

THE SECRETARY PROBLEM ON AN UNKNOWN POSET

BRYN GARROD AND ROBERT MORRIS

ABSTRACT. We consider generalizations of the classical secretary problem, also known as the problem of optimal choice, to posets where the only information we have is the size of the poset and the number of maximal elements. We show that, given this information, there is an algorithm that is successful with probability at least $\frac{1}{e}$. We conjecture that if there are k maximal elements and $k \geq 2$ then this can be improved to $\sqrt[k-1]{\frac{1}{k}}$, and prove this conjecture for posets of width k . We also show that no better bound is possible.

1. INTRODUCTION

The exact origins of the classical secretary problem are complicated (and the subject of Ferguson's history of the problem [4]), but the problem was popularized by Martin Gardner in his *Scientific American* column in February 1960, as the game *googol*. The problem itself is simple to state, and its 'secretary problem' formulation is as follows. There are n candidates to be interviewed for a position as a secretary. They are interviewed one by one and, after each interview, the interviewer must decide whether or not to accept that candidate. If the candidate is accepted then the process stops, and if the candidate is rejected then the interviewer moves on to the next candidate. The interviewer may only accept the most recently interviewed candidate. At each stage, the interviewer knows the complete ranking of the candidates interviewed so far, all of whom are comparable, but has no other measure of their ability. The interviewer is only interested in finding the very best candidate; selecting any other for the job is considered a failure. It is well-known (see [4], for example) that the interviewer has a simple strategy that is successful with probability at least $\frac{1}{e}$, and that there is no strategy achieving a better bound.

Since 1960, many generalizations of the problem have been considered. One direction has been to consider partial orders on the candidates other than a total order. In this case, the interviewer knows the poset induced by the candidates interviewed so far, and wishes to choose a candidate who is maximal in the original poset. Morayne [13] considered the case of a full binary tree of depth n , and showed that the optimal strategy is to select the maximum out of the elements seen so far when the poset induced by these elements is either linear or of length greater than $\frac{n}{2}$ or non-linear with a unique maximum. He showed that as n tends to infinity, the probability of success tends to 1. Garrod, Kubicki and Morayne [6] considered the case of n pairs of 'twins', where there are n levels with two incomparable elements on each level. They showed that the optimal strategy is to wait until a certain threshold number of levels have been seen and then to select the next element that is maximal and whose twin has already been seen. They further showed that as n tends to infinity, this threshold tends to $\sim 0.4709n$ and the probability of success to ~ 0.7680 . Calculating these asymptotic values for the natural extension to ' k -tuplets', for $k > 2$, seems to be a harder problem.

This research was supported by: (BG) EPSRC; (RM) ERC Advanced grant DMMCA, and a Research Fellowship from Murray Edwards College, Cambridge.

A further interesting generalization was an attempt to find an algorithm that was successful on *all* posets of a given size with positive probability. Surprisingly, Preater [14] proved that there is such a ‘universal’ algorithm (depending only on the size of the poset), which is successful on every poset with probability at least $\frac{1}{8}$. In this algorithm, an initial random number of elements are rejected and a subsequent element is accepted according to randomized criteria. A slightly modified version of the algorithm, also suggested by Preater, was analysed by Georgiou, Kuchta, Morayne and Niemiec [7], and gave an improved lower bound of $\frac{1}{4}$ for the probability of success. More recently, Kozik [10] introduced a ‘dynamic threshold strategy’ and showed that it was successful with probability at least $\frac{1}{4} + \varepsilon$, for some $\varepsilon > 0$ and for all sufficiently large posets. Since the best possible probability of success in the classical secretary problem, on a totally ordered set, is $\frac{1}{e}$, the best possible lower bound for a universal algorithm must lie between $\frac{1}{4} + \varepsilon$ and $\frac{1}{e}$.

In this paper, we show that given any poset there is an algorithm that is successful with probability at least $\frac{1}{e}$, so, in this sense, the total order is the hardest possible partial order. In fact, this algorithm depends only on the size of the poset and its number of maximal elements, so it is universal for any family where these are given. It is therefore natural to ask which is the hardest partial order with a given number of maximal elements. The most obvious choice is the poset consisting of k disjoint chains. We shall give an asymptotically sharp lower bound on the probability of success in the problem of optimal choice on k disjoint chains, and show that it is at least as hard as on any poset with k maximal elements and of width k , that is, whose largest antichain has size k .

More precisely, our main aim is to prove the following two theorems.

Theorem 1.1. *Let (P, \prec) be a poset with k maximal elements and of width k . Then there is an algorithm for the secretary problem on (P, \prec) that is successful with probability at least p_k , where*

$$p_k = \begin{cases} \frac{1}{e} & \text{if } k = 1 \\ k^{-1} \sqrt{\frac{1}{k}} & \text{if } k > 1, \end{cases} \quad (1.1)$$

and these are the best possible such bounds.

We emphasize that in both the theorem above and that below, the claimed algorithm is not universal, but depends on both $|P|$ and the number of maximal elements of (P, \prec) .

Theorem 1.2. *Let (P, \prec) be a poset. Then there is an algorithm for the secretary problem on (P, \prec) that is successful with probability at least $\frac{1}{e}$, and this is the best possible such bound.*

We conjecture that Theorem 1.1 can be extended to all posets with k maximal elements. It is not inconceivable that the same algorithm works; if not, it would be good to find some other algorithm dependent only on k that does so.

Conjecture 1.3. *Let (P, \prec) be a poset with k maximal elements. Then there is an algorithm for the secretary problem on (P, \prec) that is successful with probability at least p_k , where p_k is as defined in (1.1).*

Our algorithm, which gives the bound in Theorem 1.2, depends only on size of the poset and the number of maximal elements. In the original version of this paper we conjectured that the latter piece of information is not needed; a beautiful proof of this result was given around the same time by Freij and Wästlund [5].

Theorem 1.4 (Freij and Wästlund [5]). *There is an universal algorithm for the secretary problem which is successful on every poset (P, \prec) with probability at least $\frac{1}{e}$.*

We remark that the secretary problem on a poset with k maximal elements was also considered recently (and independently of this work) by Kumar, Lattanzi, Vassilvitskii and Vattani [12], who obtained similar results via a different method. The poset consisting of k disjoint chains was also studied by Kuchta and Morayne [11], but with a restriction on the order in which the elements are observed: those from the first chain all appear in a random order, then those from the second chain, and so on. This poset is also related to multicriteria extensions of the secretary problem. In the original multicriteria version, each element is ranked independently in $k > 1$ different criteria, and the selector wishes to select an element that is maximal in at least one of them. This is equivalent to the problem on k equally-sized disjoint chains with the elements appearing one at a time from each chain in the same cyclic order. This version was solved by Gneden [8]. Gneden has also produced a more general survey of multicriteria problems [9]. Interestingly, the asymptotic value of the probability of success in Theorem 1.1, $\sqrt[k-1]{\frac{1}{k}}$, is the same as in the multicriteria version.

This paper is organized as follows. In Section 2, we shall introduce the formal model and some notation. In Section 3, we shall describe a (randomized) algorithm for choosing an element of our poset, and prove lower bounds for its probability of success for various families of posets. In Section 4, we shall show that our bounds are best possible, by proving that, for the poset that consists of k disjoint chains of length n (which lies in each of these families), there is no strategy that wins with probability greater than $p_k + o(1)$ (as $n \rightarrow \infty$).

2. FORMAL MODEL AND NOTATION

We begin by defining formally the probability space in which we shall work throughout the paper. The reader who wishes to avoid technicalities on a first reading is encouraged to skip this section, since all crucial definitions will be restated when used.

Our probability space will depend on a poset (P, \prec) with $P = \{x_1, \dots, x_n\}$. Let $\max_{\prec}(P)$ denote the set of its maximal elements, that is,

$$\max_{\prec}(P) = \{x \in P : \nexists y \text{ such that } x \prec y\}.$$

We shall suppress the subscript in \max_{\prec} when it is clear from the context.

