}(\tilde{v} > i) = \sum_{i=0}^{n/2-1} \mathbb{P}(\tilde{S}_i \leq 1 - 2n^{-1/2}) = \frac{n}{2} - \sum_{i=0}^{n/2-1} \mathbb{P}(\tilde{S}_i > 1 - 2n^{-1/2}).$$

The proof of the lower bound (2.7) will then be complete once we check that

$$\sum_{i=0}^{n/2-1} \mathbb{P}(\tilde{S}_i > 1 - 2n^{-1/2}) < (\pi/6)^{1/2} n^{3/4} + \lceil n^{1/2} \rceil. \tag{2.10}$$

This bound turns out to be a reasonably easy consequence of Bernstein’s inequality (see Lugosi [11, Theorem 6]), which asserts that for *any* i.i.d. sequence $\{Z_j\}$ with the almost sure bound $|Z_j - \mathbb{E}Z_j| \leq M$, we have for all $t > 0$ that

$$\mathbb{P}\left(\sum_{j=1}^i \{Z_j - \mathbb{E}Z_j\} > t\right) \leq \exp\left\{-\frac{t^2}{2i\text{Var}[Z_1] + 2Mt/3}\right\}.$$

If we set $n^* = \lfloor n/2 - n^{1/2} - 1 \rfloor$, then Bernstein’s inequality together with the bounds (2.9) and some simplification will give us

$$\begin{aligned} \sum_{i=0}^{n/2-1} \mathbb{P}(\tilde{S}_i > 1 - 2n^{-1/2}) &\leq \lceil n^{1/2} \rceil + \sum_{i=0}^{n^*} \mathbb{P}(\tilde{S}_i > 1 - 2n^{-1/2}) \\ &\leq \lceil n^{1/2} \rceil + \sum_{i=0}^{n^*} \exp\left\{-\frac{3(-2i - 2n^{1/2} + n)^2}{8n(n^{1/2} - 1)}\right\}. \end{aligned}$$

The summands are increasing, so the sum is bounded by

$$\int_0^{n/2-n^{1/2}} \exp\left\{-\frac{3(-2u - 2n^{1/2} + n)^2}{8n(n^{1/2} - 1)}\right\} du = (2/3)^{1/2} (n^{3/2} - n)^{1/2} \int_0^{\alpha(n)} e^{-u^2} du,$$

where $\alpha(n) = (3/8)^{1/2} (n^{1/2} - 2)(n^{1/2} - 1)^{-1/2}$. Upon bounding the last integral by $\pi^{1/2}/2$, we then complete the proof of the target bound (2.10). Finally, we note that if n is odd, we can simply ignore the last observation at the cost of decreasing our lower bound by at most one.

Remark. A benefit of Bernstein’s inequality (and the slightly sharper Bennett inequality) is that one gets to take advantage of the good bound on $\text{Var}[Z_j]$. The workhorse Hoeffding inequality would be blind to this useful information.

3. Variance bound: Proof of Theorem 1.2

To prove the variance bound in Theorem 1.2 we need some of the machinery of the Bellman equation and dynamic programming. To introduce the classical backward induction, we first set $v_i(s, r)$ equal to the expected length of the longest unimodal subsequence of $\{X_i, X_{i+1}, \dots, X_n\}$ that is obtained by sequential selection when $S_{i-1} = s$ and $R_{i-1} = r$. We then have the ‘terminal conditions’

$$v_n(s, 0) = 1, \quad v_n(s, 1) = s, \quad \text{for all } s \in [0, 1],$$

and we set

$$v_{n+1}(s, r) \equiv 0 \quad \text{for all } s \in [0, 1] \text{ and } r \in \{0, 1\}.$$

