





Prophet Inequalities for Independent and Identically Distributed Random Variables from an Unknown Distribution

José Correa,^a Paul Dütting,^b Felix Fischer,^c Kevin Schewior^d

^aDepartamento de Ingeniería Industrial, Universidad de Chile, Santiago 8320000, Chile; ^bDepartment of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom; ^cSchool of Mathematical Sciences, Queen Mary University of London, London E1 4NS, United Kingdom; ^dDepartment Mathematik/Informatik, Universität zu Köln, Köln 50931, Germany

Contact: correa@uchile.cl,  <https://orcid.org/0000-0002-3012-7622> (JC); p.d.duetting@lse.ac.uk,  <https://orcid.org/0000-0002-0635-6812> (PD); felix.fischer@qmul.ac.uk,  <https://orcid.org/0000-0002-8403-9273> (FF); kschewior@gmail.com,  <https://orcid.org/0000-0003-2236-0210> (KS)

Received: July 31, 2019

Revised: October 16, 2020; February 3, 2021

Accepted: March 23, 2021

Published Online in Articles in Advance:
December 20, 2021

MSC2020 Subject Classification: Primary:
62L15, 60G40, 68W27, 91B26

<https://doi.org/10.1287/moor.2021.1167>

Copyright: © 2021 INFORMS

Abstract. A central object of study in optimal stopping theory is the single-choice prophet inequality for independent and identically distributed random variables: given a sequence of random variables X_1, \dots, X_n drawn independently from the same distribution, the goal is to choose a stopping time τ such that for the maximum value of α and for all distributions, $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max_t X_t]$. What makes this problem challenging is that the decision whether $\tau = t$ may only depend on the values of the random variables X_1, \dots, X_t and on the distribution F . For a long time, the best known bound for the problem had been $\alpha \geq 1 - 1/e \approx 0.632$, but recently a tight bound of $\alpha \approx 0.745$ was obtained. The case where F is unknown, such that the decision whether $\tau = t$ may depend only on the values of the random variables X_1, \dots, X_t , is equally well motivated but has received much less attention. A straightforward guarantee for this case of $\alpha \geq 1/e \approx 0.368$ can be derived from the well-known optimal solution to the secretary problem, where an arbitrary set of values arrive in random order and the goal is to maximize the probability of selecting the largest value. We show that this bound is in fact tight. We then investigate the case where the stopping time may additionally depend on a limited number of samples from F , and we show that, even with $o(n)$ samples, $\alpha \leq 1/e$. On the other hand, n samples allow for a significant improvement, whereas $O(n^2)$ samples are equivalent to knowledge of the distribution: specifically, with n samples, $\alpha \geq 1 - 1/e \approx 0.632$ and $\alpha \leq \ln(2) \approx 0.693$, and with $O(n^2)$ samples, $\alpha \geq 0.745 - \varepsilon$ for any $\varepsilon > 0$.

Funding: This work was funded in part by an Amazon Research Award, by the Deutscher Akademischer Austauschdienst [Grant PRIME], by the Comisión Nacional de Investigación Científica y Tecnológica [Grant AFB-170001], by the Fondo de Fomento al Desarrollo Científico y Tecnológico [Grant 190043], by the European Research Council [Grant 691672], and by the Engineering and Physical Sciences Research Council [Grant EP/T015187/1].

Keywords: optimal stopping • prophet inequalities • posted pricing • online algorithms

1. Introduction

The theory of optimal stopping is concerned with sequential decision making given imperfect information about the future, in order to maximize a reward or minimize a 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 theoretical computer science and operations research, at least in part due to their relevance to the design of posted-price mechanisms for online sales.