Given (P, \prec) , we shall work with a probability space $(\Omega_P, \mathcal{F}_P, \mathbb{P}_P)$, with \mathbb{E}_P defined in the obvious way. We shall suppress the subscripts when they are clear from the context, as they will be for the rest of this section. We define the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ as follows. Set $\Omega = S_n \times [0, 1]$, where S_n is the permutation group on $[n]$, and $\mathcal{F} = \mathcal{P}(S_n) \times \mathcal{B}$, where \mathcal{B} is the Borel σ -algebra. Let $\mathbb{P} = \mu \times \lambda$, where μ is the uniform probability measure, that is,

$$\mu(\{\sigma\}) = \frac{1}{n!}$$

for all $\sigma \in S_n$, and λ is the Lebesgue measure. In other words, $(\sigma, \delta) \in \Omega$ is picked uniformly at random. Given $(\sigma, \delta) \in \Omega$, the σ -co-ordinate will determine the order in which elements of P appear and the δ -co-ordinate will allow us to introduce randomness independent of this order into our algorithms. Specifically, the δ -co-ordinate will determine an initial number of elements to reject without considering them. The reason why we are using continuous space and Lebesgue measure, despite the fact that all of our randomized strategies pick one of a finite number of options, is that this allows them all to lie in the same probability space.

Write $P^{[n]}$ for the set of permutations of P , and let $\pi : \Omega \rightarrow P^{[n]}$ be the random variable defined by

$$\pi(\sigma, \delta)(i) = x_{\sigma(i)}.$$

Let \mathcal{P}_t denote the set of all posets with vertex set $[t] = \{1, \dots, t\}$. Let $(P_t)_{t \in [n]}$ be a family of random variables with P_t representing the poset we see at time t . Formally, $P_t : \Omega \rightarrow \mathcal{P}_t$ and each $P_t(\sigma, \delta) = ([t], \prec_t)$ is defined by

$$\forall i, j \in [t], i \prec_t j \iff \pi(i) \prec \pi(j).$$

The poset P_t is the natural description of what we see at time t as the elements of P appear one by one.

Let $(\mathcal{F}_t)_{t \in [n]}$ be the sequence of σ -algebras with each \mathcal{F}_t generated by the random variables P_1, \dots, P_t , that is,

$$\mathcal{F}_t = \sigma(P_1, \dots, P_t) = \sigma(P_t),$$

the second equality holding since P_t is a labelled poset and so its value determines the values of P_1, \dots, P_{t-1} . We think of \mathcal{F}_t as the information we know at time t about where we are in the universe Ω . Since P_t takes only finitely-many values, \mathcal{F}_t has a simple structure; it is the pre-images in Ω of the possible values of P_t and the unions of these pre-images. We call these pre-images the atoms of \mathcal{F}_t .

Let \mathcal{F}'_t be the projection of \mathcal{F}_t onto $\mathcal{P}(S_n)$. Since our definitions have so far depended only on the σ -co-ordinate of $(\sigma, \delta) \in \Omega$, we see that, for each t ,

$$\mathcal{F}_t = \{A \times [0, 1] : A \in \mathcal{F}'_t\}.$$

In other words, (σ_1, δ_1) and (σ_2, δ_2) are in the same atom of \mathcal{F}_t if and only if σ_1 and σ_2 are in the same atom of \mathcal{F}'_t , which happens if and only if the labelled posets induced by the first t elements $\pi(1), \dots, \pi(t)$ are identical.

By a *stopping time*, we mean a random variable τ taking values in $[n]$ and satisfying the property

$$\{\tau = t\} \in \mathcal{F}_t,$$

that is, our decision to stop at time t is based only on the values of P_1, \dots, P_t .

We shall need to refer to conditional expectation and probability, which in the finite world are trivial, intuitive concepts. We define a family of random variables $(Z_t)_{t \in [n]}$ by

$$Z_t = \mathbb{P}[\pi(t) \in \max(P) \mid \mathcal{F}_t],$$

that is, Z_t is the probability that the t^{th} element observed is maximal given P_1, \dots, P_t . Our aim will be to choose a stopping time τ to maximize $\mathbb{P}[\pi(\tau) \in \max(P)]$. The value of $\mathbb{P}[\pi(\tau) \in \max(P)]$ can be easily shown to be equal to $\mathbb{E}(Z_\tau)$ – see page 45 of Chow, Robbins and Siegmund [2], for example. These equivalent formulations will be useful later.

Recall that \mathcal{F}'_t is the projection of \mathcal{F}_t onto $\mathcal{P}(S_n)$. By a *randomized stopping time*, we mean a random variable τ taking values in $[n]$ and satisfying the property

$$\{\tau = t\} \in \mathcal{F}'_t \times \mathcal{B},$$

that is, our decision to stop at time t is based on the values of P_1, \dots, P_t and on some \mathcal{B} -measurable random variable. The randomized stopping times that we shall consider will be convex combinations of a finite number of true stopping times, so if such a randomized stopping time gives a certain probability of success, then there is a true stopping time with at least that probability of success.

3. LOWER BOUNDS

Throughout this section, p is a real number satisfying $0 < p < 1$. Recall that $\pi(t)$ is the t^{th} element of the poset P that we see, and that P_t is a poset with vertex set $[t]$ that is isomorphic to the poset seen at time t . We shall prove lower bounds for the probability of success of the following randomized algorithm on different families of posets.

Algorithm. Given a poset with n elements, of which k are maximal, let $X(p) \sim \text{Bin}(n, p)$. Reject the first $X(p)$ elements and accept the first subsequent element where the following condition holds: the poset induced by the elements seen so far (including the currently observed element) has at most k maximal elements and the currently observed element is one of them.

This algorithm gives rise to the following stopping time, $\tau_k(p)$.

Let $X(p) : \Omega \rightarrow \{0, \dots, n\}$ be the random variable defined by

$$X(p)(\sigma, \delta) = \min \left\{ x \geq 0 : \sum_{i=0}^x \binom{n}{i} p^i (1-p)^{n-i} \geq \delta \right\},$$

so that

$$\mathbb{P}(X(p) = x) = \binom{n}{x} p^x (1-p)^{n-x}$$

and $X(p) = X(p)(\sigma, \delta)$ is independent of σ . Then $\tau_k(p)$ is defined by

$$\tau_k(p) = \begin{cases} \min \{ t > X(p) : |\max(P_t)| \leq k \text{ and } t \in \max(P_t) \} & \text{if this exists,} \\ n & \text{otherwise.} \end{cases}$$

Given the definition of $\tau_k(p)$, it makes sense to consider another random variable, the set of $X(p)$ elements that we reject without considering. We denote this random variable by $S(p)$, where

$$S(p) = \{\pi(t) : t \leq X(p)\}.$$

We shall make use of the following simple property of $S(p)$, which is easily verified.

Lemma 3.1. *The events $\{x \in S(p)\}_{x \in P}$ are independent and $\mathbb{P}(x \in S(p)) = p$ for all $x \in P$.*

Proof. We can generate π and $X(p)$ with the required distributions in the following way. Put each element of P in $S(p)$ with probability p independently of all other elements. Let π consist of a uniformly random ordering of the elements of $S(p)$ followed by a uniformly random ordering of $P \setminus S(p)$. By symmetry, π is a uniformly random ordering of P , and $X(p) = |S(p)|$ is a binomial random variable independent of π . The events $\{x \in S(p)\}_{x \in P}$ depend only on π and $X(p)$, and by construction the properties in the statement of the lemma hold. \square

We shall also use the following standard identity; for completeness we include a proof.

Lemma 3.2. *For all integers $k \geq 1$, the following holds:*

$$\sum_{s=0}^{\infty} \binom{k+s-1}{k-1} (1-p)^s = \frac{1}{p^k}.$$

Proof. Suppose that we have a coin that comes up heads with probability p and tails with probability $1-p$, and that we toss it infinitely many times. Then, with probability 1, we shall see at least k heads, and the k^{th} head comes up in position $k+s$ for some $s \geq 0$. In this case,

then $\pi(\tau_k(p))$ will be the first element observed from the $j_1 + \dots + j_k$ elements not in $S(p)$ that are at the tops of their chains.

