Prophet Inequalities for Independent Random Variables from an Unknown Distribution

José Correa[*] Paul Dütting[†] Felix Fischer[‡] Kevin Schewior[§]

Abstract

A classic problem in optimal stopping is the single-choice prophet problem for independent, identically distributed random variables: Given a sequence of independent random variables $X_1 \sim F, \ldots, X_2 \sim F$, the goal is to come up with a stopping rule that decides on a time τ at which to stop so as to maximize the ratio $\mathbb{E}[X_\tau]/\mathbb{E}[\max_t X_t]$ between the expected reward claimed this way and the expected maximum reward in hindsight, in the worst-case over all possible distributions. For quite some time the best known bound for this problem was $1-1/e \approx 0.632$ [Hill and Kertz, 1982]. Only recently this bound was improved by Abolhassani et al. [2017], and a tight bound of ≈ 0.745 was obtained by Correa et al. [2017].

Although a very well motivated problem (see, e.g., [Azar et al., 2014]), much less is known about the case where the underlying distribution is unknown, and has to be learned. A straightforward guarantee for this problem is the $1/e \approx 0.368$ approximation that follows from applying the secretary algorithm, however, it is unclear whether one can improve upon this bound. Our main result is that this bound is tight (i.e., we provide a matching impossibility result). The proof of this impossibility result is quite intricate, and uncovers an unexpected connection between our problem and the celebrated theorem of Ramsey. Motivated by this impossibility result, we then investigate the case where the stopping rule additionally has access to samples. A straightforward extension of our main result shows that with $o(n)$ samples it is still impossible to beat $1/e \approx 0.368$, so the interesting case is when the stopping rule has access to $\Omega(n)$ samples. We show that it is possible to achieve a ratio of $1-1/e \approx 0.632$ with n samples and a ratio of $0.745-\varepsilon$ with $O(n^2)$ samples. The former result uses a clever recycling technique to bring down the number of samples, and is close to optimal as with n samples it is impossible to beat $\ln(2) \approx 0.693$. The second result applies the Correa et al. algorithm to the empirical distribution function, and uses the Dvoretzky-Kiefer-Wolfowitz inequality to bound the error of approximation. Finally, we provide evidence that any algorithm achieving the optimal bound of 0.745 with $O(n)$ samples would have to go through very different techniques.

[*]Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile. Email: correa@uchile.cl

[†]Department of Mathematics, London School of Economics, London, UK. Email: p.d.duetting@lse.ac.uk.

[‡]School of Mathematical Sciences, Queen Mary University of London, London, UK. Email: felix.fischer@qmul.ac.uk.

[§]Fakultät für Informatik, Technische Universität München, München, Germany; Département d'Informatique, École Normale Supérieure, Paris, France. Email: kschewior@gmail.com. Supported by Conicyt Grant PII 20150140 and DAAD PRIME program.

1 Introduction

The theory of optimal stopping is concerned with what a computer scientist would call online algorithms, and the basic problem is one of sequential decision making with imperfect information about the future so as to maximize some reward or minimize some cost. Two canonical problems in the field are the *secretary problem* and the *prophet problem*. Both problems have over the past few years also received considerable attention from the theoretical computer science and algorithms community, particularly since they are closely related to the design of posted-price mechanism in online sales.

In the *secretary problem* we are given n distinct, non-negative numbers from an unknown range. These numbers are presented in random order, and the goal is to stop at one of these numbers in order to maximize the probability with which we pick the maximum. The problem has a surprisingly simple, and surprisingly positive answer: by discarding a $1/e$ fraction of the numbers, and then picking the first number that is greater than any of the discarded numbers, one is guaranteed to pick the maximum with probability $1/e$ [e.g., Gilbert and Mosteller, 1966]. The guarantee of $1/e$ provided by this simple stopping rule is best possible, and remains best possible for example when numbers come from a uniform distribution with unknown and randomly chosen endpoints and are therefore correlated random variables [Berezovskiy and Gnedin, 1984; Ferguson, 1989]. When number are i.i.d. according a known distribution a better guarantee of around 0.58 can be obtained [Gilbert and Mosteller, 1966]

In the *prophet problem* we are again shown n non-negative numbers, one at a time, but now these numbers are independent draws from known distributions, and our goal is to maximize the expected value of the number on which we stop relative to the expected maximum value in hindsight. The two main results here concern the case where the distributions are distinct, and the case where they are identical. For the former a tight bound of $1/2$ was given by Krengel and Sucheston [1977, 1978]. For a long time the best known result for the case of identical distributions was a bound of $1 - 1/e$ due to [Hill and Kertz, 1982], until very recently it was improved first to 0.738 [Abolhassani et al., 2017] and then to 0.745 [Correa et al., 2017]. The latter bound is known to be tight due to a matching bound of Hill and Kertz [1982], which was analyzed by Kertz [1986].

An interesting variant of the (single-choice) prophet problem, for both identical and non-identical distributions, can be obatined by assuming that the distributions from which values are drawn are unknown. Despite being very natural [e.g., Azar et al., 2014], preciously little is known about this variant.

1.1 Our Contribution

In this paper we consider the (single-choice) prophet problem in which valuations are drawn from a single, unknown distribution; and we ask which approximation guarantees on can obtain relative to the expected maximum value in hindsight. We think this is an interesting problem specifically for i.i.d. distributions, as here one could hope to learn something about future values from the earlier ones. More generally, this seems to be a challenging question, as unlike in the case where the distributions are known, where one can use backward induction, it is unclear how an optimal algorithm would look like.

It is in fact relatively straightforward to see that one can achieve a $1/e$ approximation by applying the secretary algorithm. Using the secretary algorithm we are guaranteed to stop on the maximum with probability at least $1/e$, and one can show that this implies a $1/e$ approximation relative to the expected maximum in hindsight. This analysis, however, seems very crude and in particular does not take into account that we are also rewarded if we do not stop on the maximum value. Indeed, one would expect that the prophet objective is easier to achieve than the objective of the secretary problem.

Interestingly, our main result is an impossibility result of $1/e$ that shows that the straightforward guarantee of $1/e$ that follows from applying the secretary algorithm is in fact best possible in the prophet setting. The main difficulty in showing this result is that we want to show an impossibility result that applies to all possible algorithms, not just a subclass such as, for example, threshold rules that use one of the seen numbers as a threshold. In fact, a priori, the class of stopping rules (algorithms) is extremely rich and, unlike in the case of known distributions, it is essentially non-understood. The key to our proof is an in-depth understanding of the types of stopping rules that need to be considered, and — somewhat surprisingly — the main tool in showing that we can in fact restrict attention to this class of stopping rules is Ramsey's theorem [1930].

Motivated by this impossibility result we then turn to the case where the stopping rule additionally has access to samples. A straightforward extension of our main result shows that with $o(n)$ samples it is still impossible to beat $1/e$. The interesting domain is therefore when the stopping rule has access to $\Omega(n)$ samples. Our first result for this setting (and our main result for the case with samples) is a simple (but clever) algorithm that achieves a $1-1/e$ approximation with $n-1$ samples. This algorithm proceeds by drawing $n-1$ samples and using the maximum of the samples as a threshold for the first random variable. If the first random variable exceeds the threshold, the algorithm stops here. Otherwise, the algorithm first adds the newly sampled random variable to the set of samples, and then it removes a random element form the resulting set. Afterwards, it uses the maximum of that set as a threshold for the second random variable, and so on.

While this algorithm is easy to state, it is rather tricky to analyze. The key step in the proof is to show that the sets of random variables used to set the thresholds behave like $n-1$ fresh samples. So in particular, the probability of accepting a random variable conditioned on reaching it is exactly $1/n$, and the expected value that we collect from each random variable conditioned on accepting it equals the expected maximum value of n independent draws from the distribution.

We do not know whether the $1-1/e$ approximation obtained by this algorithm is best possible for algorithms that have access to n samples, but we show that any such algorithm would have to use different techniques, as by setting the same random threshold at each step (as in our algorithm) one cannot hope to beat $1-1/e \approx 0.632$ even if the distributions are known. We also provide an impossibility result showing that our algorithm makes almost optimal use of the samples it is given. Namely, we show that with n samples it is impossible to beat $\ln(2) \approx 0.693$.

Finally, we show how to get arbitrarily close to the optimal bound of 0.745 with $O(n^2)$ samples. The basic idea here is to mimic the optimal algorithm for known distributions, which uses a decreasing sequence of thresholds as determined by conditional acceptance probabilities ("quantiles"), which are increasing over time. Our algorithm mirrors this approach using the corresponding quantiles of the empirical distribution function. It additionally skips a constant fraction of values at the beginning, and in our analysis we use the Dvoretzky-Kiefer-Wolfowitz inequality [1956] to show simultaneous concentration of all empirical quantiles. These two steps allow us to reduce the number of required samples from $O(n^4)$ to $O(n^2)$, relative to the obvious approach which uses all random variables and uses Chernoff and union bound to show concentration. We provide some evidence that any algorithm that achieves the optimal bound with $o(n^2)$ samples would have to use very different techniques.