In the *secretary problem*, we are given n distinct, nonnegative 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 select the maximum. The problem has a surprisingly simple, and surprisingly positive, answer: by discarding a $1/e$ fraction of the numbers, and then selecting the first number that is greater than any of the discarded numbers, one is guaranteed to select the maximum with probability $1/e$ (e.g., Gilbert and Mosteller [29]). The guarantee of $1/e$ achieved by this simple stopping rule is best possible, and it 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 [8], Ferguson [28]). When numbers are independent and identically distributed (i.i.d.) from a known distribution, a better guarantee of around 0.58 can be obtained (Gilbert and Mosteller [29]), and this bound is again tight.

In the *prophet problem*, we are again shown n nonnegative 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 the sequence. Two central results for this problem 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 [38], [39] and Samuel-Cahn [44]. For the latter, a lower bound of $1 - 1/e \approx 0.632$ due to Hill and Kertz [33], corresponding to a stopping rule with this guarantee, was improved only very recently, first to around 0.738 by Abolhassani et al. [1] and then to around 0.745 by Correa et al. [15]. The lower bound of Correa et al. [15] is in fact known to be tight due to an impossibility result of Hill and Kertz [33] and Kertz [36] that implies a matching upper bound.

A natural variant of the prophet problem, for both identical and nonidentical distributions, can be obtained if we assume that the distributions from which values are drawn are unknown. Despite its obvious appeal, which was noted for example by Azar et al. [6], precious little is known about this variant.

1.1. Our Contribution

We consider the prophet problem in which values are drawn independently from a *single unknown* distribution, and we ask which approximation guarantees can be obtained relative to the expected maximum value in hindsight. It is worth pointing out that, in contrast to the case where the distribution is known and an optimal stopping rule can be obtained via backward induction, there is no clear candidate for an optimal stopping rule. The case of identical distributions seems particularly interesting, as here one may hope to be able to learn something about later values from earlier ones.

A guarantee of $1/e$ for the problem can be obtained in a relatively straightforward way from the well-known optimal stopping rule for the secretary problem (see Theorem 1 in Section 3). The rule is guaranteed to stop on the maximum value with probability at least $1/e$, and one can show that this implies a $1/e$ -approximation relative to the expected maximum. Note that such an analysis, however, does not take into account that all values come from the same distribution and thus ignores any possibility of the aforementioned learning.

We show that no learning of the distribution is possible and that the straightforward guarantee of $1/e$ is in fact best possible in the prophet setting (see Theorem 2 in Section 3). The main difficulty in showing an impossibility result of this kind is that the set of stopping rules to which it applies is very rich. We will see, however, that for every stopping rule there exists a set $V \subseteq \mathbb{N}$ of arbitrary size and with an arbitrary gap between the largest and second-largest element on which the stopping rule is what we call *value-oblivious*: for random variables X_1, \dots, X_n supported on V , the decision to stop at X_i when $X_i > \max\{X_1, \dots, X_{i-1}\}$ does not depend on the values of the random variables X_1, \dots, X_i but depends only on whether X_i is the largest among these values. We will then construct a distribution F with support V such that n values drawn independently from F are pairwise distinct with probability arbitrarily close to 1 and the expectation of their maximum is dominated by the largest value in V . The objective of the prophet problem on F is thus identical, up to a small error, to that of the secretary problem, and any stopping rule with a guarantee better than $1/e$ for the former would yield such a stopping rule for the latter. To understand why stopping rules must be value-oblivious, it is useful to consider the special case where $n = 2$. In this case, we may focus on rules that always stop at X_2 whenever they have not stopped at X_1 , and every such stopping rule can be described by a function $p: \mathbb{R} \rightarrow [0, 1]$ such that $p(x)$ is the probability of stopping at X_1 when $X_1 = x$. By the Bolzano–Weierstrass theorem, the infinite sequence $(p(n))_{n \in \mathbb{N}}$ contains a monotone subsequence and thus, for some $q \in [0, 1]$ and every $\varepsilon > 0$, a subsequence of values contained in the interval $[q - \varepsilon, q + \varepsilon]$. For random variables that only take values in the index set of that latter subsequence, the stopping rule will therefore stop at the first random variable with what is essentially a fixed probability. When $n > 2$, the set of possible stopping rules becomes much richer, and identifying a set V on which a particular stopping rule is value-oblivious becomes much more challenging. Rather than the Bolzano–Weierstrass theorem, our proof uses the infinite version of Ramsey’s theorem [40] to establish the existence of such a set.