The events $\{A_{j_1, \dots, j_k} : 0 \leq j_1 \leq m_1, \dots, 0 \leq j_k \leq m_k\}$ partition the whole space. Thus, writing $Q_k(p)$ for $\mathbb{P}[\pi(\tau_k(p)) \in \max(P)]$,

$$\begin{aligned} Q_k(p) &= \sum_{0 \leq j_1 \leq m_1, \dots, 0 \leq j_k \leq m_k} \mathbb{P}[\pi(\tau_k(p)) \in \max(P) \mid A_{j_1, \dots, j_k}] \cdot \mathbb{P}[A_{j_1, \dots, j_k}] \\ &> \sum_{\substack{0 \leq j_1 \leq m_1, \dots, 0 \leq j_k \leq m_k \\ (j_1, \dots, j_k) \neq (0, \dots, 0)}} \frac{|\{i : j_i > 0\}|}{j_1 + \dots + j_k} (1-p)^{j_1 + \dots + j_k} p^{|\{i : j_i < m_i\}|}. \end{aligned} \quad (3.1)$$

Since $1 + (1-p) + (1-p)^2 + \dots = \frac{1}{p}$, this can be written as

$$\begin{aligned} Q_k(p) &> \sum_{\substack{0 \leq j_1 \leq m_1, \dots, 0 \leq j_k \leq m_k \\ (j_1, \dots, j_k) \neq (0, \dots, 0)}} \frac{|\{i : j_i > 0\}|}{j_1 + \dots + j_k} (1-p)^{j_1 + \dots + j_k} p^k (1 + (1-p) + (1-p)^2 + \dots)^{|\{i : j_i = m_i\}|} \\ &= \sum_{\substack{j_1, \dots, j_k \geq 0 \\ (j_1, \dots, j_k) \neq (0, \dots, 0)}} \frac{|\{i : j_i > 0\}|}{\min\{j_1, m_1\} + \dots + \min\{j_k, m_k\}} (1-p)^{j_1 + \dots + j_k} p^k \\ &> \sum_{\substack{j_1, \dots, j_k \geq 0 \\ (j_1, \dots, j_k) \neq (0, \dots, 0)}} \frac{|\{i : j_i > 0\}|}{j_1 + \dots + j_k} (1-p)^{j_1 + \dots + j_k} p^k. \end{aligned} \quad (3.2)$$

To see the equality, simply note that the term corresponding to (j_1, \dots, j_k) on the right-hand side appears on the left-hand side by choosing the term $(1-p)^{j_i - m_i}$ in the sum whenever $j_i \geq m_i$.

We now rewrite (3.2) as a sum over $r = |\{i : j_i > 0\}|$ and $s = j_1 + \dots + j_k$, and obtain

$$Q_k(p) > \sum_{r=1}^k \sum_{s=r}^{\infty} \left| \left\{ (j_1, \dots, j_k) : |\{i : j_i > 0\}| = r \text{ and } j_1 + \dots + j_k = s \right\} \right| \cdot \frac{r}{s} (1-p)^s p^k.$$

The rest of the proof is a straightforward calculation. To calculate $|\{(j_1, \dots, j_k) : |\{i : j_i > 0\}| = r \text{ and } j_1 + \dots + j_k = s\}|$, we note that there are $\binom{k}{r}$ ways of choosing the indices i with $j_i > 0$ and there are then $\binom{s-1}{r-1}$ ways for r non-zero numbers to add up to s . Thus

$$Q_k(p) > \sum_{r=1}^k \sum_{s=r}^{\infty} \binom{k}{r} \binom{s-1}{r-1} \frac{r}{s} (1-p)^s p^k = kp^k \sum_{r=1}^k \sum_{s=r}^{\infty} \binom{k-1}{r-1} \binom{s-1}{r-1} \frac{1}{s} (1-p)^s.$$

Reversing the order of summation,

$$Q_k(p) > kp^k \sum_{s=1}^{\infty} \frac{1}{s} (1-p)^s \sum_{r=1}^{\min\{k, s\}} \binom{s-1}{r-1} \binom{k-1}{k-r}.$$

The second sum is easily evaluated (a result known as Vandermonde's identity) to give

$$Q_k(p) > kp^k \sum_{s=1}^{\infty} \frac{1}{s} (1-p)^s \binom{k+s-2}{k-1}. \quad (3.3)$$

Theorem 3.4. *Let (P, \prec) be a poset with n elements. Suppose that (P, \prec) has k maximal elements and that none of its antichains has size greater than k . Then*

$$\mathbb{P}[\pi(\tau_k(p)) \in \max(P)] > \begin{cases} p \log \frac{1}{p} & \text{if } k = 1, \\ \frac{k}{k-1} p(1 - p^{k-1}) & \text{if } k > 1. \end{cases}$$

Proof. By Dilworth's theorem, we see that P takes the form of k chains with some comparisons in between them. Clearly, the k elements of $\max(P)$ lie at the top of the k chains. The proof therefore proceeds in an almost identical manner to that of Theorem 3.3. The only difference is that the denominator in each term of (3.1) is now at most, rather than equal to, $j_1 + \dots + j_k$ (see Figure 2), so the expression in this line is still a lower bound. The calculations that make up the remainder of the proof of Theorem 3.3 therefore follow in the same way. \square

The values that maximize the function in Theorem 3.4 are

$$p_k = \begin{cases} \frac{1}{e} & \text{if } k = 1, \\ k^{-1} \sqrt[k]{\frac{1}{k}} & \text{if } k > 1. \end{cases} \quad (3.4)$$

This gives us the following corollary and the lower bounds in Theorem 1.1.

Corollary 3.5. *Let (P, \prec) be a poset with n elements. Suppose that (P, \prec) has k maximal elements and that none of its antichains has size greater than k . Then*

$$\mathbb{P}[\pi(\tau_k(p_k)) \in \max(P)] > p_k.$$

It is interesting to note that the expected proportion of elements that we reject without considering is the same as the probability of success.

We now wish to prove the following theorem, which, with the right choice of p , will give us a lower bound of $\frac{1}{e}$ for all posets, as in Theorem 1.2.

Theorem 3.6. *Let (P, \prec) be a poset with n elements. Suppose that (P, \prec) has k maximal elements. Then*

$$\mathbb{P}[\pi(\tau_k(p)) \in \max(P)] > kp^k \log \frac{1}{p}.$$

The proof will use two simple lemmas. The first states that the linear order is the hardest of all posets with a unique maximal element.

Lemma 3.7. *Let (P, \prec) be a poset with n elements. Suppose that (P, \prec) has exactly one maximal element. Then the probability that $\tau_1(p)$ is successful on (P, \prec) is at least the probability that it is successful on a linear ordering of P , and hence*

$$\mathbb{P}[\pi(\tau_1(p)) \in \max(P)] > p \log \frac{1}{p}.$$

Proof. We begin by taking an arbitrary linear extension of \prec , that is, a partial order \prec' such that any two elements are comparable and such that $x \prec y \Rightarrow x \prec' y$. (It is clear that such a partial order exists.) We denote the unique element in $\max_{\prec}(P) = \max_{\prec'}(P)$ by x_{\max} .

Given this new poset, (P, \prec') , we define random variables π' , $X'(p)$, $S'(p)$ and $\tau'_1(p)$ in the same way as π , $X(p)$, $S(p)$ and $\tau_1(p)$ were defined given (P, \prec) . We couple the random variables $(\pi, X(p), S(p), \tau_1(p))$ and $(\pi', X'(p), S'(p), \tau'_1(p))$ in the obvious way; we set $\pi' = \pi$ and $X'(p) = X(p)$, and hence $S'(p) = S(p)$. This means that the elements appear in the same order in both instances, and the same set $S(p)$ is rejected in both cases. The induced posets observed in the process on (P, \prec') are linear extensions of those observed in the process on (P, \prec) . We show

that if $\pi(\tau_1(p)) \neq x_{\max}$ then $\pi'(\tau'_1(p)) \neq x_{\max}$, that is, if $\tau_1(p)$ fails in the process on (P, \prec) then $\tau'_1(p)$ fails on (P, \prec') . From this, the result follows, since the probability of success is therefore at least as large on (P, \prec) as on (P, \prec') , and Theorem 3.3 applied to (P, \prec') gives the lower bound.

If we reach x_{\max} then it will be accepted, since it must be the unique maximal element in the poset induced by the elements observed so far. Thus $\pi(\tau_1(p)) \neq x_{\max}$ if either

- (i) $x_{\max} \in S(p)$ or
- (ii) after rejecting $S(p)$, we accept an element that appears earlier than x_{\max} .

In case (i), $\tau'_1(p)$ must fail on (P, \prec') for the same reason, since $S'(p) = S(p)$. In case (ii), such an element must be the unique maximal element of the poset induced by what we have seen so far, and this is still the case in any linear extension. Therefore, with $\tau'_1(p)$, if this element is observed then it must be accepted, and so we still accept an element that appears earlier than x_{\max} . It follows that in either case $\pi'(\tau'_1(p)) \neq x_{\max}$. \square

The next lemma gives a lower bound for the probability of success restricting our attention to the case when all but one of the maximal elements of our poset are in $S(p)$. This turns out to be enough to prove Theorem 3.6.