1.2 Further Related Work

For early work on the classic (single-choice) prophet inequality in Mathematics we refer the reader to a survey of Hill and Kertz [1992]. Starting with work by Hajiaghayi et al. [2007] prophet inequalities, in particular extensions to richer feasibility domains, have seen a surge of interest in Theoretical Computer Science over the past decade. Examples include [Chawla et al., 2010; Alaei, 2014; Kleinberg and Weinberg, 2012; Feldman et al., 2015, 2016; Rubinstein, 2016; Rubinstein and Singla, 2017; Dütting et al., 2017].

In Theoretical Computer Science there is a relatively thin (but important) base of prior work on the unknown distributions case. Most relevant to us is the aforementioned paper by Azar et al. [2014], which focuses on richer feasibility structures, such as matching constraints and matroids. Also relevant to us is even earlier work by Babaioff et al. [2011], who consider a similar setting as we do, but focus on a different objective (revenue maximization), which leads to very different results and requires different techniques.

Related learning problems have also been studied in the Operations Research and Management Science community (see, e.g., the recent survey of den Boer [2015]), but the type of problems, objectives, and techniques are fairly different and typically involve regret minimization problems.

2 Preliminaries

Let $k \in \mathbb{N}_0$ and $n \in \mathbb{N}$. Throughout the paper we typically deal with algorithms that have access to k samples of some distribution and then sequentially observe n values taken from the same distribution. Our algorithms are online in that the (possibly randomized) decision to stop at a given value may only depend on the k samples and on the values observed until that point. The basic underlying notion that captures such algorithms is that of *stopping rule*, while the implied random variable that captures the time at which the algorithm stops is called a *stopping time*. More precisely, a function $r : \mathbb{R}^{k+n} \to [0,1]^n$ is a (k,n)-stopping rule if $\sum_{i=1}^{n}(r(\mathbf{v}))_i \leq 1$ for all $\mathbf{v} \in \mathbb{R}^{k+n}$, and $(r(\mathbf{v}))_i = (r(\mathbf{w}))_i$ for all $i \in \{1,\ldots,n\}$ and $\mathbf{v},\mathbf{w} \in \mathbb{R}^{k+n}$ such that $\mathbf{v}_j = \mathbf{w}_j$ for all $j \leq k+i$. The stopping time τ of a (k,n)-stopping rule r, given random variables $S_1,\ldots,S_k$ and $X_1,\ldots,X_n$, is the random variable with support $\{1,\ldots,n\} \cup \{\infty\}$ such that $\Pr[\tau = i \mid S_1 = s_1,\ldots,S_k = s_k, X_1 = x_1,\ldots,X_n = x_n] = r(s_1,\ldots,s_k,x_1,\ldots,x_n)$. We will consider random variables $S_1,\ldots,S_k$ and $X_1,\ldots,X_n$ that are independent and identically distributed, and respectively denote by f and F the probability density function and cumulative distribution function of their distribution. For a given stopping rule we will be interested in the expected value $\mathbb{E}[X_\tau]$ of the variable at which it stops, where we use the convention that $X_\infty = 0$, and will measure its performance relative to the expected maximum $\mathbb{E}[\max\{X_1,\ldots,X_n\}]$ of the random variables $X_1,\ldots,X_n$. We will say that a stopping rule achieves guarantee α, for $\alpha \leq 1$, if for any distribution, $\mathbb{E}[X_\tau] \geq \alpha\mathbb{E}[\max\{X_1,\ldots,X_n\}]$. For $i \in \mathbb{N}$ we write $[i] = \{1,\ldots,i\}$ for the set of the first i natural numbers and $\mathcal{S}_i$ for the set of permutations of $[i]$.

For ease of exposition we will assume continuity of F in proving lower bounds and use discrete distributions to prove upper bounds. All results can be shown to hold in general by standard arguments, to break ties among random variables and to approximate a discrete distribution by a continuous one.

3 Prophet Inequalities via Learning

In this section we investigate the classic problem of single-item prophet inequalities from a learning perspective, for a situation where the only information available about the random variables besides their values is that they are drawn independently from the same distribution. In Section 3.1 we will see that a straightforward baseline can be obtained from the optimal solution to the secretary problem, which skips a $1/e$ fraction of the values and then accepts the first value that exceeds the maximum of the skipped values. The algorithm is guaranteed to stop at the maximum of the sequence with probability $1/e$, and can be shown to also provide a $1/e$ approximation for our objective. This analysis seems crude and in particular does not account for the fact that the prophet inequality is rewarded even when it does not stop on the maximum value of the sequence. Indeed the objective of the prophet inequality seems easier to achieve than that of the secretary problem, and one would expect to be able to improve on the bound of $1/e$. Our main result, proved in Section 3.2, shows that this is not the case: the $1/e$ bound is in fact best possible.

3.1 A Straightforward Baseline

In the appendix, we prove formally the claim that the guarantee of $1/e$ for the secretary problem translates into a prophet inequality for unknown distributions.

Theorem 1. *Let $X_1, X_2, \ldots, X_n$ be drawn independently from an unknown distribution. Then there exists an algorithm for choosing a stopping time τ such that*

$$\mathbb{E}\left[X_\tau\right] \geq \frac{1}{e} \cdot \mathbb{E}\left[\max\{X_1, X_2, \ldots, X_n\}\right].$$

3.2 A Matching Impossibility Result

We proceed to show our main result: perhaps surprisingly, it is impossible to improve on the straightfoward bound of $1/e$.

Theorem 2. *Let τ be a stopping time for n random variables $X_1, \ldots, X_n$. Let $\delta > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that when $X_1, \ldots, X_n$ are drawn i.i.d. from F,*

$$\mathbb{E}[X_\tau] \leq \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}[\max\{X_1, \ldots, X_n\}].$$

The main difficulty in showing an impossibility result of this kind is that it applies to the set of all possible stopping times, which a priori is very rich. Indeed we can view any stopping time τ as a family of functions $(p_i^\tau)_{i\in[n]}$, where for all $i \in [n]$, $p_i^\tau : \mathbb{R}_+^i \to [0,1]$ and $p_i^\tau(x_1, \ldots, x_i) = \Pr\left[\tau = i \mid X_1 = x_1, \ldots, X_i = x_i, \tau \geq i\right]$ is the conditional probability of accepting the ith value given realizations for the first i values and given that none of the first $i-1$ values has been accepted.

Our main structural insight will be that we can in fact restrict our attention to stopping rules that are rather limited, in that their decision to stop does not depend on the magnitute of the value of a random variable but only on whether it exceeds those observed previously. As a vehicle for the proof, we also introduce other types of limited stopping rules.

Definition 1. *Consider a stopping time τ. Let $\varepsilon > 0$, $i \in [n]$, and $V \subseteq \mathbb{R}_+$. Then τ is called*

(a) (ε, i)-value-oblivious *on V if there exists $q \in [0,1]$ such that, for all $v_1, \ldots, v_i \in V$ with $v_i > \max\{v_1, \ldots, v_{i-1}\}$, it holds that $p_i^\tau(v_1, \ldots, v_i) \in [q+\varepsilon, q-\varepsilon)$;*

(b) ε-value-oblivious *on V if it is (ε, j)-value-oblivious for all $j \in [n]$;*

(c) weakly (ε, i)-value-oblivious *on V if, for all $v_1, \ldots, v_{i-1} \in V$, there exists a $q \in [0,1]$ such that, for all $v_i \in V$, it holds that $p_i^\tau(v_1, \ldots, v_n) \in [q+\varepsilon, q-\varepsilon)$;*

(d) order-oblivious *if for all $j \in [n]$, all $v_1, \ldots, v_j \in \mathbb{R}_+$ and all permutations $\pi \in \mathcal{S}_{j-1}$, $p_i^\tau(v_1, \ldots, v_j) = p_i^\tau(v_{\pi(1)}, \ldots, v_{\pi(j-1)}, v_i)$.*

Lemma 3 (Main structural lemma)**.** *Let $\varepsilon > 0$. If there exists a stopping time with guarantee α, then there exists a stopping time τ with guarantee α and an infinite set $S \subseteq \mathbb{N}$ such that τ is ε-value-oblivious on S.*

To prove Lemma 3 we will successively restrict the class of stopping rules that need to be considered, first to those whose decision to stop does not take into account the order of the values observed previously. The proof of this part is given in the appendix.

Lemma 4. *If there exists a stopping time with guarantee α, then there exists a stopping time with guarantee α that is order-oblivious.*

The proof of Lemma 3 now proceeds to restrict order-oblivious stopping rules further, first to a weak value-obliviousness and then to value-obliviousness. Both steps of the proof of Lemma 3 rely on the infinite version of Ramsey's Theorem.