Motivated by this impossibility result, we then turn to the case where the stopping rule has access to a limited number of additional samples from the distribution, which it may use in deciding when to stop. An extension of our impossibility result shows that $o(n)$ samples are not enough to improve on the bound of $1/e$. The interesting case, therefore, is the one with $\Omega(n)$ samples, and we show that a simple stopping rule achieves a guarantee of $1 - 1/e \approx 0.632$ with $n - 1$ samples (see Theorem 3 in Section 4). The rule starts by drawing $n - 1$ samples. Then, when considering the i th random variable for $i \geq 1$, it also considers a random subset of size $n - 1$ drawn uniformly from the $n - 1$ initial samples and the $i - 1$ random variables seen so far. If the i th random variable is greater than the maximum of that random subset, then the rule stops; otherwise, it continues with the next random variable. Whereas the stopping rule itself is easy to describe, its analysis relies on an insight that is somewhat subtle. Indeed, each of the sets of random variables used to set a threshold for acceptance is distributed like

a set of $n - 1$ fresh samples from the distribution. The expected value collected from each random variable, conditioned on its acceptance, thus equals the expected maximum value of n independent draws from the distribution, and the probability of accepting a random variable conditioned on reaching it is exactly $1/n$. The approximation guarantee is then equal to the overall probability of stopping, which is at least $1 - 1/e$. By a straightforward extension, Corollary 2 in Section 4, we obtain a lower bound of $\frac{1+\gamma}{2}(1 - 1/e)$ on the guarantee achievable with γn samples for any $\gamma \in [0, 1]$.

We complement the lower bound of $1 - 1/e$ with matching upper bounds for two different classes of stopping rules that share specific properties of the aforementioned stopping rule. These bounds limit the types of approaches that could conceivably be used to go beyond a performance guarantee of $1 - 1/e$. We then give a parametric upper bound that applies to any stopping rule with access to γn samples for $\gamma \geq 0$, see Theorem 4 in Section 4. For rules that use at most n samples this upper bound is equal to $\ln(2) \approx 0.693$ and thus nearly tight.

We finally show that $O(n^2)$ samples are enough to get arbitrarily close to the optimal guarantee of around 0.745 that is attainable when the distribution is known (see Theorem 5 in Section 5). This is achieved by mimicking the stopping rule that attains that bound, which uses a decreasing sequence of thresholds corresponding to conditional acceptance probabilities that increase over time, but using quantiles of the empirical distribution rather than the actual one. By discarding a constant initial fraction of the values, and using the inequality of Dvoretzky et al. [21] to show simultaneous concentration of all empirical quantiles, we reduce the number of required samples from $O(n^4)$ to $O(n^2)$ relative to the obvious approach that potentially stops on any of the values and uses Chernoff and union bounds to show concentration.

Taken together, our results reveal a phase transition from secretary-like to prophet-like behavior when going from $o(n)$ to $\Omega(n)$ samples and show that $O(n^2)$ samples are equivalent to full knowledge of the distribution. We give an overview of the results in Figure 1.

1.2. Follow-Up Work

Rubinstein et al. [43] subsequently showed that a guarantee arbitrarily close to 0.745 can in fact already be achieved with $O(n)$ samples. Kaplan et al. [35] and Correa et al. [14] studied generalizations of the problem that we consider here. The results of Kaplan et al. imply a lower bound for our problem with γn samples of $e^{-e^{-\gamma}}$ when $0 \leq \gamma \leq 0.567$ and of around $\gamma(1 - \gamma - e^{-\gamma})$ when $0.567 \leq \gamma \leq 1$. This improves on our lower bound when $\gamma < 1$ and matches our lower bound when $\gamma = 1$. The results of Correa et al., on the other hand, imply an improved lower bound of 0.635 when $\gamma = 1$, that is, for a situation with n samples.