For $1 \leq i \leq n - 1$ we have the *Bellman equation*:

$$v_i(s, r) = \begin{cases} \int_0^s \max\{v_{i+1}(s, 0), 1 + v_{i+1}(x, 1)\} dx \\ \quad + \int_s^1 \max\{v_{i+1}(s, 0), 1 + v_{i+1}(x, 0)\} dx & \text{if } r = 0, \\ (1 - s)v_{i+1}(s, 1) \\ \quad + \int_0^s \max\{v_{i+1}(s, 1), 1 + v_{i+1}(x, 1)\} dx & \text{if } r = 1. \end{cases} \tag{3.1}$$

One should note that the map $s \mapsto v_i(s, 0)$ is continuous and strictly decreasing on $[0, 1]$ for $1 \leq i \leq n - 1$ with $v_n(s, 0) = 1$ for all $s \in [0, 1]$. In addition, the map $s \mapsto v_i(s, 1)$ is continuous and strictly increasing on $[0, 1]$ for all $1 \leq i \leq n$.

If we now define $a^* : \{1, 2, \dots, n\} \times [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ by setting

$$a^*(i, s, r) = \inf\{x \in [0, s] : v_{i+1}(s, r) \leq 1 + v_{i+1}(x, 1)\}, \tag{3.2}$$

then we have $a^* \in \mathcal{A}$. Similarly, if we define $b^* : \{1, 2, \dots, n\} \times [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ by setting

$$b^*(i, s, r) = \begin{cases} \sup\{x \in [s, 1] : v_{i+1}(s, 0) \leq 1 + v_{i+1}(x, 0)\} & \text{if } r = 0, \\ s & \text{if } r = 1. \end{cases} \tag{3.3}$$

then we have $b^* \in \mathcal{B}$. Here, $a^*(i, s, r)$ and $b^*(i, s, r)$ are state-dependent thresholds for which one is indifferent between (i) selecting the current observation x , adjusting r to r' as in (2.1), and continuing to act optimally with new state pair (x, r') , or (ii) rejecting the current observation, x , and continuing to act optimally with unchanged state pair, (s, r) .

By the Bellman equation (3.1) and the continuity and monotonicity properties of the value function, the values a^* and b^* provide us with a unique acceptance interval for all $1 \leq i \leq n$ and all pairs (s, r) . The policy π_n^* associated with a^* and b^* then accepts X_i at time $1 \leq i \leq n$ if and only if

$$X_i \in \Delta_i^*(S_{i-1}, R_{i-1}) \equiv [a^*(i, S_{i-1}, R_{i-1}), b^*(i, S_{i-1}, R_{i-1})],$$

where, as in Section 2, S_{i-1} is the value of the last observation selected up to and including time $i - 1$, and R_{i-1} tracks the direction of the monotonicity of the subsequence selected up to and including time $i - 1$. In Section 4 we will prove that this policy is indeed the unique optimal policy for the sequential selection of a unimodal subsequence.

We do not need a detailed analysis of a^* and b^* , but it is useful to collect some facts. In particular, one should note that $a^*(i, s, r) = 0$ whenever $v_{i+1}(s, r) \leq 1$ and $b^*(i, s, 0) = 1$ whenever $v_{i+1}(s, 0) \leq 1$. In addition, the difference $b^*(i, s, r) - a^*(i, s, r)$ provides us with an explicit bound on the increments of the value function $v_i(s, r)$, as the following lemma suggests.

Lemma 3.1. *For all $s \in [0, 1]$, $r \in \{0, 1\}$ and $1 \leq i \leq n$, we have*

$$0 \leq v_i(s, r) - v_{i+1}(s, r) \leq b^*(i, s, r) - a^*(i, s, r) \leq 1. \tag{3.4}$$

Proof. The lower bound is trivial and it follows by the fact that $v_i(s, r)$ is strictly decreasing in i for each $(s, r) \in [0, 1] \times \{0, 1\}$.