Lemma 5 (Ramsey [1930])**.** *Let $c, d \in \mathbb{N}$, and let H be a complete d-uniform hypergraph whose hyperedges are colored with c colors. Then there exists an infinite complete d-uniform sub-hypergraph of H that is monochromatic.*

Proof of Lemma 3. By Lemma 4 it suffices to prove the lemma for order-oblivious stopping times. Let τ be such a stopping time. We fix $\varepsilon > 0$ for the entire proof and show by induction on $j \in [n]$ that there exists an infinite set $S^j \subseteq \mathbb{N}$ such that, for all $i \in [j]$, τ is (ε, i)-value-oblivious on S^j. Note that showing this statement for $j = n$ is indeed sufficient: according to Definition 1 it is identical to the statement of the lemma for this value of j. As $S^1 = \mathbb{N}$ satisfies the induction hypothesis for $j = 1$, we proceed to show it for $j = k > 1$ assuming that it is true for $j < k$.

First observe that we only need to find an infinite set $S^k \subseteq S^{k-1}$ such that τ is (ε, k)-value-oblivious on S^k, because it follows from the induction hypothesis that S^k, as a subset of S^{k-1}, is (ε, i)-value-oblivious on S^i for all $i \in [k-1]$. We construct such a set in two steps, starting with the following claim.

Claim 1. There exists an infinite set $R^k \subseteq S^{k-1}$ such that τ is weakly $(\varepsilon/2, k)$-value-oblivious on R^k.

We construct $R^k = \{r_1^k, r_2^k, \dots\}$ by inductively defining, for any $\ell \geq 1$, a set $P_\ell^k \subseteq S^{k-1}$ and then setting $r_\ell^k = \min P_{\ell-1}^k$. Let $P_0^k = S^{k-1}$. For $\ell > 1$, P_ℓ^k is itself determined inductively. Let $h = \binom{\ell}{k-1}$. We construct $h+1$ sets $P_{\ell,0}^k \supseteq P_{\ell,1}^k \supseteq \cdots \supseteq P_{\ell,h}^k$, starting from $P_{\ell,0}^k = P_{\ell-1}^k \setminus \{r_\ell^k\}$, and then set $P_\ell^k = P_{\ell,h}^k$. Note that $h = 0$ while $\ell < k-1$, and we simply set $P_\ell^k = P_{\ell,0}^k$ for these values of ℓ. Denote by $T_1, \dots, T_h$ the subsets of $\{r_1^k, r_2^k, \dots, r_\ell^k\}$ of cardinality $k-1$ that contain r_ℓ^k. To obtain $P_{\ell,i}^k$ for $i \in [h]$ consider $T_i = \{t_1, \dots, t_{k-1}\}$. Toward the first, and rather trivial, application of Lemma 5 consider an edge-colored complete 1-uniform hypergraph H_i with vertex set $P_{\ell,i-1}^k$, where for all $v \in P_{\ell,i-1}^k$, the singleton hyperedge containing v is colored with color $\lfloor p_k^\tau(t_1, \dots, t_{k-1}, v)/\varepsilon \rfloor$. Note that at most $\lfloor 1/\varepsilon \rfloor + 1$ colors are used in total. Now, by Lemma 5 with $c = \lfloor 1/\varepsilon \rfloor + 1$ and $d = 1$, there exists an infinite set of vertices that induces a monochromatic subgraph of H_i, and we define $P_{\ell,i}^k$ to be such a set of vertices.

Note that for every $(k-1)$-element set $\{s_1, \dots, s_{k-1}\} \subseteq \{r_1^k, r_2^k, \dots, r_\ell^k\}$ there exists $i \in [h]$ such that $\{s_1, \dots, s_{k-1}\} = T_i$. Thus, by construction of $P_{\ell,i}^k \supseteq P_{\ell,i}^k$ and since τ is order-oblivious, there exists $q \in [0,1]$ such that, for all $s_k \in P_\ell^k$, $p_k^\tau(s_1, \dots, s_k) \in [q - \varepsilon/2, q + \varepsilon/2)$. Since $r_m^k \in P_\ell^k$ for all $m > \ell$, τ is weakly (ε, k)-value-oblivious on R^k. Since $R^k \subseteq S^{k-1}$, we have proved of the claim.

We now complete the induction step by proving the following claim.

Claim 2. There is an infinite set $S^k \subseteq R^k$ such that τ is (ε, k)-value-oblivious

Toward the second, and less trivial, application of Lemma 5, we construct a complete $(k-1)$-uniform hypergraph H with vertex set R^k. Consider any set $\{s_1, \dots, s_{k-1}\} \subseteq R^k$ of cardinality $k-1$. By Claim 1, there exists $q \in [0,1]$ such that, for all $s_k \in R^k$, $p_k^\tau(s_1, \dots, s_k) \in [q - \varepsilon/2, q + \varepsilon/2)$. Note that there exists a unique $r \in \mathbb{N}$ such that $r \cdot \varepsilon \in [q - \varepsilon/2, q + \varepsilon/2)$, and color the hyperedge $\{s_1, \dots, s_{k-1}\}$ of H with color r.

Note that for every color r, for every hyperedge $\{s_1, \dots, s_{k-1}\}$ with color r, and for every $s_k \in R^k$, $p_i^\tau(s_1, \dots, s_{k-1}, s_k) \in [r \cdot \varepsilon - \varepsilon, r \cdot \varepsilon + \varepsilon)$. For any set $V \subseteq R^k$ that induces a monochromatic sub-hypergraph of H we thus have that τ is (ε, k)-value-oblivious on V. By Lemma 5 with $c = \lfloor 1/\varepsilon \rfloor + 1$ and $d = k-1$ there exists an infinite set that induces a monochromatic sub-hypergraph, and letting S^k be such a set completes the proof of the claim, the induction step, and the lemma. $\square$

With Lemma 3 at hand, we are now ready to prove the theorem.

Proof of Theorem 2. Consider some stopping time τ and let $n \in \mathbb{N}$. Set $\varepsilon = 1/n^2$. By Lemma 3 there exists an infinite set $V \subseteq \mathbb{N}$ on which τ is ε-value-oblivious. Let $v_1, \dots, v_{n^3}, u \in V$ be pairwise distinct such that

$u \geq n^3 \max v_1, v_{n^3}$. For each $i \in [n]$, let

$$X_i = \begin{cases} v_1 & \text{w.p. } \frac{1}{n^3} \cdot (1 - \frac{1}{n^2}) \\ \vdots & \\ v_{n^3} & \text{w.p. } \frac{1}{n^3} \cdot (1 - \frac{1}{n^2}) \\ u & \text{w.p. } \frac{1}{n^2} \end{cases}.$$

We proceed to bound $\mathbb{E}[\max\{X_1, \ldots, X_n\}]$ from below and $\mathbb{E}[X_\tau]$ from above. To this end, let $X_{(i)}$ denote i-th order statistic of $X_1, \ldots, X_n$, such that $X_{(n)} = \max\{X_1, \ldots, X_n\}$. Then

$$\mathbb{E}[\max\{X_1, \ldots, X_n\}] \geq \Pr[X_{(n)} = u] \cdot u = \frac{1 - o(1)}{n} \cdot u.$$

On the other hand,

$$\begin{aligned}
\mathbb{E}[X_\tau] = {} & \Pr[X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \\
& + \Pr[X_{(n)} = u \wedge X_{(n-1)} = u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} = u \wedge X_{(n-1)} = u] \\
& + \Pr[X_{(n)} \neq u] \cdot \mathbb{E}[X_\tau \mid X_{(n)} \neq u] \\
\leq {} & \frac{1}{n} \cdot \Big(\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot u \\
& \qquad + \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot O(n^{-3}) \cdot u \Big) \\
& + O(n^{-2}) \cdot u + 1 \cdot O(n^{-3}) \cdot u \\
\leq {} & \frac{1 + o(1)}{n} \cdot \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u] \cdot u \\
\leq {} & \frac{1 + o(1)}{n} \cdot \Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \ldots, X_n \text{ are distinct}] \cdot u.
\end{aligned}$$

To complete the proof we argue that

$$\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \ldots, X_n \text{ are distinct}] \leq 1/e + o(1).$$

To see this let $w_1, \ldots, w_n \in \{v_1, \ldots, v_{n^3}, u\}$ such that $w_1 < \cdots < w_n = u$. Under the condition that $\{X_1, \ldots, X_n\} = \{w_1, \ldots, w_n\}$, the values appear in a random order. Moreover, τ is ε-value-oblivious on V. For each $i \in [n]$ and up to an error probability of $n \cdot 2 \cdot \varepsilon = O(1/n)$, it must therefore select X_i with the same probability if $X_i > \max\{X_1, \ldots, X_{i-1}\}$. Up to the same error τ thus faces an instance of the secretary problem, for which the maximum w_n can be selected with probability at most $1/e + o(1)$. Since this is true for every sequence of values $w_1, \ldots, w_n$ as above, we are done. $\square$

4 Prophet Inequalities with Additional Samples

The previous section has revealed a fairly strong and perhaps suprising impossibility: when stopping on i.i.d. random variables rather than arbitrary chosen values like in the secretary problem, it is impossible to improve on the well-known guarantee of $1/e \approx 0.368$ for the latter. While we may have hoped to learn something about the distribution and thus about future values from the random variables we observe, nothing of value can in fact be learned. It is natural to ask whether this impossibility continues to hold when we are given a bit more information about the distribution, in the form of additional samples on which we may not stop but which may be used to aid our decision when to stop on the actual random variables. It is a straightforward extension of Theorem 2 that, even for $o(n)$ samples, the guarantee $1/e$ is best-possible.