Lemma 3.8. *Let (P, \prec) be a poset with n elements. Suppose that (P, \prec) has k maximal elements. Then*

$$\mathbb{P}[\pi(\tau_k(p)) \in \max(P) \mid |\max(P) \cap S(p)| = k - 1] > \frac{p}{1-p} \log \frac{1}{p}.$$

Proof. We first observe that we may assume that $k = 1$, for the following reason. The condition that $|\max(P) \cap S(p)| = k - 1$ means that the $k - 1$ maximal elements in $\max(P) \cap S(p)$ will be maximal for the remainder of the process, so when using $\tau_k(p)$ we may ignore these and all elements dominated by at least one of these, and wait for a unique maximal element from the remaining elements. Those elements form a poset (P', \prec') with a unique maximal element x_{\max} , which is not in $S(p)$. Since all elements are in $S(p)$ with probability p independently of the others, the situation is the same as if we were working with (P', \prec') and conditioning on $x_{\max} \notin S(p)$. Looking for one of at most k maximal elements in P using $\tau_k(p)$ is the same as looking for a unique maximal element in P' using $\tau_1(p)$.

We assume from now on that $k = 1$; we shall use Lemma 3.7 to prove the result in this case. Lemma 3.7 used the bound from Theorem 3.3, and we recall from the proof of that theorem that the lower bound for $\mathbb{P}[\pi(\tau_1(p)) = x_{\max}]$ is in fact a lower bound for $\mathbb{P}[(S(p) \neq P) \wedge (\pi(\tau_1(p)) = x_{\max})]$.

Let us write M for the event that $x_{\max} \in S(p)$ and W for the event $(S(p) \neq P) \wedge (\pi(\tau_1(p)) = x_{\max})$. We note that if $S(p) \neq P$ and $x_{\max} \in S(p)$ then $\pi(\tau_1(p)) \neq x_{\max}$ and hence $W = W \wedge M^c$. Therefore,

$$\mathbb{P}(W) = \mathbb{P}(W \wedge M^c) = \mathbb{P}(M^c) \mathbb{P}(W \mid M^c) = (1-p) \mathbb{P}(W \mid M^c).$$

Since, by the bound from Lemma 3.7,

$$\mathbb{P}(W) > p \log \frac{1}{p},$$

and the quantity that we are interested in is $\mathbb{P}(W \mid M^c)$, the result follows. \square

We are now in a position to prove our theorem.

Proof of Theorem 3.6. We have

$$\mathbb{P}[|\max(P) \cap S(p)| = k - 1] = kp^{k-1}(1-p).$$

Thus, for general k ,

$$\begin{aligned} \mathbb{P}[\pi(\tau_k(p)) \in \max(P)] &> \mathbb{P}[\pi(\tau_k(p)) \in \max(P) \mid |\max(P) \cap S(p)| = k-1] \\ &\quad \times \mathbb{P}[|\max(P) \cap S(p)| = k-1] \\ &> \frac{p}{1-p} \log \frac{1}{p} \cdot kp^{k-1}(1-p) = kp^k \log \frac{1}{p}, \end{aligned}$$

as required. \square

This gives us the following corollary. The probability $e^{-\frac{1}{k}}$ is chosen to maximize the function in Theorem 3.6 and gives the lower bound in Theorem 1.2. As mentioned earlier, it is well-known that $\frac{1}{e}$ is the best possible lower bound for the probability of success in the classical secretary problem, and so this completes Theorem 1.2.

Corollary 3.9. *Let (P, \prec) be a poset with n elements of which k are maximal. Then*

$$\mathbb{P}\left[\pi\left(\tau_k\left(e^{-\frac{1}{k}}\right)\right) \in \max(P)\right] > \frac{1}{e}.$$

4. UPPER BOUND

In this section, we show that the bound in Corollary 3.5 is best possible. The proof of Theorem 3.3 shows that the probability of success of the stopping time $\tau_k(p_k)$ on k disjoint chains decreases towards the given lower bound as the chains increase in length. This might suggest that the probability of success of an optimal strategy is reduced as the chains increase in length and thus, to prove that the bounds are best possible, we would consider chains with length tending to infinity. The main theorem in this section, Theorem 4.1, does just that; for sufficiently long chains, the probability of success of an optimal stopping time can be made arbitrarily close to that in Corollary 3.5, and so $\tau_k(p_k)$ is asymptotically optimal. Since the poset consisting of k disjoint chains satisfies the conditions of Corollary 3.5, the bounds given are the best possible such bounds.

We define $D_k(x)$ to be the poset consisting of k disjoint chains, each of size x . It might be useful at this point to recall some definitions from Section 2. The probability space associated with the poset $D_k(x)$ is denoted by $(\Omega_{D_k(x)}, \mathcal{F}_{D_k(x)}, \mathbb{P}_{D_k(x)})$, but we suppress the subscripts when they are clear from the context. The poset induced by the first t elements that we observe is isomorphic to the random variable P_t , which is a poset on vertex set $[t]$, and \mathcal{F}_t is the σ -algebra generated by P_1, \dots, P_t , which represents what we know at time t . A stopping time is a random variable τ taking values in $[n]$ and satisfying the property

$$\{\tau = t\} \in \mathcal{F}_t.$$

We shall use the notation $\mathcal{C}_{(\mathcal{F}_t)}$ to denote the class of all such stopping times, and extend this notation to any sequence of σ -algebras in the analogous way.

We are trying to find an upper bound for $\mathbb{E}(Z_\tau)$ that holds for all stopping times τ , where

$$Z_t = \mathbb{P}[\pi(t) \in \max(D_k(x)) \mid \mathcal{F}_t].$$

Theorem 4.1 states that, as $x \rightarrow \infty$, the limit of the probability of success of the optimal stopping time on $D_k(x)$ is no greater than p_k . Since Corollary 3.5 showed the existence of a stopping time with probability of success at least p_k , Theorem 4.1 shows that this is the best possible such bound and so gives Theorem 1.1.

Theorem 4.1. *Let p_k be as defined in (1.1). Then*

$$\lim_{x \rightarrow \infty} \sup_{\tau \in \mathcal{C}(\mathcal{F}_t)} \mathbb{E}_{D_k(x)}(Z_\tau) \leq p_k.$$

Note that the supremum is over stopping times in $\mathcal{C}(\mathcal{F}_t)$, which means that we are allowed to use the extra information from the structure of the posets, not just the pay-offs that we are offered.

The following observation is important, so we record it as a lemma.

Lemma 4.2. *When $(P, \prec) = D_k(x)$, we have*

$$Z_t = \begin{cases} y/x & \text{if } \pi(t) \in \max(P_t), \pi(t) \in C \text{ and } |C \cap \{\pi(1), \dots, \pi(t)\}| = y, \\ 0 & \text{if } \pi(t) \notin \max(P_t), \end{cases}$$

where C is one of the k chains in $D_k(x)$.