For the upper bound, we first assume that $r = 0$. Then, subtracting $v_{i+1}(s, 0)$ on both sides of equation (3.1) when $r = 0$ and using the definition of a^* and b^* , we obtain

$$\begin{aligned} v_i(s, 0) - v_{i+1}(s, 0) &= -(b^*(i, s, r) - a^*(i, s, r))v_{i+1}(s, 0) \\ &\quad + \int_{a^*(i, s, r)}^s (1 + v_{i+1}(x, 1)) dx + \int_s^{b^*(i, s, r)} (1 + v_{i+1}(x, 0)) dx. \end{aligned}$$

Recalling the monotonicity property of $s \mapsto v_{i+1}(s, r)$, we then have

$$\begin{aligned} v_i(s, 0) - v_{i+1}(s, 0) &\leq -(b^*(i, s, r) - a^*(i, s, r))v_{i+1}(s, 0) \\ &\quad + (s - a^*(i, s, r))(1 + v_{i+1}(s, 1)) + (b^*(i, s, r) - s)(1 + v_{i+1}(s, 0)), \end{aligned}$$

and since $v_{i+1}(s, 1) \leq v_{i+1}(s, 0)$, we finally obtain

$$v_i(s, 0) - v_{i+1}(s, 0) \leq b^*(i, s, r) - a^*(i, s, r) \leq 1,$$

as (3.4) requires. The proof for $r = 1$ is very similar and it is therefore omitted. □

We now come to the main lemma of this section.

Lemma 3.2. *The process defined by*

$$Y_i = U_i^o(\pi_n^*) + v_{i+1}(S_i, R_i), \quad \text{for all } 0 \leq i \leq n,$$

is a martingale with respect to the natural filtration $\{\mathcal{F}_i\}_{0 \leq i \leq n}$. Moreover, for the martingale difference sequence $d_i = Y_i - Y_{i-1}$, we have that

$$|d_i| = |Y_i - Y_{i-1}| \leq 1 \quad \text{for all } 1 \leq i \leq n.$$

Proof. We first note that Y_i is \mathcal{F}_i -measurable and bounded. Then, from the definition of $v_i(s, r)$ we have that $v_i(S_{i-1}, R_{i-1}) = \mathbb{E}[U_n^o(\pi_n^*) - U_{i-1}^o(\pi_n^*) \mid \mathcal{F}_{i-1}]$. Thus,

$$Y_i = U_i^o(\pi_n^*) + \mathbb{E}[U_n^o(\pi_n^*) - U_i^o(\pi_n^*) \mid \mathcal{F}_i] = \mathbb{E}[U_n^o(\pi_n^*) \mid \mathcal{F}_i],$$

which is clearly a martingale.

To see that the martingale differences are bounded let

$$W_i = v_{i+1}(S_{i-1}, R_{i-1}) - v_i(S_{i-1}, R_{i-1})$$

represent the change in Y_i if we do not select X_i , and let

$$Z_i = (1 + v_{i+1}(X_i, \mathbb{1}(X_i < S_{i-1})) - v_{i+1}(S_{i-1}, R_{i-1}))\mathbb{1}(X_i \in \Delta_i^*(S_{i-1}, R_{i-1}))$$

represent the change when we do select X_i . We then have that

$$d_i = W_i + Z_i,$$

and by our Lemma 3.1 we know that $-1 \leq W_i \leq 0$. Moreover, the definition of the threshold functions a^* and b^* and the monotonicity property of $s \mapsto v_{i+1}(s, r)$ give us that $0 \leq Z_i \leq 1$, so that $|d_i| \leq 1$, as desired. □

3.1. Final argument for the variance bound

For the martingale differences $d_i = Y_i - Y_{i-1}$ we have

$$Y_n - Y_0 = \sum_{i=1}^n d_i, \quad \text{and} \quad \text{Var}[Y_n] = \mathbb{E} \left[\sum_{i=1}^n d_i^2 \right],$$

and we also have the initial representation

$$Y_0 = U_0^o(\pi_n^*) + v_1(S_0, R_0) = v_1(0, 0) = \mathbb{E} [U_n^o(\pi_n^*)]$$