4.1 Warm-Up: A $1/2$-Approximation with $n-1$ Samples

To gain some intuition let us first consider the natural approach to sample $n-1$ values $S_1, \dots, S_{n-1}$ from F and to use the maximum of these samples as a uniform threshold for all of the random variables $X_1, \dots, X_n$. It is not difficult to see that the expected value we collect from any random variable X_t conditioned on stopping at that random variable is $\mathbb{E}[\max\{X_1, \dots, X_n\}]$, since under this condition X_t is the maximum of n i.i.d. random variables. We can thus understand the approximation guarantee provided by this approach by understanding the probability that it stops on some random variable. It turns out that this probability, and hence the approximation guarantee, is $1/2 + 1/(4n-2)$. A more detailed analysis, which we provide in Appendix C, also reveals that an improvement over the bound of rounghly $1/2$ is impossible with a uniform threshold but could potentially be achieved by increasing the probability of stopping while maintaining the guarantee when we do stop.

4.2 A $(1-1/e)$-Approximation with $n-1$ Samples

We proceed to show that it is indeed possible to obtain an improved bound of $1-(1-1/n)^n \geq 1-1/e \approx 0.632$ with just $n-1$ samples.

Theorem 6. *Let $X_1, X_2, \dots, X_n$ be drawn independently from an unknown distribution. Then there exists an algorithm for choosing a stopping time τ that uses $n-1$ independent samples from the same distribution with*

$$\mathbb{E}[X_\tau] = \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

Let us first convince ourselves that the bound would be straightforward to achieve if we were given access to $n \cdot (n-1) \in \Theta(n^2)$ samples, by partitioning the samples into n sets of size $n-1$ each and using the maximum of the ith set as a threshold for the ith random variable. Upon acceptance of any random variable, that random variable would have a value equal to the expected maximum of n i.i.d. random variables, which is equal to $\mathbb{E}[\max\{X_1, \dots, X_n\}]$. Conditioned on reaching the ith random variable it would be accepted with probability $1/n$, for an overall probability of acceptance of $\sum_{i=1}^{n}(1-1/n)^{i-1} \cdot 1/n = 1-(1-1/n)^n$.

Algorithm 1: Fresh Looking Samples

Data: Sequence of i.i.d. random variables $X_1, \dots, X_n$ sampled from an unknown distribution F, sample access to F

Result: Stopping time τ

$\tau \longleftarrow n+1$

$S_1, \dots, S_{n-1} \longleftarrow n-1$ independent samples from F

$S \longleftarrow \{S_1, \dots, S_{n-1}\}$

for $t = 1, \dots, n$ **do**

  if $X_t \geq \max\{S\}$ **then**

    $\tau \longleftarrow t$

    Break

  else

    $S \longleftarrow S \cup \{X_t\}$

    $Z \longleftarrow$ value from S, chosen uniformly at random

    $S \longleftarrow S \setminus \{Z\}$

return τ

Algorithm 1 mimics this approach, but instead of using $n-1$ fresh samples for each of the n random variables constructs $n-1$ fresh-looking samples for each of the n random variables from a *single* set $\{S_1,\ldots,S_{n-1}\}$ of $n-1$ samples. For the first random variable X_1 the algorithm uses a threshold equal to the maximum of the $n-1$ samples. If $X_1 \geq \max\{S_1,\ldots,S_{n-1}\}$, the algorithm stops. Otherwise it adds X_1 to the set of samples, picks one of the elements in $\{S_1,\ldots,S_{n-1},X_1\}$ uniformly at random, and drops this element from the set. The algorithm then continues in the same way, by using the maximum of the set thus obtained as a threshold for the next random variable, and updating the set when a random variable fails to exceed its threshold.

To analyze the algorithm it will be useful to consider a sequence $j_1,\ldots,j_{n-1}$ of random variables drawn independently and uniformly from $[n]$. Then, for $i=1,\ldots,n$, define variables $R^i_1,\ldots,R^i_n$ recursively as follows

$$R^1_\ell = \begin{cases} S_\ell & \text{for } \ell = 1,\ldots,n-1 \\ X_1 & \text{for } \ell = n \end{cases} \qquad \text{and} \qquad R^i_\ell = \begin{cases} R^{i-1}_\ell & \text{for } \ell = 1,\ldots,j_{i-1}-1 \\ R^{i-1}_{\ell+1} & \text{for } \ell = j_{i-1},\ldots,n-1 \\ X_i & \text{for } \ell = n \end{cases} \quad \text{for } t>1.$$

Then the random variable $X_i = R^i_n$ and the threshold that we set for this random variable is $\max\{R^i_1,\ldots,R^i_{n-1}\}$. Denote by ξ_i the event that random variable X_i exceeds the threshold that we set for it, that is $R^i_n = \max\{R^i_1,\ldots,R^i_n\}$. An important observation is that ξ_i is independent from $\xi_1,\ldots,\xi_{i-1}$.

Lemma 7. *For every $i \in [n]$, $Pr[\xi_i \cap (\bigcap_{j<i} \neg\xi_j)] = Pr[\xi_i] \cdot \prod_{j<i} Pr[\neg\xi_j]$.*

Proof. It suffices to show that for all $t=1,\ldots,i-1$, the event $\xi_i \cap (\bigcap_{j=t+1}^{i-1} \neg\xi_j)$ is independent of the event $\neg\xi_t$, i.e.,

$$\Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, \neg\xi_t\right] = \Pr\left[\xi_i \cap \bigcap_{j=t+1}^{i-1} \neg\xi_j\right].$$

We claim that

$$\begin{aligned}
\Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right)\right] &= \sum_{\ell=1}^{n} \Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, R^t_\ell = \max\{R^t_1,\ldots,R^t_n\}\right] \\
&\qquad \cdot \Pr\left[R^t_\ell = \max\{R^t_1,\ldots,R^t_n\}\right] \\
&= \frac{n-1}{n} \Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, R^t_n < \max\{R^t_1,\ldots,R^t_n\}\right] \\
&\quad + \frac{1}{n} \Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, R^t_n = \max\{R^t_1,\ldots,R^t_n\}\right] \\
&= \Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, R^t_n < \max\{R^t_1,\ldots,R^t_n\}\right] \\
&= \Pr\left[\xi_i \cap \left(\bigcap_{j=t+1}^{i-1} \neg\xi_j\right) \,\middle|\, \neg\xi_t\right].
\end{aligned}$$

Indeed, the first equality can be obtained by distinguishing the index where the maximum is attained. For the second equality observe that the first probability on its left-hand side is the same for all values of ℓ because the events are independent of the choice of ℓ and conditioning is symmetric, and that the maximum is attained with probability $1/n$ at each index. This establishes the claim. $\square$

Proof of Theorem 6. The value $\mathbb{E}[X_\tau]$ obtained by Algorithm 1 can be written by summing over all possible stopping times $i = 1, \dots, n$ the product of the probability of stopping at $X_i = R_n^i$ and the expectation of X_i upon stopping, i.e.,

$$\mathbb{E}[X_\tau] = \sum_{i=1}^{n} \Pr[\xi_i \wedge \neg\xi_{i-1} \wedge \cdots \wedge \neg\xi_1] \cdot \mathbb{E}\left[R_n^i \mid \xi_i \wedge \neg\xi_{i-1} \wedge \cdots \wedge \neg\xi_1\right].$$

By Lemma 7, and since for each $i \in [n]$, the set $\{R_1^i, \dots, R_n^i\}$ is a set of n i.i.d. random variables,

$$\Pr[\xi_i \wedge \neg\xi_{i-1} \wedge \cdots \wedge \neg\xi_1] = \left(1 - \frac{1}{n}\right)^{i-1} \frac{1}{n}.$$

Since R_n^i is independent of $\xi_1, \dots, \xi_{i-1}$, and using again that $\{R_1^i, \dots, R_n^i\}$ is a set of n i.i.d. random variables,

$$\mathbb{E}\left[R_n^i \mid \xi_i \wedge \neg\xi_{i-1} \wedge \cdots \wedge \neg\xi_1\right] = \mathbb{E}\left[R_n^i \mid \xi_i\right] = \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

Thus

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \sum_i \left(1 - \frac{1}{n}\right)^{i-1} \frac{1}{n} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}] \\
&= \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],
\end{aligned}$$

as claimed. □

We do not know if the bound of Theorem 6 can be improved further, but any such an improvement is likely to require a substantially different technique. To this end note that Algorithm 1 uses the same random threshold for each of the random variables $X_1, \dots, X_n$, namely the maximum of $n-1$ random variables. The following result proven in the appendix shows that with a single random threshold it is impossible to achieve a bound greater than $1 - 1/e$, even when the distribution is known.