Proof. The maximal element of a chain C is equally likely to be at any position in the order in which its x elements are observed. Therefore, when y elements have been observed from this chain, the probability that one of them is the maximal element is $\frac{y}{x}$, independent of the most recently observed element being maximal. \square

The proof of Theorem 4.1 will proceed roughly as follows. At time t we expect to have seen approximately $\frac{t}{k}$ elements from each chain. Therefore, since all orders are equally likely, the t^{th} element that we observe is maximal in what we have seen so far with probability approximately $\frac{1}{t/k} = \frac{k}{t}$. By Lemma 4.2, if this happens then it is a maximal element of $D_k(x)$ with probability approximately $\frac{t}{kx}$. We conclude that Z_t is approximately distributed as

$$Z_t = \begin{cases} \frac{t}{kx} & \text{with probability } \frac{k}{t}, \\ 0 & \text{with probability } 1 - \frac{k}{t}. \end{cases} \quad (4.1)$$

If Z_t were distributed exactly like this with the Z_t all independent of each other, then the proof would not be difficult to complete. Since the potential non-zero value of Z_t increases with t , it is straightforward to show (as in the classical secretary problem; more details will be given later in this section) that the optimal strategy is to ignore the first I elements and accept the next non-zero Z_t . Let us denote the associated stopping time by τ_I and make some rough calculations. This is only an outline of the more precise argument that will be given later; in particular, \approx is only intended to have an intuitive meaning and does not stand for any well-defined relation. We find that

$$\begin{aligned} \mathbb{E}(Z_{\tau_I}) &= \sum_{t=I+1}^{kx} \mathbb{P}(\tau_I = t) \mathbb{E}(Z_t | \tau_I = t) \\ &\approx \sum_{t=I+1}^{kx} \mathbb{P}\left[(Z_i = 0 \ \forall i \in \{I+1, \dots, t-1\}) \wedge (Z_t > 0)\right] \cdot \frac{t}{kx} \\ &\approx \sum_{t=I+1}^{kx} \left(\prod_{i=I+1}^{t-1} \left(1 - \frac{k}{i}\right) \right) \cdot \frac{k}{t} \cdot \frac{t}{kx} \\ &= \frac{1}{x} \sum_{t=I+1}^{kx} \prod_{i=I+1}^{t-1} \left(1 - \frac{k}{i}\right). \end{aligned}$$

We shall apply this formula in the case where k is much smaller than I , so we can approximate $1 - \frac{k}{i}$ by $e^{-\frac{k}{i}}$, and twice approximate sums by integrals to obtain

$$\begin{aligned} \mathbb{E}(Z_{\tau_I}) &\approx \frac{1}{x} \sum_{t=I+1}^{kx} e^{-\sum_{i=I+1}^{t-1} \frac{k}{i}} \approx \frac{1}{x} \sum_{t=I+1}^{kx} e^{-k \log(\frac{t}{I})} \\ &= \frac{I^k}{x} \sum_{t=I+1}^{kx} t^{-k} \approx \begin{cases} \frac{I}{x} \cdot \log\left(\frac{x}{I}\right) & \text{if } k = 1, \\ \frac{k}{k-1} \cdot \frac{I}{kx} \cdot \left(1 - \left(\frac{I}{kx}\right)^{k-1}\right) & \text{if } k > 1. \end{cases} \end{aligned}$$

This is the formula in Theorem 3.4 with $p = \frac{I}{kx}$ and is thus maximized, as in Corollary 3.5, when $\frac{I}{kx} = p_k$, in which case $\mathbb{E}(Z_{\tau_I}) \approx p_k$. Therefore the bounds in Corollary 3.5 are best possible, and if these calculations had been exact then the proof of Theorem 1.1 would be complete.

Unfortunately, Z_t is not distributed exactly as in (4.1). In order to conclude that the optimal stopping time is of the simple form above, we should like to use the principle of *backward induction*, described later (see also [2], Theorem 3.2). This formalizes the intuitive principle that, in a finite game, the optimal strategy is simply to analyse at each step whether or not we expect our situation to improve by continuing, and to do so if and only if this is the case.

The reason why the sequence of random variables $(Z_t)_{t \in [n]}$ is difficult to analyse is that the values they can take vary depending on how the process unfolds. However, it is very likely that at any time we shall have seen approximately the same number of elements from each chain. The proof will therefore proceed by defining a sequence of random variables $(Y_t)_{t \in [n]}$, which act as asymptotically almost sure upper bounds for Z_t and are easier to analyse. To obtain these bounds, we shall split each chain into m segments, each of length ℓ , and split the process into m sets of $k\ell$ observations. These lengths ℓ are margins of error beyond which we do not expect the number of elements observed from a chain to deviate. Initially, we shall fix m and let $\ell \rightarrow \infty$ to find an upper bound for $\mathbb{E}(Y_\tau)$ and hence $\mathbb{E}(Z_\tau)$ in terms of m . Letting $m \rightarrow \infty$ will then give us a best possible result.

This means that the precise statement we shall prove for Theorem 4.1 is in fact

$$\lim_{m \rightarrow \infty} \lim_{\ell \rightarrow \infty} \sup_{\tau \in \mathcal{C}(\mathcal{F}_t)} \mathbb{E}_{D_k(\ell m)}(Z_\tau) \leq p_k.$$

However, this is purely a matter of convenience; it is clear that the proof can be extended to posets $D_k(x)$ where x is not a multiple of m by dividing each chain into m almost equal rather than exactly equal segments.

We need to show that the process behaves in this approximately uniform manner with high probability as $\ell \rightarrow \infty$. We shall first define what it means to be approximately uniform in one particular chain C at time $k\ell s$, an event we call $U_{C,s}$ (see Figure 3), and then what it means to be approximately uniform everywhere at all times, an event we call U .

Given one of the chains, C , and for all $s \in \{0, \dots, m\}$, let $U_{C,s}$ be the event that when we have observed $k\ell s$ elements in total we have observed between $\ell(s-1)$ and $\ell(s+1)$ elements from C , that is,

$$U_{C,s} = \{\ell(s-1) \leq |C \cap \{\pi(1), \dots, \pi(k\ell s)\}| \leq \ell(s+1)\}.$$

For all $t \in \{0, \dots, k\ell m\}$, let $s(t)$ be the unique integer s such that $k\ell(s-1) < t \leq k\ell s$, that is,

$$s(t) = \left\lceil \frac{t}{k\ell} \right\rceil. \quad (4.2)$$

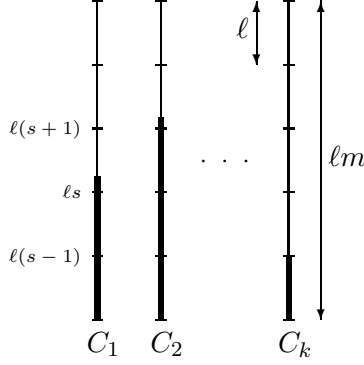


FIGURE 3. This figure shows the number of elements observed from each chain. In this example, after a total of $k\ell s$ elements have been observed we see that $U_{C_1,s}$ and $U_{C_k,s}$ hold but $U_{C_2,s}$ does not.

Let U be the event that for all t when we have observed t elements in total we have observed between $\ell(s(t) - 2)$ and $\ell(s(t) + 1)$ elements from each chain, that is,

$$U = \bigcap_{i,t} \{ \ell(s(t) - 2) \leq |C_i \cap \{\pi(1), \dots, \pi(t)\}| \leq \ell(s(t) + 1) \}.$$

We shall use Lemma 4.4, which follows easily from Lemma 4.3. It states that the process is approximately uniform with high probability.

Lemma 4.3. *Let $m \geq 1$ be an integer, let C be one of the chains in $D_k(\ell m)$ and let $s \in \{0, \dots, m\}$. Then*

$$\lim_{\ell \rightarrow \infty} \mathbb{P}_{D_k(\ell m)}(U_{C,s}) = 1.$$

Proof. We show that the probability that we have observed more than $\ell(s + 1)$ or fewer than $\ell(s - 1)$ elements tends to zero as $\ell \rightarrow \infty$. (If $s = 0$ or $s = m$ then we need consider only one of these tails.)

Assume C and s are given. Let N be the number of elements we have observed from chain C when we have observed $k\ell s$ elements in total. It is straightforward to check that

$$\mathbb{P}(N = x) = \frac{\binom{\ell m}{x} \binom{(k-1)\ell m}{k\ell s - x}}{\binom{k\ell m}{k\ell s}}$$

is increasing for $x < \ell s$ and decreasing for $x > \ell s$, and that

$$\left| \frac{\mathbb{P}(N = x + 1)}{\mathbb{P}(N = x)} - 1 \right| > c > 0$$

if $x \notin [\ell(s - 1), \ell(s + 1)]$, for some absolute $c > 0$. It follows that $\mathbb{P}(N \notin [\ell(s - 1), \ell(s + 1)]) = O\left(\frac{1}{\ell}\right) \rightarrow 0$ as $\ell \rightarrow \infty$. \square

Lemma 4.4. *Let $m \geq 1$ be an integer. Then*

$$\lim_{\ell \rightarrow \infty} \mathbb{P}_{D_k(\ell m)}(U) = 1.$$

Proof. This lemma follows simply from the previous lemma: choose ℓ sufficiently large that each of the $k(m+1)$ events $U_{C_{i,s}}$ occurs with probability at least $1 - \delta$. Then, trivially, all $k(m+1)$ events hold with probability at least $1 - k(m+1)\delta$. It is easy to see that

$$\bigcap_{C_{i,s}} U_{C_{i,s}} \subseteq U,$$

since the events $U_{C_{i,s(t)-1}}$ and $U_{C_{i,s(t)}}$ imply that

$$\ell(s(t) - 2) \leq |C_i \cap \{\pi(1), \dots, \pi(t)\}| \leq \ell(s(t) + 1),$$

and so this holds for every t and i , as required. \square

The next lemma states that in order to prove Theorem 4.1, it suffices to show that its statement is true if we condition on U holding. This formalizes the intuition that, since the process is asymptotically almost surely uniform (that is, since $\lim_{\ell \rightarrow \infty} \mathbb{P}_{D_k(\ell m)}(U) \rightarrow 1$), we may assume that it is uniform.