and the terminal identity

$$Y_n = U_n^o(\pi_n^*) + v_{n+1}(S_n, R_n) = U_n^o(\pi_n^*).$$

We now recall the decomposition $d_i = W_i + Z_i$ introduced in the proof of Lemma 3.2, where

$$W_i = v_{i+1}(S_{i-1}, R_{i-1}) - v_i(S_{i-1}, R_{i-1})$$

and

$$Z_i = (1 + v_{i+1}(X_i, \mathbb{1}(X_i < S_{i-1})) - v_{i+1}(S_{i-1}, R_{i-1}))\mathbb{1}(X_i \in \Delta_i^*(S_{i-1}, R_{i-1})).$$

Since W_i is \mathcal{F}_{i-1} -measurable, we have

$$\mathbb{E}[d_i^2 \mid \mathcal{F}_{i-1}] = \mathbb{E}[Z_i^2 \mid \mathcal{F}_{i-1}] + 2 W_i \mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] + W_i^2.$$

We also have $0 = \mathbb{E}[d_i \mid \mathcal{F}_{i-1}] = W_i + \mathbb{E}[Z_i \mid \mathcal{F}_{i-1}]$ so

$$\mathbb{E}[d_i^2 \mid \mathcal{F}_{i-1}] = \mathbb{E}[Z_i^2 \mid \mathcal{F}_{i-1}] - W_i^2. \tag{3.5}$$

Finally, from the definition of Z_i , a^* and b^* we obtain

$$\begin{aligned} \mathbb{E}[Z_i^2 \mid \mathcal{F}_{i-1}] &= \int_{a^*(i, S_{i-1}, R_{i-1})}^{b^*(i, S_{i-1}, R_{i-1})} (1 + v_{i+1}(x, \mathbb{1}(x < S_{i-1})) - v_{i+1}(S_{i-1}, R_{i-1}))^2 dx \\ &\leq b^*(i, S_{i-1}, R_{i-1}) - a^*(i, S_{i-1}, R_{i-1}), \end{aligned}$$

since the integrand is bounded by 1. Summing (3.5), applying the last bound, and taking expectations gives us

$$\text{Var}[U_n^o(\pi_n^*)] \leq \sum_{i=1}^n \mathbb{E}[b^*(i, S_{i-1}, R_{i-1}) - a^*(i, S_{i-1}, R_{i-1})] = \mathbb{E}[U_n^o(\pi_n^*)],$$

where the last equality follows from our basic representation (2.2).

4. Intermezzo: Optimality and uniqueness of interval policies

The unimodal sequential selection problem is a finite-horizon Markov decision problem with bounded rewards and finite action space, and for such a problem it is known that there exists a non-randomized Markov policy π_n^* that is optimal (see Bertsekas and Shreve [2, Corollary 8.5.1]). This amounts to saying that there exists an optimal strategy π_n^* such that for each i , S_{i-1} and R_{i-1} , there is a Borel set $D_i^*(S_{i-1}, R_{i-1}) \subseteq [0, 1]$ such that X_i is accepted if and only if $X_i \in D_i^*(S_{i-1}, R_{i-1})$. Here we just want to show that the Borel sets $D_i^*(S_{i-1}, R_{i-1})$ are actually intervals (up to null sets).

Given the optimal acceptance sets $D_i^*(S_{i-1}, R_{i-1})$, $1 \leq i \leq n$, we now set

$$v_i(S_{i-1}, R_{i-1}) = \mathbb{E}\left[\sum_{k=i}^n \mathbb{1}(X_k \in D_k^*(S_{k-1}, R_{k-1})) \mid \mathcal{F}_{i-1}\right],$$

so we have the recursion

$$v_i(S_{i-1}, R_{i-1}) = \mathbb{E}[\mathbb{1}(X_i \in D_i^*(S_{i-1}, R_{i-1})) + v_{i+1}(S_i, R_i) \mid \mathcal{F}_{i-1}], \tag{4.1}$$

where $v_i(s, r)$ is simply the optimal expected number of selections made from the subsample $\{X_i, X_{i+1}, \dots, X_n\}$ given that $S_{i-1} = s$ and $R_{i-1} = r$. We then note that $v_n(s, 0) = 1$ for all $s \in [0, 1]$, and one can check by induction on i that the map $s \mapsto v_i(s, 0)$ is continuous and strictly decreasing in s for $1 \leq i \leq n - 1$. A similar argument also gives that the map $s \mapsto v_i(s, 1)$ is continuous and strictly increasing in s for all $1 \leq i \leq n$.