Proposition 8. *Let $\varepsilon > 0$. Then there exists a distribution and $n \in \mathbb{N}$ such that for a sequence of i.i.d. random variables $X_1, \dots, X_n$ from this distribution and any stopping rule τ that sets the same random threshold for each random variable,*

$$\mathbb{E}[X_\tau] \leq \left(1 - \frac{1}{e} + \varepsilon\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

While an improvement over the bound of $1 - 1/e \approx 0.632$ remains possible via more complicated stopping rules, such an improvement cannot go beyond $\ln(2) \approx 0.693$. This is a consequence of the following strenthening of Theorem 2, which also shows that at least a linear number of samples in n is required to improve over the bound of $1/e$ for the secretary problem. We give a proof in the appendix.

Theorem 9. *Let r be an $(n-1, n)$-stopping rule and $\delta > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that when $S_1, \dots, S_{n-1}, X_1, \dots, X_n$ are drawn i.i.d. from F,*

$$\mathbb{E}[X_\tau] \leq (\ln(2) + \delta) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where τ is the stopping time induced by r on $S_1, \dots, S_{n-1}, X_1, \dots, X_n$.

4.3 A $(\mathbf{0.745}-\varepsilon)$-Approximation with $\mathbf{O(n^2)}$ Samples

Our final result is that it is in fact possible to get arbitrarily close to the optimal approximation guarantee of a stopping algorithm that knows the distribution [Correa et al., 2017], if we have access to $O(n^2)$ samples from the distribution.

We first recall that the optimal algorithm for known distributions computes a decreasing sequence $x_i = y(i/n)^{1/(n-1)}$ for $i \in [n]$, where y is the unique solution (which turns out to be decreasing and convex) to the following ordinary differential equation

$$y' = y(\ln(y)-1)-(\beta-1) \quad \text{and} \quad y(0)=1,$$

where $\beta \approx 1.3414 \approx 1/0.745$. Then conditional on reaching random variable X_i, the algorithm accepts it with probability essentially equal to $\varepsilon_i = 1-x_i$; where for consistency we have have that $0=\varepsilon_0<\varepsilon_1<\ldots<\varepsilon_{n-1}<\varepsilon_n=1$. An important property of the algorithm, obtained by linearity of expectation, is that its revenue can be computed as the sum of the step by step expected revenue. So calling r the stopping time implied by the algorithm and letting $R(q)=\int_0^q F^{-1}(1-\theta)d\theta$, we have that.

$$\mathbb{E}\left[X_r\right] = \sum_{i=1}^n \rho_i \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2}R(q)\mathrm{d}q\,.$$

And the terms ρ_i are chosen so that they are all equal (and equal to $0.745n$). Therefore,

$$\mathbb{E}\left[X_r\right] = n\cdot 0.745\cdot \int_0^1 (n-1)(1-q)^{n-2}R(q)\mathrm{d}q = n\cdot 0.745\cdot \mathbb{E}\left[\max\{X_1,\ldots,X_n\}\right],$$

which is where the approximation factor follows from (for details see Correa et al. [2017]).

Theorem 10. *Let $X_1,\ldots,X_n$ be drawn independently from an unknown distribution. Then for every $\varepsilon>0$ and all $n\geq n_\varepsilon$ there exists an algorithm for choosing a stopping time τ that uses $O(n^2)$ samples from the same distribution with*

$$\mathbb{E}\left[X_\tau\right] \geq (0.745-\varepsilon)\cdot \mathbb{E}\left[\max\{X_1,\ldots,X_n\}\right].$$

Proof sketch. The algorithm that achieves the claimed bound starts by skipping some random variables until the acceptance probability of the optimal algorithm ε_i becomes sufficiently large, say δ/n, say this happens at step k. Afterwards, it uses the empirical distribution function of the samples to estimate the quantiles $\varepsilon_{k+1},\ldots,\varepsilon_n$ used by the optimal algorithm that knows the distribution on the remaining random variables. The algorithm then accepts random variable X_i conditional on reaching it with probability $\tilde{\varepsilon}_i$, where $\tilde{\varepsilon}_i$ is its estimate of ε_i. The reason why we skip the first few random variables is because the initial acceptance probability of the optimal algorithm is of the order of $1/n^2$, therefore with n^2 sample we cannot get a reliable estimate of the corresponding threshold.

To see that the algorithm satisfies the claimed property we first observe that by skipping the first few random variables, until the acceptance probability if δ/n we only loose a small revenue. Indeed, the revenue of the algorithm until the acceptance probability becomes δ/n is given by:

$$n\cdot 0.745\int_0^{\delta/n}(n-1)(1-q)^{n-2}R(q)\mathrm{d}q\,.$$

Since $R(q)$ is monotone, this revenue is at most

$$n\cdot 0.745\int_0^{\delta/n}(n-1)(1-q)^{n-2}\mathrm{d}q\cdot R(\delta/n) \leq n\cdot 0.745(1-e^{-\delta})\cdot R(\delta/n) \leq n\cdot 0.745\delta\cdot R(\delta/n).$$

On the other hand, between this time and the end the algorithm's expected revenue is

$$n \cdot 0.745 \int_{\delta/n}^{1} (n-1)(1-q)^{n-2} R(q) \mathrm{d}q \geq n \cdot 0.745 e^{-\delta} \cdot R(\delta/n) \geq n \cdot 0.745(1-\delta) \cdot R(\delta/n).$$

Therefore the ratio between the two is $\delta/(1-\delta)$, so that by picking a small value δ, the loss can be made arbitrarily small.

The next observation is that because we have $O(n^2)$ samples we can use the Dvoretzky-Kiefer-Wolfowitz inequality [1956] to argue that with probability at least $1-\alpha$ all quantiles of the empirical distribution function will be within a β/n band of the actual quantiles, with β an arbitrarily small constant.

Now conditioned on the fact that all our estimated quantiles lie within the respective error band the probabilities with which our algorithm stops on each of the random variables $X_{k+1}, \ldots, X_n$ is arbitrarily close to the corresponding acceptance probabilities in the optimal algorithm that knows the distribution. Indeed as the error within each step is at most δ/n for the first k random variables and β/n for the random variables starting from k, the cumulative error until step i is at most $1-(1-\max\{\delta,\beta\}/n)^i \approx \max\{\delta,\beta\}$, which by the choice of δ can be made arbitrarily small. The latter shows that the distribution of the stopping time of our algorithm and that of the optimal algorithm are essentially the same.

The last ingredient in the proof is to notice that conditional on stopping at a given time both algorithms get roughly the same amount. This is quite clear since in general for a random variable X and two thresholds τ_1 and τ_2 such that $F(\tau_1)$ and $F(\tau_2)$ are close then $\mathbb{E}[X|X > \tau_1]$ is close to $\mathbb{E}[X|X > \tau_2]$.

In summary our algorithm skips the first few random variables by loosing only a small fraction of the revenue, then stops essentially at the same (random) time than the optimal algorithm, and finally conditional on stopping at time i it obtains essentially the same reward as the optimal algorithm. $\square$

Our argument improves upon the obvious approach that shows simultaneous concentration of all quantiles via Chernoff and union bound by a factor of $\Theta(n^2)$. The first factor n is saved by avoiding having to estimate too small quantiles by skipping a constant fraction of the random variables, the second factor n is saved by exploiting the DKW inequality which shows uniform concentration at the same cost as a single application of a Chernoff-type bound.

We conclude with an argument that suggests that any approach that achieves the optimal 0.745 approximation with $o(n^2)$ samples would have to go through very different techniques. To this end we will show that even if the random variables $X_1, \ldots, X_n$ were uniform on $[0,1]$ in order to get concentration of the median around its expectation within a $O(1/n)$ band at least $\Omega(n^2)$ samples are needed.