Recall that $\mathcal{C}_{(\mathcal{F}_t)}$ is the class of all stopping times relative to the σ -algebras $\mathcal{F}_t = \sigma(P_1, \dots, P_t)$, that is, the decision to stop at time t depends only on P_1, \dots, P_t .

Lemma 4.5. *For all m ,*

$$\lim_{\ell \rightarrow \infty} \sup_{\tau \in \mathcal{C}_{(\mathcal{F}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau) \leq \lim_{\ell \rightarrow \infty} \sup_{\tau \in \mathcal{C}_{(\mathcal{F}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U).$$

Proof. By Lemma 4.4, for all $\varepsilon > 0$ we may choose ℓ sufficiently large that $\mathbb{P}_{D_k(\ell m)}(U^c) \leq \varepsilon$. We also know that $Z_t \leq 1$ for all t . Therefore, for all τ ,

$$\begin{aligned} \mathbb{E}_{D_k(\ell m)}(Z_\tau) &= \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) \mathbb{P}_{D_k(\ell m)}(U) + \mathbb{E}_{D_k(\ell m)}(Z_\tau | U^c) \mathbb{P}_{D_k(\ell m)}(U^c) \\ &\leq \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) + \varepsilon. \end{aligned}$$

We now take suprema to obtain

$$\sup_{\tau \in \mathcal{C}_{(\mathcal{F}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau) \leq \sup_{\tau \in \mathcal{C}_{(\mathcal{F}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) + \varepsilon.$$

Since ε is arbitrary, the result follows. \square

Next, we define the random variables Y_t that act as upper bounds for the Z_t and are easier to analyse. These random variables are not strict upper bounds, in the sense that the random variables are not coupled in any way. However, conditioned on U occurring, Z_t is less than the potential non-zero value of Y_t and the probability that Z_t is non-zero is less than the probability that Y_t is non-zero, and we shall be able to show that the optimal strategy for the game on Z_t has a lower expected pay-off than the optimal strategy for the game on Y_t .

We shall often need to refer to the potential non-zero value of Y_t and the probability that it takes this value, so let

$$y_t = \frac{s(t) + 1}{m} \quad \text{and} \quad \tilde{p}_t = \begin{cases} \frac{1}{\ell(s(t)-2)} & \text{if } s(t) \geq 3, \\ 1 & \text{if } s(t) \leq 2, \end{cases} \quad (4.3)$$

where $s(t)$ is as defined in (4.2). We now define a sequence of independent random variables $(Y_t)_{t \in [n]}$ by

$$Y_t = \begin{cases} y_t & \text{with probability } \tilde{p}_t, \\ 0 & \text{with probability } 1 - \tilde{p}_t. \end{cases}$$

We shall not define these explicitly on any probability space as there is no need to do so, although it is of course straightforward to do it on Ω .

The next lemma states the intuitive principle that we expect to do at least as well in the game with the random variables Y_t as in the game with the random variables Z_t conditioned on U occurring. Analogously to \mathcal{F}_t for Z_t , let $(\mathcal{G}_t)_{t \in [n]}$ be defined by $\mathcal{G}_t = \sigma(Y_1, \dots, Y_t)$.

Lemma 4.6. *For all $\ell, m \in \mathbb{N}$ with $m \geq 3$,*

$$\sup_{\tau \in \mathcal{C}(\mathcal{F}_t)} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) < \sup_{\tau \in \mathcal{C}(\mathcal{G}_t)} \mathbb{E}_{D_k(\ell m)}(Y_\tau).$$

Putting Lemmas 4.5 and 4.6 together tells us that

$$\lim_{\ell \rightarrow \infty} \sup_{\tau \in \mathcal{C}(\mathcal{F}_t)} \mathbb{E}_{D_k(\ell m)}(Z_\tau) \leq \lim_{\ell \rightarrow \infty} \sup_{\tau \in \mathcal{C}(\mathcal{G}_t)} \mathbb{E}_{D_k(\ell m)}(Y_\tau).$$

In order to prove this lemma, we need a precise statement of what backward induction tells us is the optimal stopping time in a finite process. We define a new random variable for each t , the *value* of the game at time t . This is the expected pay-off ultimately accepted given what has happened so far. We calculate these values inductively, starting at the end. The value of the game at the final step is just the final pay-off offered. The value of the game at each earlier step is the maximum of the currently offered pay-off and the expected value of the game at the next step. The optimal strategy is to stop when the currently offered pay-off is at least the expected value at the next step.

In the backward induction theorem below, the pay-offs offered are the W_t and the values at each step are the γ_t . The σ -algebras \mathcal{A}_t represent what we know at time t . We remind the reader that being \mathcal{A}_t -measurable means that $\sigma(W_t) \subset \mathcal{A}_t$, that is, the value of W_t is determined by what we know at time t or, in the finite world, W_t is constant on each atom of \mathcal{A}_t . In fact, the nested condition means that $\mathcal{A}_t \supset \sigma(W_1, \dots, W_t)$. The statement of the theorem is that the strategy that stops at the first t when $W_t = \gamma_t$ (or, equivalently, when W_t is at least as large as the expected value of γ_{t+1} given \mathcal{A}_t) is indeed a stopping time and achieves the optimal value. For more details, see [2, Theorem 3.2].

Backward induction. *Let $\mathcal{A}_1 \subset \dots \subset \mathcal{A}_n$ be a nested sequence of σ -algebras and let W_1, \dots, W_n be a sequence of random variables with each W_t being \mathcal{A}_t -measurable. Let $\mathcal{C}_{(\mathcal{A}_t)}$ be the class of stopping times relative to $(\mathcal{A}_t)_{t \in [n]}$ and let v^* be given by*

$$v^* = \sup_{\tau \in \mathcal{C}_{(\mathcal{A}_t)}} \mathbb{E}(W_\tau).$$

Define successively $\gamma_n, \gamma_{n-1}, \dots, \gamma_1$ by setting

$$\begin{aligned} \gamma_n &= W_n, \\ \gamma_t &= \max \{W_t, \mathbb{E}(\gamma_{t+1} | \mathcal{A}_t)\}, \quad t = n-1, \dots, 1. \end{aligned}$$

Let

$$\tau^* = \min\{t : W_t = \gamma_t\}.$$

Then $\tau^ \in \mathcal{C}_{(\mathcal{A}_t)}$ and*

$$\mathbb{E}(W_{\tau^*}) = \mathbb{E}(\gamma_1) = v^* \geq \mathbb{E}(W_\tau) \text{ for all } \tau \in \mathcal{C}_{(\mathcal{A}_t)}.$$

This theorem provides the machinery we need to prove our lemma.

Proof of Lemma 4.6. We shall apply backward induction to the sequences Y_t and Z_t , conditioned on U occurring, to show that the optimal expected pay-off for the Y_t is at least as large as for the Z_t . For convenience, we shall continue to use n to denote the number of elements in $D_k(\ell m)$, that is, $n = k\ell m$.