If we now set

$$a(i, S_{i-1}, R_{i-1}) = \text{ess inf } D_i(S_{i-1}, R_{i-1}) \quad \text{and} \\ b(i, S_{i-1}, R_{i-1}) = \text{ess sup } D_i(S_{i-1}, R_{i-1}),$$

then we want to show for all $1 \leq i \leq n$ and all (S_{i-1}, R_{i-1}) that we have

$$\mathbb{P}(\{D_i(S_{i-1}, R_{i-1})^c \cap [a(i, S_{i-1}, R_{i-1}), b(i, S_{i-1}, R_{i-1})]\}) = 0.$$

To argue by contradiction, we suppose that there is an $1 \leq i \leq n$ and an acceptance set $D_i^* \equiv D_i^*(S_{i-1}, R_{i-1})$ that is not equivalent to an interval; *i.e.*, we suppose

$$\mathbb{P}(\{D_i^{*c} \cap [a^*(i, S_{i-1}, R_{i-1}), b^*(i, S_{i-1}, R_{i-1})]\}) > 0. \tag{4.2}$$

We then consider the sets

$$L_i = [0, S_{i-1}] \cap D_i^* \quad \text{and} \quad U_i = [S_{i-1}, 1] \cap D_i^*,$$

and we introduce the intervals

$$\tilde{L}_i = [S_{i-1} - |L_i|, S_{i-1}] \quad \text{and} \quad \tilde{U}_i = [S_{i-1}, S_{i-1} + |U_i|],$$

where $|A|$ denotes the Lebesgue measure of a set A . The set $\tilde{D}_i = \tilde{L}_i \cup \tilde{U}_i$ is also an interval and $|\tilde{D}_i| = |D_i^*|$, so, if we can show that

$$\mathbb{E}[\mathbb{1}(X_i \in D_i^*) + v_{i+1}(S_i, R_i)] < \mathbb{E}[\mathbb{1}(X_i \in \tilde{D}_i) + v_{i+1}(S_i, R_i)], \tag{4.3}$$

then the representation (4.1) tells us that policy π_n^* is not optimal, a contradiction.

To prove the bound (4.3), we note that

$$\begin{aligned} & \mathbb{E}[\mathbb{1}(X_i \in \tilde{D}_i) + v_{i+1}(S_i, R_i) | \mathcal{F}_{i-1}] - \mathbb{E}[\mathbb{1}(X_i \in D_i^*) + v_{i+1}(S_i, R_i) | \mathcal{F}_{i-1}] \\ &= \mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in \tilde{D}_i) | \mathcal{F}_{i-1}] - \mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in D_i^*) | \mathcal{F}_{i-1}] \end{aligned}$$

since \tilde{D}_i and D_i^* are \mathcal{F}_{i-1} -measurable and $\mathbb{E}[\mathbb{1}(X_i \in \tilde{D}_i) | \mathcal{F}_{i-1}] = \mathbb{E}[\mathbb{1}(X_i \in D_i^*) | \mathcal{F}_{i-1}]$. By our construction, we also have the identities

$$\mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in \tilde{D}_i) | \mathcal{F}_{i-1}] = \int_{\tilde{L}_i} v_{i+1}(x, 1) dx + \int_{\tilde{U}_i} v_{i+1}(x, 0) dx, \tag{4.4}$$

and

$$\mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in D_i^*) | \mathcal{F}_{i-1}] = \int_{L_i} v_{i+1}(x, 1) dx + \int_{U_i} v_{i+1}(x, 0) dx. \tag{4.5}$$