Proposition 11. *Let X denote the median of $f(n)$ samples $X_1, \ldots, X_{f(n)}$ drawn independently from the uniform distribution on $[0,1]$. Then, for any constant $\varepsilon > 0$, $f(n)$ must be in the order of $\Omega(n^2)$ to have*

$$Pr\left[|X - \mathbb{E}[X]| \leq \frac{1}{n}\right] \leq \varepsilon.$$

APPENDIX

A Proofs from Subsection 3.1

Proof of Theorem 1. Let τ be the stopping time corresponding to the optimal stopping rule for the secretary problem, which rejects a certain fraction of the random variables and uses thei maximum as a threshold for the remaining ones. Since $X_1, X_2, \ldots, X_n$ are drawn independently from the same distribution, we can assume that their realizations are obtained by independently drawing n values from the distribution and then ordering them according to a random permutation π. Denoting the density of the distribution from which $X_1, \ldots, X_n$ are drawn by f,

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \mathbb{E}_\pi[v_{\pi(\tau)}] \,\mathrm{d}v_1 \cdots \mathrm{d}v_n \\
&\geq \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \Pr_\pi[v_{\pi(\tau)} = \max\{v_1, \ldots, v_n\}] \cdot \max\{v_1, \ldots, v_n\} \,\mathrm{d}v_1 \cdots \mathrm{d}v_n \\
&\geq \frac{1}{e} \cdot \int_0^\infty \cdots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \max\{v_1, \ldots, v_n\} \,\mathrm{d}v_1 \cdots \mathrm{d}v_n \\
&= \frac{1}{e} \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right],
\end{aligned}$$

where the second inequality holds because the values $v_1, \ldots, v_n$ have been randomly ordered and τ is thus guaranteed to select $\max\{v_1, \ldots, v_n\}$ with probability at least $1/e$ for any realization [Ferguson, 1989]. This proves the claim. □

B Proofs from Subsection 3.2

Proof of Lemma 4. For $i \in [n]$, let $\sim_i$ be the equivalence relation on $\mathbb{R}_+^i$ such that $(v_1, \ldots, v_i) \sim_i (w_1, \ldots, w_i)$ if $v_1, \ldots, v_{i-1}$ is a permutation of $w_1, \ldots, w_{i-1}$ and $v_i = w_i$. Note that a stopping time τ is order-oblivious if and only if for all $i \in [n]$ and $v_1, \ldots, v_i, w_1, \ldots, w_i \in \mathbb{R}_+$ $p_i^\tau(v_1, \ldots, v_i) = p_i^\tau(w_1, \ldots, w_i)$ whenever $(v_1, \ldots, v_i) \sim_i (w_1, \ldots, w_i)$. We will refer to the equivalence classes of $\sim_i$ as *states*, and will say that a stopping time τ *arrives at* $s \in \mathbb{R}_+^i / \sim_i$ in the event that $\tau \geq i$ and $X_1 = v_1, \ldots, X_{i-1} = v_{i-1}$ where $[v_1, \ldots, v_i]_{\sim_i} = s$.

Let τ be an arbitrary stopping time, and consider a stopping time σ such that $p_1^\sigma(v_1) = p_1^\tau(v_1)$ and for all $i \in \{2, \ldots, n\}$ and $v_1, \ldots, v_i \in \mathbb{R}_+$ with $\Pr[\tau \text{ arrives at } [v_1, \ldots, v_i]_{\sim_i}] > 0$,

$$p_i^\sigma(v_1, \ldots, v_i) = \Pr\left[\tau = i \mid \tau \text{ arrives at } [v_1, \ldots, v_i]_{\sim_i}\right].$$

Since $[v_1, \ldots, v_i]_{\sim_i}$ is invariant under permutations of the sequence $v_1, \ldots, v_{i-1}$, σ is indeed order-oblivious. It remains to be shown that σ provides guarantee α.

As an intermediate step we show by induction that for all $i \in [n]$ and $s \in \mathbb{R}_+^i / \sim_i$,

$$\Pr[\tau \text{ arrives at } s] = \Pr[\sigma \text{ arrives at } s]. \tag{1}$$

This holds trivially for $i=1$, so we assume that it holds $i=k-1\geq 1$ and show it for $i=k$. Indeed, for any $v_1,\ldots,v_k\in\mathbb{R}_+$ and $s=[v_1,\ldots,v_k]_{\sim_k}$,

$$\begin{aligned}
\Pr[\tau \text{ arrives at } s] &= \sum_{j=1}^{k-1}\Pr\bigl[\tau \text{ arrives at } [v_1,\ldots,v_{j-1},v_{j+1},\ldots,v_{k-2},v_j]_{\sim_{k-1}}\bigr]\\
&\qquad\cdot\Pr\bigl[\tau\neq i\mid \tau \text{ arrives at } [v_1,\ldots,v_{j-1},v_{j+1},\ldots,v_{k-2},v_j]_{\sim_{k-1}}\bigr]\cdot\Pr[X_k=v_k]\\
&= \sum_{j=1}^{k-1}\Pr\bigl[\sigma \text{ arrives at } [v_1,\ldots,v_{j-1},v_{j+1},\ldots,v_{k-2},v_j]_{\sim_{k-1}}\bigr]\\
&\qquad\cdot\Pr\bigl[\sigma\neq i\mid \sigma \text{ arrives at } [v_1,\ldots,v_{j-1},v_{j+1},\ldots,v_{k-2},v_j]_{\sim_{k-1}}\bigr]\cdot\Pr[X_k=v_k]\\
&= \Pr[\sigma \text{ arrives at } s],
\end{aligned}$$

where the first and last equalities hold by definition of $\sim_{k-1}$ and the second equality by the induction hypothesis and by definition of σ.

We now claim that

$$\begin{aligned}
\mathbb{E}\left[X_\tau\right] &= \sum_{i=1}^{n}\mathbb{E}\left[X_i\mid\tau=i\right]\cdot\Pr\left[\tau=i\right]\\
&= \sum_{i=1}^{n}\int_0^\infty\cdots\int_0^\infty\prod_{j=1}^{i}f(v_j)\cdot v_i\\
&\qquad\cdot\frac{1}{(i-1)!}\cdot\sum_{\pi\in\mathcal{S}_{i-1}}\Pr\left[\tau=i\mid X_1=v_{\pi(1)},\ldots,X_{i-1}=v_{\pi(i-1)},X_i=v_i\right]\,\mathrm{d}v_1\ldots\mathrm{d}v_i\\
&= \sum_{i=1}^{n}\int_0^\infty\cdots\int_0^\infty\prod_{j=1}^{i}f(v_j)\cdot v_i\cdot\Pr\bigl[\tau=i\mid\tau \text{ arrives at } [v_1,\ldots,v_i]_{\sim_i}\bigr]\\
&\qquad\qquad\cdot\Pr\bigl[\tau \text{ arrives at } [v_1,\ldots,v_i]_{\sim_i}\bigr]\,\mathrm{d}v_1\ldots\mathrm{d}v_i\\
&= \sum_{i=1}^{n}\int_0^\infty\cdots\int_0^\infty\prod_{j=1}^{i}f(v_j)\cdot v_i\cdot p_i^\sigma(v_1,\ldots,v_i)\cdot\Pr\bigl[\sigma \text{ arrives at } [v_1,\ldots,v_i]_{\sim_i}\bigr]\,\mathrm{d}v_1\ldots\mathrm{d}v_i\\
&= \sum_{i=1}^{n}\mathbb{E}\left[X_i\mid\sigma=i\right]\cdot\Pr\left[\sigma=i\right]=\mathbb{E}\left[X_\sigma\right].
\end{aligned}$$

Indeed, the second equality can be seen to hold by imagining that $X_1,\ldots,X_i$ are drawn by first drawing i values independently and then permuting the first $i-1$ of these values uniformly at random. The fourth equality holds by definition of σ and by (1). This completes the proof. □

C A 1/2-approximation with $n-1$ samples

In this appendix we formalize the discussion in Section 4.1. We show that if the stopping rule has access to $n-1$ samples, then we can simply take the maximum of these samples as a single, non-adaptive threshold for all random variables to obtain a factor 1/2-approximation.

Theorem 12. *Let $X_1,X_2,\ldots,X_n$ be a sequence of i.i.d. random variables from an unknown distribution. Then there exists an algorithm for choosing a stopping time τ that uses $n-1$ independent samples from F with*

$$\mathbb{E}\left[X_\tau\right]\geq\frac{1}{2}\cdot\mathbb{E}\left[\max\{X_1,\ldots,X_n\}\right].$$

To prove Theorem 12 we will analyze a slight variation of the algorithm described above, Algorithm 2, which only uses the maximum of $n-1$ samples as a threshold for the first $n-1$ random variables and stops on the n-th random variable with certainty. The advantage of this is that it becomes even clearer when and why our analysis is tight.