Let us first consider what happens with the sequence Y_t . Recall that $(\mathcal{G}_t)_{t \in [n]}$ are defined by $\mathcal{G}_t = \sigma(Y_1, \dots, Y_t)$ and let $(\alpha_t)_{t \in [n]}$ be defined for $(Y_t)_{t \in [n]}$ as $(\gamma_t)_{t \in [n]}$ were for $(W_t)_{t \in [n]}$ in the backward induction theorem, that is,

$$\begin{aligned}\alpha_n &= Y_n, \\ \alpha_t &= \max \{Y_t, \mathbb{E}(\alpha_{t+1} | \mathcal{G}_t)\}, \quad t = n-1, \dots, 1.\end{aligned}$$

Since the random variables Y_t are independent, the values of Y_1, \dots, Y_t give no information about the values of Y_{t+1}, \dots, Y_n , and therefore $\mathbb{E}(\alpha_{t+1} | \mathcal{G}_t)$ is constant on \mathcal{G}_t and equal to $\mathbb{E}(\alpha_{t+1})$. Therefore we may define the function $v : [n] \rightarrow \mathbb{R}$ by

$$v(t) = \mathbb{E}(\alpha_t)$$

and note that backward induction tells us that the stopping time that stops at the first t such that $Y_t \geq \mathbb{E}(\alpha_{t+1} | \mathcal{G}_t) = v(t+1)$ is optimal. By definition,

$$v(t) = \mathbb{E}(\max\{Y_t, v(t+1)\}) \geq v(t+1),$$

whereas y_t , the potential non-zero value of Y_t , is a non-decreasing function of t . We conclude that there exists I such that

$$\begin{aligned}y_t &< v(t+1) && \text{if } t \leq I, \\ y_t &\geq v(t+1) && \text{if } t > I,\end{aligned}\tag{4.4}$$

and therefore an optimal strategy for the game on Y_t is ‘reject the first I elements, and accept the next with a non-zero pay-off.’

Recall that the distribution of Y_t is given by

$$Y_t = \begin{cases} y_t & \text{with probability } \tilde{p}_t, \\ 0 & \text{with probability } 1 - \tilde{p}_t. \end{cases}$$

We deduce that, since $m \geq 3$ and $s(n) = m$,

$$v(n) = \mathbb{E}(\alpha_n) = \mathbb{E}(Y_n) = \tilde{p}_n y_n = \frac{m+1}{m} \cdot \frac{1}{\ell(m-2)} > \frac{1}{\ell m} = \frac{k}{n}\tag{4.5}$$

and that, for $1 \leq t \leq n-1$,

$$v(t) = \mathbb{E}(\alpha_t) = \mathbb{E}(\max\{Y_t, \mathbb{E}(\alpha_{t+1})\}) = \begin{cases} \tilde{p}_t y_t + (1 - \tilde{p}_t)v(t+1) & \text{if } t > I, \\ v(t+1) & \text{if } t \leq I. \end{cases}\tag{4.6}$$

We now turn our attention to the sequence of random variables Z_t . Since we are conditioning on U , we introduce a new sequence of σ -algebras \mathcal{H}_t defined by

$$\mathcal{H}_t = \sigma(\mathcal{F}_t \cup \{U\})$$

and shall consider only $\omega \in U$. Analogously to γ_t and α_t , let

$$\beta_n = Z_n$$

and, for $1 \leq t \leq n-1$, let

$$\beta_t = \max \{Z_t, \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)\}.$$

Recalling that $\mathcal{C}_{(\mathcal{F}_t)}$ denotes the class of stopping times relative to \mathcal{F}_t , observe that we have $\mathcal{F}_t \subset \mathcal{H}_t$, and hence $\mathcal{C}_{(\mathcal{F}_t)} \subset \mathcal{C}_{(\mathcal{H}_t)}$ and

$$\sup_{\tau \in \mathcal{C}_{(\mathcal{F}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) \leq \sup_{\tau \in \mathcal{C}_{(\mathcal{H}_t)}} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U).$$

Note that intuitively this is obvious: it just says that having extra information (that the event U occurs) can only help us in choosing our stopping time τ .

Recall that $\mathbb{E}(\beta_t | \mathcal{H}_{t-1})(\omega)$ denotes the expected value of the game at time t , if we have so far seen the first $t - 1$ elements of P and are told whether or not U holds (that is, whether or not $\omega \in U$). The lemma follows from the following claim by the Backward Induction Theorem.

Claim. *For all $\omega \in U$ and for all $t \in [n]$,*

$$\mathbb{E}(\beta_t | \mathcal{H}_{t-1})(\omega) < v(t),$$

where $\mathcal{H}_0 = \{\emptyset, U, U^c, \Omega\}$ and so $\mathbb{E}(\beta_1 | \mathcal{H}_0)(\omega) = \mathbb{E}(\beta_1 | U)$.

Proof of claim. We shall prove the claim by induction on $n - t$, using (4.5) and (4.6). First, recall that $\mathbb{E}(\beta_n | \mathcal{H}_{n-1})(\omega)$ is just the probability that the final element of ω is maximal in P . Thus, by (4.5),

$$\mathbb{E}(\beta_n | \mathcal{H}_{n-1}) = \frac{k}{n} < v(n),$$

which proves the case $n - t = 0$.

So let $1 \leq t \leq n - 1$, and assume that the result holds for $t + 1$. We claim that

$$\mathbb{E}(\beta_t | \mathcal{H}_{t-1})(\omega) < \begin{cases} \tilde{p}_t y_t + (1 - \tilde{p}_t)v(t + 1) & \text{if } t > I, \\ v(t + 1) & \text{if } t \leq I, \end{cases} \quad (4.7)$$

and hence that $\mathbb{E}(\beta_t | \mathcal{H}_{t-1})(\omega) < v(t)$, by (4.6).

In order to prove (4.7), let $\omega \in U$ and consider the atom $A \in \mathcal{H}_{t-1}$ which contains ω . Partitioning the space according to whether or not $Z_t \geq \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)$, we obtain

$$\begin{aligned} \mathbb{E}(\beta_t | \mathcal{H}_{t-1})(\omega) &= \mathbb{E}(\beta_t | A) = \mathbb{P}(Z_t \geq \mathbb{E}(\beta_{t+1} | \mathcal{H}_t) | A) \cdot \mathbb{E}(Z_t | (Z_t \geq \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)) \cap A) \\ &\quad + \mathbb{P}(Z_t < \mathbb{E}(\beta_{t+1} | \mathcal{H}_t) | A) \cdot \mathbb{E}(\mathbb{E}(\beta_{t+1} | \mathcal{H}_t) | (Z_t < \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)) \cap A), \end{aligned} \quad (4.8)$$

since if $Z_t \geq \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)$ then the payoff is Z_t , and otherwise the payoff is $\mathbb{E}(\beta_{t+1} | \mathcal{H}_t)$.

Now, by the induction hypothesis we have

$$\mathbb{E}(\beta_{t+1} | \mathcal{H}_t)(\omega') < v(t + 1)$$

for every $\omega' \in (Z_t < \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)) \cap A$, and therefore

$$\mathbb{E}(\mathbb{E}(\beta_{t+1} | \mathcal{H}_t) | (Z_t < \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)) \cap A) < v(t + 1). \quad (4.9)$$

Moreover, by Lemma 4.2 and (4.3), and since $\omega \in U$, we have

$$Z_t(\omega) \leq y_t \quad \text{and} \quad \mathbb{P}(Z_t > 0 | \mathcal{H}_{t-1})(\omega) \leq \tilde{p}_t.$$

Hence

$$\mathbb{E}(Z_t | (Z_t \geq \mathbb{E}(\beta_{t+1} | \mathcal{H}_t)) \cap A) \leq y_t, \quad (4.10)$$

and

$$\mathbb{P}(Z_t < \mathbb{E}(\beta_{t+1} | \mathcal{H}_t) | A) \geq 1 - \tilde{p}_t. \quad (4.11)$$

Finally, recall that I was chosen so that $y_t < v(t+1)$ if and only if $t \leq I$. Now (4.7) follows easily from (4.8), (4.9), (4.10) and (4.11). This completes the induction step, and hence proves the claim. \square

The lemma follows from the claim, since, by the backward induction theorem, we have

$$\sup_{\tau \in \mathcal{C}(\mathcal{H}_t)} \mathbb{E}_{D_k(\ell m)}(Z_\tau | U) = \mathbb{E}(\beta_1 | U) < v(1) = \mathbb{E}(\alpha_1) = \sup_{\tau \in \mathcal{C}(\mathcal{G}_t)} \mathbb{E}_{D_k(\ell m)}(Y_\tau),$$

as required. \square

In the final lemma before the proof of Theorem 4.1, we use backward induction to show that an optimal stopping time for the process with the Y_t takes the simple form of rejecting the first $k\ell u^*$ elements, for some integer u^* , and accepting the next non-zero pay-off.