Now, since $|L_i| = |\tilde{L}_i|$ implies that $|\tilde{L}_i \cap L_i^c| = |L_i \cap \tilde{L}_i^c|$, we can write

$$\begin{aligned} \int_{\tilde{L}_i} v_{i+1}(x, 1) dx - \int_{L_i} v_{i+1}(x, 1) dx &= \int_{\tilde{L}_i \cap L_i^c} v_{i+1}(x, 1) dx - \int_{L_i \cap \tilde{L}_i^c} v_{i+1}(x, 1) dx \\ &= (\beta_i - \alpha_i)|\tilde{L}_i \cap L_i^c|, \end{aligned} \tag{4.6}$$

where $\alpha_i = \alpha_i(S_{i-1}, R_{i-1})$, and $\beta_i = \beta_i(S_{i-1}, R_{i-1})$ are chosen according to the mean value theorem for integrals. The sets $\tilde{L}_i \cap L_i^c$ and $L_i \cap \tilde{L}_i^c$ are almost surely disjoint since $\tilde{L}_i \cap L_i^c \subset [S_{i-1} - |L_i|, S_{i-1}]$ and $L_i \cap \tilde{L}_i^c \subset [0, S_{i-1} - |L_i|]$. So, we find that $\alpha_i < \beta_i$ since $v_{i+1}(x, 1)$ is strictly decreasing in x .

A perfectly analogous argument tells us that we can write

$$\int_{\tilde{U}_i} v_{i+1}(x, 1) dx - \int_{U_i} v_{i+1}(x, 1) dx = (\delta_i - \gamma_i)|\tilde{U}_i \cap U_i^c|, \tag{4.7}$$

where $\gamma_i < \delta_i$ and γ_i and δ_i depend on (S_{i-1}, R_{i-1}) . If we now set

$$c_i(S_{i-1}, R_{i-1}) = \min\{\beta_i - \alpha_i, \delta_i - \gamma_i\},$$

then the identities (4.4) and (4.5) and the differences (4.6) and (4.7) give us the bound

$$c_i(S_{i-1}, R_{i-1})|\tilde{D}_i \cap D_i^{*c}| \leq \mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in \tilde{D}_i) - v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in D_i^*) | \mathcal{F}_{i-1}].$$

Since $c_i(S_{i-1}, R_{i-1}) > 0$, the assumption (4.2) implies that the left-hand side above is strictly positive. When we take total expectation we get

$$0 < \mathbb{E}[v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in \tilde{D}_i) - v_{i+1}(X_i, R_i)\mathbb{1}(X_i \in D_i^*)].$$

In view of the recursion (4.1), this contradicts the optimality of π^* . This completes the proof of (4.3), and, in summary, we have the following proposition.

Proposition 4.1. *If π_n^* is an optimal non-randomized Markov policy for the unimodal sequential selection problem, then, up to sets of measure zero, π^* is an interval policy. \square*

Corollary 4.2. *There is a unique policy $\pi_n^* \in \Pi(n)$ that is optimal.*

To prove the corollary one combines the optimality of the interval policy given by Proposition 4.1 with the monotonicity properties of the Bellman equation (3.1). Specifically, the map $s \mapsto v_i(s, 0)$ is strictly decreasing in s for all $1 \leq i \leq n - 1$ and the map $s \mapsto v_i(s, 1)$ is strictly increasing in s for all $1 \leq i \leq n$, so the equations (3.2) and (3.3) determine the values $a^*(\cdot)$ and $b^*(\cdot)$ uniquely.