Algorithm 2: Single threshold algorithm

Data: Sequence of i.i.d. random variables $X_1, \ldots, X_n$ sampled from an unknown distribution F, sample access to F

Result: Stopping time τ

```
τ ⟵ n
S_1, ..., S_{n-1} ⟵ n − 1 samples from F
for t = 1, ..., n − 1 do
    if X_t ≥ max{S_1, ..., S_{n-1}} then
        τ ⟵ t
        Break
return τ
```

Proof of Theorem 12. The expected value achieved by Algorithm 2 is the sum over all time steps $i = 1, \ldots, n$ of the product of the probability of stoping at this time step and the expected value of the random variable conditioned on being above the threshold

$$\begin{aligned}\mathbb{E}[X_\tau] &= \sum_{i=1}^{n-1} \Big(\mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \Big) + \mathbb{E}[X_n] \cdot \Pr[\tau = n] \\ &\geq \sum_{i=1}^{n-1} \Big(\mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \Big). \end{aligned} \tag{2}$$

We stop at time step i if the maximum among the $n-1$ samples and the first i random variables happens to be the i-th random variable, and if, conditioned on this, the second maximum is among the $n-1$ samples and not the other $i-1$ random variables. Hence,

$$\Pr[\tau = i] = \frac{1}{n-1+i} \cdot \frac{n-1}{n-2+i}$$

Summing this over all i from 1 to $n-1$ shows that the probability of stopping at one of the first $n-1$ random variables is precisely

$$\sum_{i=1}^{n-1} \Pr[\tau = i] = \sum_{i=1}^{n-1} \frac{1}{n-1+i} \cdot \frac{n-1}{n-2+i} = \frac{1}{2}. \tag{3}$$

We conclude the proof by showing that for all $i = 1, \ldots, n-1$ the conditional expectation $\mathbb{E}[X_i \mid \tau = i]$ is at least $\mathbb{E}[\max\{X_1, \ldots, X_n\}]$. Let $T = \max\{S_1, \ldots, S_n\}$. The algorithm stops at time step i if $X_i \geq T > \max\{X_1, \ldots, X_{i-1}\}$. So under this event X_i is the maximum of $n-1+i$ random variables. And so

$$\mathbb{E}[X_i \mid \tau = i] = \mathbb{E}[\text{max of } n-1+i \text{ i.i.d. RVs}] \geq \mathbb{E}[\max\{X_1, \ldots, X_n\}]. \tag{4}$$

Substituting (3) and (4) into (2) completes the proof. □

As argued in the proof of Theorem 12 the probability that Algorithm 2 stops on one of the first $n-1$ variables is precisely $1/2$. The two potentially lossy steps are that we dropped the contribution from the final random variable, and that we lower bounded the contribution from each of the first $n-1$ random variables by $\mathbb{E}[\max\{X_1,\ldots,X_n\}]$.

It turns out that both of the potentially lossy steps are in fact lossless in the limit as $n \to \infty$ if F is the exponential distribution.

Proposition 13. *Let $X_1,\ldots,X_n$ be a sequence of n i.i.d. random variables sampled from the exponential distribution $F = 1-e^{-x}$. Then for stopping times τ as determined by Algorithm 2*

$$\lim_{n\to\infty} \frac{\mathbb{E}[X_\tau]}{\mathbb{E}[\max\{X_1,\ldots,X_n\}]} = \frac{1}{2}.$$

Proof. It is not super straightforward to compute, but a well-known fact that the maximum of n independent, exponentially distributed random variables $X_1,\ldots,X_n$ is equal to the n-th Harmonic number. That is,

$$\mathbb{E}[\max\{X_1,\ldots,X_n\}] = H_n$$

As argued in the proof of Theorem 12 we can express the expected value obtained by Algorithm 2 as follows

$$\mathbb{E}[X_\tau] = \sum_{i=1}^{n-1}\left(H_{n-1+i}\cdot\frac{1}{n-1+i}\cdot\frac{n-1}{n-2+i}\right)+\frac{1}{2}.$$

Tedious calculations allow to express the expected value via the digamma function $\psi^{(0)}$ and the Euler-Mascheroni constant γ as follows

$$\mathbb{E}[X_\tau] = \psi^{(0)}(n) - \frac{1}{2}H_{2n-2} + \gamma + 1.$$

This can then be used to show that

$$\lim_{n\to\infty} \frac{\mathbb{E}[X_\tau]}{\mathbb{E}[\max\{X_1,\ldots,X_n\}]} = \lim_{n\to\infty} \frac{\psi^{(0)}(n) - \frac{1}{2}H_{2n-2} + \gamma + 1}{H_n} = \frac{1}{2},$$

which proves the claim. □

We conclude by showing that our analysis of Algorithm 2 is tight, and indeed that any stopping rule that beats the bound of $1/2$ has to use a different approach.

Proposition 14. *Suppose $X_1,\ldots,X_n$ are a sequence of i.i.d. random variables, each sampled from a known distribution. Then for any algorithm that determines a stopping time τ by setting the same deterministic threshold to each of the random variables*

$$\mathbb{E}[X_\tau] \le \left(\frac{1}{2}+\varepsilon\right)\cdot\mathbb{E}\left[\max_t X_t\right].$$

Proof. Take n copies of the random variable X that is

$$X_n = \begin{cases} n & \text{with probability } 1/n^2\text{, and} \\ 1 & \text{with probability } 1-1/n^2. \end{cases}$$

Then $\mathbb{E}[\max\{X_1,\ldots,X_n\}] = (1-(1-1/n^2)^n)\cdot n + (1-1/n^2)^n \cdot 1$ which tends to 2 as $n\to\infty$. The only two sensible thresholds are $T = n$ or $T = 1$. For $T = n$ we obtain $\mathbb{E}[X_\tau] = (1-(1-1/n^2)^n)\cdot n$ which tends to 1 as $n \to \infty$. For $T = 1$ we obtain $\mathbb{E}[X_\tau] = 1$. □

D Proofs from Subsection 4.2

Proof of Proposition 8. For each $i \in [n]$, let

$$X_i = \begin{cases} \frac{\sqrt{n}}{e-2} & \text{w.p. } \frac{1}{n^{3/2}}, \\ 1 & \text{w.p. } \frac{1}{\sqrt{n}}, \\ 0 & \text{otherwise.} \end{cases}$$

It is not difficult to see that, when n is large enough,

$$\mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right] \geq \frac{1}{e-2} + 1 - \varepsilon.$$

On the other hand, we can restrict our attention to stopping rules that always accept a value of $\frac{\sqrt{n}}{e-2}$ and never accept a value of 0. A stopping rule that uses a single random threshold is thus characterized by a probability α with which it accepts a value of 1. Denote by τ_α the stopping rule that accepts a value of 1 with probability α. Then, when n is large enough,

$$\mathbb{E}\left[X_{\tau_\alpha}\right] = \mathbb{E}\left[X_{\tau_\alpha} \mid X_{\tau_\alpha} > 0\right] \cdot \Pr\left[X_{\tau_\alpha} > 0\right],$$

where

$$\begin{aligned} \mathbb{E}\left[X_{\tau_\alpha} \mid X_{\tau_\alpha} > 0\right] &= \frac{1/n^{3/2}}{1/n^{3/2} + \alpha/\sqrt{n}} \cdot \frac{\sqrt{n}}{e-2} + \frac{\alpha/\sqrt{n}}{1/n^{3/2} + \alpha/\sqrt{n}} \cdot 1 \\ &= \frac{1/(n(e-2)) + \alpha/\sqrt{n}}{1/n^{3/2} + \alpha/\sqrt{n}} \end{aligned}$$

and

$$\begin{aligned} \Pr\left[X_{\tau_\alpha} > 0\right] &= 1 - \left(1 - \frac{1}{n^{3/2}} - \frac{\alpha}{\sqrt{n}}\right)^n \\ &\leq 1 - e^{-(\alpha\sqrt{n} + 1/\sqrt{n})} + \varepsilon. \end{aligned}$$

Thus

$$\mathbb{E}\left[X_{\tau_\alpha}\right] \leq \frac{1/(e-2) + \alpha\sqrt{n}}{1/\sqrt{n} + \alpha\sqrt{n}} \cdot \left(1 - e^{-(\alpha\sqrt{n} + 1/\sqrt{n})} + \varepsilon\right).$$

Behavior of the expression on the right-hand side is determined by the value of $\alpha\sqrt{n}$. If $\alpha\sqrt{n}$ increases in n then $\mathbb{E}\left[X_{\tau_\alpha}\right]$ tends to 1. If $\alpha\sqrt{n}$ decreases in n then $e^{-(\alpha\sqrt{n}+1/\sqrt{n})} \geq 1 - (\alpha\sqrt{n} + 1/\sqrt{n}) + \varepsilon$, so that $\mathbb{E}\left[X_{\tau_\alpha}\right]$ tends to $1/(e-2)$. The case where $\alpha\sqrt{n}$ is constant can finally be solved by considering the maximum of $((1/(e-2)+x)/x) \cdot (1 - e^{-x})$, which occurs at $x = 1$.