Lemma 4.7. *For $u^* \in \{0, \dots, m-1\}$, let*

$$\tau_{u^*} = \begin{cases} \min \{t > k\ell u^* : Y_t > 0\} & \text{if this exists,} \\ n & \text{otherwise.} \end{cases}$$

Then

$$\sup_{\tau \in \mathcal{C}(\mathcal{G}_t)} \mathbb{E}_{D_k(\ell m)}(Y_\tau) = \sup_{u^* \in \{0, \dots, m-1\}} \mathbb{E}_{D_k(\ell m)}(Y_{\tau_{u^*}})$$

Proof. We have already shown in the proof of Lemma 4.6 that an optimal strategy takes the form ‘ignore the first I , and accept the next non-zero pay-off.’ In fact, we see that I must be a multiple of $k\ell$: suppose, for contradiction, that $I = k\ell u^* + r$, where $r \in \{1, \dots, k\ell - 1\}$. Then $s(I) = s(I+1)$ by (4.2); recall from (4.4) that $y_t = (s(t) + 1)/m < v(t+1)$ if and only if $t \leq I$. It follows that

$$v(I+2) \leq \frac{s(I+1) + 1}{m} = \frac{u^* + 2}{m},$$

since we would be willing to stop at time $I+1$. But then, by (4.6),

$$v(I+1) \leq \max \{y_{I+1}, v(I+2)\} = \frac{u^* + 2}{m} = \frac{s(I) + 1}{m}.$$

Thus we would in fact be willing to stop at time I , which is a contradiction. \square

We are now in a position to complete the proof of the main theorem in this section.

Proof of Theorem 4.1. All that remains is to calculate and maximize $\mathbb{E}(Y_{\tau_{u^*}})$ over u^* , where τ_{u^*} is the smallest $t > k\ell u^*$ such that $Y_t > 0$. These calculations are very similar to those on page 13, but also include error terms which tend to zero as $\ell, m \rightarrow \infty$.

We shall assume first that $u^* \rightarrow \infty$ as $m \rightarrow \infty$, and deduce from our calculation that this is a valid assumption. Indeed, if $t = o(n)$ then $y_t = o(1)$, whereas we shall obtain a probability of success that is separated from zero. We should therefore never accept a payoff for $t = o(n)$. In particular, for sufficiently large m we have $u^* \geq 3$, and so

$$\mathbb{E}(Y_{\tau_{u^*}}) = \sum_{u=u^*}^{m-1} \sum_{r=1}^{k\ell} \mathbb{P}(Y_{k\ell u^*+1} = 0, \dots, Y_{k\ell u^*+r-1} = 0, Y_{k\ell u^*+r} > 0) y_{k\ell u^*+r}.$$

Recall that the Y_t are independent, and that $s(k\ell u + h) = u + 1$ when $1 \leq h \leq k\ell$. Hence, using the convention that the empty product takes the value 1, we obtain

$$\begin{aligned}
\mathbb{E}(Y_{\tau_{u^*}}) &= \sum_{u=u^*}^{m-1} \sum_{r=1}^{k\ell} \left(\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \left(1 - \frac{1}{\ell(u-1)} \right)^{r-1} \cdot \frac{1}{\ell(u-1)} \cdot \frac{u+2}{m} \\
&= \sum_{u=u^*}^{m-1} \left(\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \left(1 - \left(1 - \frac{1}{\ell(u-1)} \right)^{k\ell} \right) \cdot \frac{u+2}{m} \\
&= \sum_{u=u^*}^{m-1} \left(\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \cdot \frac{u+2}{m} \\
&\quad - \sum_{u=u^*+1}^m \left(\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \cdot \frac{u+1}{m} \\
&= \frac{u^*+2}{m} - \left(\prod_{q=u^*+1}^m \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \cdot \frac{m+1}{m} \\
&\quad + \sum_{u=u^*+1}^{m-1} \left(\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \right) \cdot \frac{1}{m}.
\end{aligned}$$

Now, let $\varepsilon > 0$ be arbitrary, choose $m = m(\varepsilon)$ and $\ell = \ell(m, \varepsilon)$ sufficiently large, and recall that therefore $u^* = u^*(\varepsilon)$ may be chosen to be sufficiently large also. Since $(1 - \frac{1}{n})^n < \frac{1}{e} < (1 - \frac{1}{n})^{n-1}$ for all $n \geq 2$, we have

$$\prod_{q=u^*+1}^m \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \geq \exp \left(- \left(\frac{k\ell+1}{\ell} \right) \sum_{q=u^*+1}^m \frac{1}{q-2} \right) \geq \left(\frac{u^*}{m} \right)^k - \frac{\varepsilon}{2},$$

and similarly

$$\prod_{q=u^*+1}^u \left(1 - \frac{1}{\ell(q-2)} \right)^{k\ell} \leq \exp \left(-k \sum_{q=u^*+1}^u \frac{1}{q-2} \right) \leq \left(\frac{u^*}{u} \right)^k.$$

Setting $p = \frac{u^*}{m}$, we obtain

$$\mathbb{E}(Y_{\tau_{u^*}}) \leq \begin{cases} p - p + p \log \frac{1}{p} + \varepsilon & \text{if } k = 1, \\ p - p^k + \frac{p}{k-1} (1 - p^{k-1}) + \varepsilon & \text{if } k > 1. \end{cases}$$

As before, these expressions are maximized when $p = p_k$, and thus

$$\lim_{m \rightarrow \infty} \lim_{\ell \rightarrow \infty} \mathbb{E}(Y_{\tau_{u^*}}) \leq p_k + \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, this completes the proof. \square

Putting Corollary 3.5 and Theorem 4.1 together gives Theorem 1.1.

ACKNOWLEDGEMENTS

We should like to thank Béla Bollobás and Graham Brightwell for their several significant contributions to this paper. In particular, we are grateful for important ideas relating to an earlier proof of Theorem 1.2, for further productive discussions about the problem, and for their comments on the paper.

REFERENCES

- [1] B. Bollobás, *Modern graph theory*, Springer Verlag, 1998.
- [2] Y.S. Chow, H. Robbins and D. Siegmund, *The theory of optimal stopping*, Dover, 1991.
- [3] R.P. Dilworth, A decomposition theorem for partially ordered sets, *Annals of Math.*, **51** (1950), 161–166.
- [4] T.S. Ferguson, Who solved the secretary problem?, *Statistical Science*, **4** (1989), 282–289.
- [5] R. Freij and J. Wästlund, Partially ordered secretaries, *Electron. Comm. Probab.*, **15** (2010), 504–507.
- [6] B. Garrod, G. Kubicki and M. Morayne, How to choose the best twins, *SIAM J. Discrete Math.*, **26** (2012), 384–398.
- [7] N. Georgiou, M. Kuchta, M. Morayne and J. Niemiec, On a universal best choice algorithm for partially ordered sets, *Random Struct. Algorithms*, **32** (2008), 263–273.
- [8] A.V. Gnedin, A multicriteria problem of optimal stopping of a selection process, *Automation and Remote Control*, **42** (1981), 981–986.
- [9] A.V. Gnedin, Multicriteria extensions of the best choice problem: Sequential selection without linear order, *Strategies for Sequential Search and Selection in Real Time*, Contemporary Mathematics, vol. 125, American Mathematical Society, Providence, Rhode Island, 1992, pp. 153–172.
- [10] J. Kozik, Dynamic threshold strategy for universal best choice problem, *DMTCS Proceedings, 21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms*, (2010), 439–452.
- [11] M. Kuchta and M. Morayne, A secretary problem with many lives, submitted.
- [12] R. Kumar, S. Lattanzi, S. Vassilvitskii and A. Vattani, Hiring a Secretary from a Poset, In *Proceedings of ACM Conference on Electronic Commerce* (2011), 39–48.
- [13] M. Morayne, Partial-order analogue of the secretary problem: the binary tree case, *Discrete Math.*, **184** (1998), 165–181.
- [14] J. Preater, The best-choice problem for partially ordered objects, *Operations Research Letters*, **25** (1999), 187–190.

DEPARTMENT OF PURE MATHEMATICS AND MATHEMATICAL STATISTICS, CENTRE FOR MATHEMATICAL SCIENCES, UNIVERSITY OF CAMBRIDGE, WILBERFORCE ROAD, CAMBRIDGE, CB3 0WB, UK

E-mail address: b.garrod@dpms.cam.ac.uk

IMPA, ESTRADA DONA CASTORINA 110, JARDIM BOTÂNICO, RIO DE JANEIRO, RJ, BRASIL

E-mail address: rob@impa.br