5. Generalizations and specializations: d -modal subsequences

There are natural analogues of Theorems 1.1 and 1.2 for ‘ d -modal subsequences’, by which we mean subsequences that are allowed to make ‘ d -turns’ rather than just one. Equivalently these are subsequences that are the concatenation of (at most) $d + 1$ monotone subsequences. If we let $U_n^{o,d}(\pi_n^*)$ denote the analogue of $U_n^o(\pi_n^*)$ when the selected subsequence is d -modal, then the arguments of the preceding sections may be adapted to provide information on the expected value of $U_n^{o,d}(\pi_n^*)$ and its variance. Here one should keep in mind that the case $d = 0$ is *not* excepted; the arguments of the preceding sections do indeed apply to the selection of monotone subsequences.

Theorem 5.1 (Expected length of optimal d -modal subsequences). *If $\Pi(n)$ denotes the class of feasible policies for the d -modal subsequence selection problem, then there is a unique $\pi_n^* \in \Pi(n)$ such that*

$$\mathbb{E}[U_n^{o,d}(\pi_n^*)] = \sup_{\pi_n \in \Pi(n)} \mathbb{E}[U_n^{o,d}(\pi_n)].$$

Moreover, for all $n \geq 1$ and $d \geq 0$ we have

$$c(d)^{1/2}n^{1/2} - c(d)^{3/4}(\pi/3)^{1/2}n^{1/4} - O(1) < \mathbb{E}[U_n^{o,d}(\pi_n^*)] < c(d)^{1/2}n^{1/2}, \quad (5.1)$$

where $c(d) = 2(d + 1)$. In particular, we have

$$\mathbb{E}[U_n^{o,d}(\pi_n^*)] \sim \{2(d + 1)\}^{1/2}n^{1/2} \quad \text{as } n \rightarrow \infty.$$

One should note that the case $d = 0$ corresponds to the *monotone subsequence selection problem* studied by Samuels and Steele [13] and more recently by Gnedin [10]. The monotone selection problem is also equivalent to certain bin packing problems studied by Bruss and Robertson [7] and Rhee and Talagrand [12].

In the special case of $d = 0$, our upper bound (5.1) agrees with that of Bruss and Robertson [7] as well as with the result of Gnedin [10]. Our lower bound (5.1) on the mean for $d = 0$ turns out to be slightly worse than that of Rhee and Talagrand [12], since our constant for the $n^{1/4}$ term is $2^{3/4}(\pi/3)^{1/2} \sim 1.72$, while theirs is $8^{1/4} \sim 1.68$.

For the d -modal problem, one can also prove a variance bound that generalizes Theorem 1.2 in a natural way.

Theorem 5.2 (Variance bound for d -modal subsequences). *For the unique optimal policy $\pi_n^* \in \Pi(n)$ we have the bound*

$$\text{Var}[U_n^{o,d}(\pi_n^*)] \leq \mathbb{E}[U_n^{o,d}(\pi_n^*)].$$

Chebyshev’s inequality and Theorem 5.2 now combine as usual to provide a weak law for $U_n^{o,d}(\pi_n^*)$. Even for $d = 0$ this variance bound is new.

6. Two conjectures

Numerical studies for small d and moderate n support the conjecture that we have the asymptotic relation

$$\text{Var}[U_n^{o,d}(\pi_n^*)] \sim \frac{1}{3} \mathbb{E}[U_n^{o,d}(\pi_n^*)] \quad \text{as } n \rightarrow \infty. \tag{6.1}$$

As observed by an anonymous referee, the methods of Section 3 and the concavity of the value function established in Samuels and Steele [13] are in fact enough to prove an appropriate lower bound:

$$\frac{1}{3} \mathbb{E}[U_n^{o,d}(\pi_n^*)] - 2 < \text{Var}[U_n^{o,d}(\pi_n^*)] \quad \text{where } d = 0. \tag{6.2}$$

Here one should now be able to prove an upper bound on $\text{Var}[U_n^{o,d}(\pi_n^*)]$ that is strong enough to establish the case $d = 0$ of the conjecture (6.1), but confirmation of this has eluded us.