In summary

$$\frac{\mathbb{E}\left[X_{\tau_\alpha}\right]}{\mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right]} \leq \left(1 - \frac{1}{e} + \varepsilon\right),$$

which shows the claim. ☐

Proof of Theorem 9. Consider any $n \in \mathbb{N}$ and $(n-1,n)$-stopping rule r and let τ be its stopping time. Set $\varepsilon = 1/n^2$. By Lemma 3 there exists an infinite set $V \subseteq \mathbb{N}$ on which r is ε-value-oblivious. Let $v_1,\dots,v_{n^3},u \in V$ be pairwise distinct such that $u \geq n^3 \max\{v_1,\dots,v_{n^3}\}$. For each $i \in [n]$, let

$$X_i = \begin{cases} v_1 & \text{w.p. } \frac{1}{n^3}\cdot(1-\frac{1}{n^2}) \\ \vdots & \\ v_{n^3} & \text{w.p. } \frac{1}{n^3}\cdot(1-\frac{1}{n^2}) \\ u & \text{w.p. } \frac{1}{n^2} \end{cases}$$

We proceed to bound $\mathbb{E}[\max\{X_1,\dots,X_n\}]$ from below and $\mathbb{E}[X_\tau]$ from above. For $i \in [2n-1]$, let

$$R_i = \begin{cases} S_i & \text{if } i \leq n-1 \\ X_{i-n+1} & \text{otherwise.} \end{cases}$$

Let $\tau' = \tau + n - 1$ be the stopping time of r on $R_1,\dots,R_{2n-1}$. Let $X_{(i)}$ denote i-th order statistic of $X_1,\dots,X_n$, such that $X_{(n)} = \max\{X_1,\dots,X_n\}$. Then

$$\mathbb{E}[\max\{X_1,\dots,X_n\}] \geq \Pr[X_{(n)} = u]\cdot u = \frac{1-o(1)}{n}\cdot u.$$

On the other hand,

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \Pr[R_{(2n-1)} = u \wedge R_{(2n-2)} \neq u]\cdot \mathbb{E}[R_\tau \mid R_{(2n-1)} = u \wedge R_{(2n-2)} \neq u] \\
&\quad + \Pr[R_{(2n-1)} = u \wedge R_{(2n-2)} = u]\cdot \mathbb{E}[R_{\tau'} \mid R_{(2n-1)} = u \wedge R_{(2n-2)} = u] \\
&\quad + \Pr[R_{(2n-1)} \neq u]\cdot \mathbb{E}[R_{\tau'} \mid R_{(2n-1)} \neq u] \\
&\leq \frac{2}{n}\cdot \Big(\Pr[R_{\tau'} = R_{(2n-1)} \mid R_{(2n-1)} = u \wedge R_{(2n-2)} \neq u]\cdot u \\
&\qquad + \Pr[R_{\tau'} = R_{(2n-1)} \mid R_{(2n-1)} = u \wedge R_{(2n-2)} \neq u]\cdot O(n^{-3})\cdot u\Big) \\
&\quad + O(n^{-2})\cdot u + 1\cdot O(n^{-3})\cdot u \\
&\leq \frac{2+o(1)}{n}\cdot \Pr[R_{\tau'} = R_{(n)} \mid R_{(2n-1)} = u \wedge R_{(2n-2)} \neq u]\cdot u \\
&\leq \frac{2+o(1)}{n}\cdot \Pr[R_{\tau'} = R_{(2n-1)} \mid R_{(2n-1)} = u \wedge R_1,\dots,R_{2n-1} \text{ are distinct}]\cdot u.
\end{aligned}$$

Let

$$p = \Pr[R_{\tau'} = R_{(2n-1)} \mid R_{(2n-1)} = u \wedge R_1,\dots,R_{2n-1} \text{ are distinct}].$$

To complete the proof we argue that $p \leq \ln(2)/2 + o(1)$. To see this let $w_1,\dots,w_{2n-1} \in \{v_1,\dots,v_{n^3},u\}$ such that $w_1 < \cdots < w_{2n-1} = u$. Under the condition that $\{R_1,\dots,R_{2n-1}\} = \{w_1,\dots,w_{2n-1}\}$, the values appear in a random order. Moreover, r is ε-value-oblivious on V. For each $i \in [n]$ and up to an error probability of $(2n-1)\cdot 2\cdot\varepsilon = O(1/n)$, it must therefore select X_i with the same probability q_i if $X_i > \max\{X_1,\dots,X_{i-1}\}$. But since $R_1,\dots,R_{n-1}$ are in fact samples, $q_i = 0$ for all $i = n,\dots,2n-1$. We are thus faced with an instance of the secretary problem with $2n-1$ values under the additional constraint that the first $n-1$ values have to be rejected. The optimal stopping rule for this problem is known to set, for some $x \in [0,1]$, $q_i = 0$ for all $i < x\cdot(2n-1)$ and $q_i = 1$ for all $i \geq x\cdot(2n-1)$ [Gilbert and Mosteller, 1966]. Then $p = -x\cdot\log x + o(1)$, which subject to $x \geq (n-1)/(2n-1)$ is maximized for $x = (n-1)/(2n-1)$, and thus $p \leq \ln(2)/2 + o(1)$. Since this is true for every sequence of values $w_1,\dots,w_n$ as above, we are done. □

E Proofs from Subsection 4.3

Proof of Proposition 11. The k-th smallest of n samples from a uniform distribution follows a $\mathsf{Beta}(n-k+1,k)$ distribution. So if we are interested in the median of n^2 samples, then for n large, the distribution of the median is well approximated by $X \sim \mathsf{Beta}(n^2/2, n^2/2)$.

The expectation of a random variable drawn from $\mathsf{Beta}(\alpha,\beta)$ is $\alpha/(\alpha+\beta)$. So when $\alpha=\beta$ as in our case this is simply $1/2$. The variance of such a random variable is $\alpha\cdot\beta/[(\alpha+\beta)^2\cdot(\alpha+\beta+1)]$. So for $\alpha=\beta=n^2/2$ it is $1/[4(n^2+1)]$. Now we want to show concentration within a $1/n$ band around the expectation, which is $1/2$. For simplicity, we will look at the one sided error only. So we seek to bound

$$\Pr\left[X \geq \frac{1}{2}+\frac{1}{n}\right].$$

To compute this probability, we will use that for $\alpha=\beta$ large, we can approximate the Beta distribution with a normal distribution. More formally, for $\alpha=\beta=n^2/2$ large the random variable $Y = 2\cdot\sqrt{n^2+1}\cdot(X-1/2)$ has probability density function

$$f_Y(y) = \frac{1}{\sqrt{\pi}}\cdot e^{-\frac{y^2}{2}}.$$

With this we obtain

$$\begin{aligned}\Pr\left[X \geq \frac{1}{2}+\frac{1}{n}\right] &\approx \Pr\left[Y \geq 2\cdot\sqrt{n^2-1}\cdot\frac{1}{n}\right]\\ &= \frac{1}{2}-\frac{1}{2}\cdot\operatorname{erf}\left(2\cdot\sqrt{n^2-1}\cdot\frac{1}{n}\cdot\frac{1}{\sqrt{2}}\right).\end{aligned}$$

Now the argument of the Erlang function tends to $\sqrt{2}$ as n tends to ∞, and $\operatorname{erf}(\sqrt{2}) \approx 0.954$. So the probability that the one sided error is at most $1/n$ is close to 1, and we can make it arbitrarily close to 1 by using $c\cdot n^2$ samples instead.

If on the other hand we used $o(n^2)$ samples, the argument of the Erlang function would tend to zero and hence the probability would not vanish as desired. □

Acknowledgements. We gratefully acknowledge a discussion with Fábio Botler that pointed us to Ramsey's Theorem.

References

M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. D. Kleinberg, and B. Lucier. Beating $1 - 1/e$ for ordered prophets. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pages 61–71, 2017.

S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1377, 2014.

M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. Posting prices with unknown distributions. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, pages 166–178, 2011.

B. A. Berezovskiy and A. V. Gnedin. *The Best Choice Problem (In Russian)*. Nauka, Moscow, 1984.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 311–320, 2010.

J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 169–186, 2017.

A. V. den Boer. Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1):1 – 18, 2015.

P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, pages 540–551, 2017.

A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Annals of Mathematical Statistics*, 27:642–669, 1956.

M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135, 2015.

M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms*, pages 1014–1033, 2016.

T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.

J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.

M. Hajiaghayi, R. Kleinberg, and T. W. Sandholm. Automated mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 58–65, 2007.

T. Hill and R. P. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics*, 125:191–207, 1992.

T. P. Hill and R. P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *The Annals of Probability*, 10:336–345, 1982.

R. P. Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis*, 19:88–112, 1986.

R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th ACM Symposium on Theory of Computing Conference*, pages 123–136, 2012.

U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83:745–747, 1977.

U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Probability Related Topics*, 4:197–266, 1978.

F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 30:264–286, 1930.

A. Rubinstein. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th ACM Symposium on Theory of Computing*, 2016. 324–332.

A. Rubinstein and S. Singla. Combinatorial prophet inequalities. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687, 2017.