Also, by numerical calculations of the optimal policy π_n^* and by subsequent simulations of $U_n^{o,d}(\pi_n^*)$ for $d = 0, d = 1$, and modest values of n , it seems likely that the random variable $U_n^{o,d}(\pi_n^*)$ obeys a central limit theorem. Specifically, the natural conjecture is that for all $d \geq 0$ we have

$$\frac{\sqrt{3}(U_n^{o,d}(\pi_n^*) - \sqrt{2(d+1)n})}{(2(d+1)n)^{1/4}} \implies N(0, 1) \quad \text{as } n \rightarrow \infty. \tag{6.3}$$

Implicit in this conjecture is the belief that the lower bound (5.1) can be improved to $\{2(d+1)n\}^{1/2} - o(n^{1/4})$, or better.

So far, the only central limit theorem available for a sequential selection problem is that obtained by Bruss and Delbaen [5, 6] for a Poissonized version of the monotone subsequence problem. Given the sequential nature of the problem, it appears to be difficult to de-Poissonize the results of Bruss and Delbaen [6] to obtain conclusions about the distribution of $U_n^{o,d}(\pi_n^*)$ even for $d = 0$.

For completeness, we should note that even for the *off-line* unimodal subsequence problem, not much more is known about the random variable U_n than its asymptotic expected value (1.1). Here one might hope to gain some information about the distribution of U_n by the methods of Bollobás and Brightwell [3] and Bollobás and Janson [4], and it is even feasible – but only remotely so – that one could extend the famous distributional results of Baik, Deift and Johansson [1] to unimodal subsequences. More modestly, one certainly should be able to prove that the distribution of U_n is *not* asymptotically normal.

One motivation for going after such a result would be to underline how the restriction to sequential strategies can bring one back to the domain of the central limit theorem.

Acknowledgement

We are grateful to an insightful referee who outlined the proof of the bound (6.2) and who suggested the conjecture (6.1) for $d = 0$.

References

- [1] 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** 1119–1178.
- [2] Bertsekas, D. P. and Shreve, S. E. (1978) *Stochastic Optimal Control: The Discrete Time Case*, Vol. 139 of *Mathematics in Science and Engineering*, Academic Press.
- [3] Bollobás, B. and Brightwell, G. (1992) The height of a random partial order: Concentration of measure. *Ann. Appl. Probab.* **2** 1009–1018.
- [4] Bollobás, B. and Janson, S. (1997) On the length of the longest increasing subsequence in a random permutation. In *Combinatorics, Geometry and Probability: Cambridge 1993*, Cambridge University Press, pp. 121–128.
- [5] Bruss, F. T. and Delbaen, F. (2001) Optimal rules for the sequential selection of monotone subsequences of maximum expected length. *Stochastic Process. Appl.* **96** 313–342.
- [6] 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** 287–311.
- [7] Bruss, F. T. and Robertson, J. B. (1991) ‘Wald’s lemma’ for sums of order statistics of i.i.d. random variables. *Adv. Appl. Probab.* **23** 612–623.
- [8] Chung, F. R. K. (1980) On unimodal subsequences. *J. Combin. Theory Ser. A* **29** 267–279.
- [9] Erdős, P. and Szekeres, G. (1935) A combinatorial problem in geometry. *Compositio Math.* **2** 463–470.
- [10] Gnedin, A. V. (1999) Sequential selection of an increasing subsequence from a sample of random size. *J. Appl. Probab.* **36** 1074–1085.
- [11] Lugosi, G. (2009) Concentration-of-measure inequalities. www.econ.upf.edu/~lugosi/anu.pdf.
- [12] Rhee, W. and Talagrand, M. (1991) A note on the selection of random variables under a sum constraint. *J. Appl. Probab.* **28** 919–923.
- [13] Samuels, S. M. and Steele, J. M. (1981) Optimal sequential selection of a monotone sequence from a random sample. *Ann. Probab.* **9** 937–947.
- [14] Steele, J. M. (1981) Long unimodal subsequences: A problem of F. R. K. Chung. *Discrete Math.* **33** 223–225.