1.3. Further Related Work

For early work on the classic single-choice prophet inequality in mathematics, the reader is referred to a survey of Hill and Kertz [34]. Starting from the work of Hajiaghayi et al. [32], the prophet problem and extensions to richer feasibility conditions have seen a surge of interest in theoretical computer science. This has produced prophet inequalities for matroids and polymatroids (e.g., Alaei [2], Anari et al. [4], Chawla et al. [12], Dütting and Kesselheim [18], Feldman et al. [26], Kleinberg and Weinberg [37], Rubinstein and Singla [42]), settings where feasible solutions are given by an arbitrary downward-closed set system (e.g., Rubinstein [41], Rubinstein and Singla [42]), matching problems (e.g., Chawla et al. [12], Ezra et al. [24], Gravin and Wang [31], Kleinberg and Weinberg [37]), knapsack constraints (e.g., Dütting et al. [20], Feldman et al. [26]), resource-allocation problems involving intervals and paths (e.g., Chawla et al. [11]), and combinatorial auctions (e.g., Feldman et al. [25]; Dütting et al. [19], [20]).

There also exists a relatively small but important body of prior work on the case of unknown distributions. Most relevant for us is the aforementioned work by Azar et al. [6], which focuses on richer feasibility structures such as matchings and matroids, and the work by Babaioff et al. [7], who consider a setting similar to ours but focus on a different objective (revenue maximization), apply different techniques, and obtain results that are qualitatively different from ours. Independently from our work, Rubinstein et al. [43] studied the case of random variables that are drawn independently from unknown *nonidentical* distributions and showed that a single sample from each distribution is enough to achieve a guarantee of $1/2$, which matches the best possible guarantee that can be achieved when the distributions are known.

In another variant of the prophet problem, the so-called prophet secretary problem, random variables are drawn from known nonidentical distributions and observed in random order. The goal is again to immediately and irrevocably choose random variables with large value. Esfandiari et al. [23] gave a lower bound of $1 - 1/e$ for the version of the problem where a single random variable must be chosen, which was subsequently improved to $1 - 1/e + 1/400$ by Azar et al. [5]. The best bounds currently known for the single-choice version are a lower bound of $1 - 1/e + 1/27 \approx 0.669$ and an upper bound of $\sqrt{3} - 1 \approx 0.732$ due to Correa et al. [13]. Combinatorial

versions of the prophet secretary problem have been studied by Ehsani et al. [22]. The aforementioned bounds of Kaplan et al. [35] and Correa et al. [14] also apply to a single-choice prophet secretary problem with unknown distributions and additional samples.

Optimal stopping with unknown distributions has also been studied in operations research and management science, but the types of problems, objectives, and techniques differ significantly from ours and typically involve regret minimization (see, e.g., the results of Goldenshluger and Zeevi [30] and the recent survey of den Boer [16]). The literature in operations research and management science, moreover, contains results on a broad range of stochastic optimization problems, which share certain features of the basic prophet inequality problem and its combinatorial extensions. These problems include constrained Bayesian online selection (e.g., Anari et al. [4]), Bayesian assortment optimization (e.g., Feng et al. [27]), and certain models of network revenue management (e.g., Alijani et al. [3]). To the best of our knowledge, they have not been studied from the perspective of sample complexity.

A final line of related work, in economics and theoretical computer science, has studied posted pricing and prophet inequalities with inaccurate priors (Bergemann and Schlag [9], [10]; Dütting and Kesselheim [17]). This line of work assumes access to prior distributions that are close in terms of some metric to the actual distributions and seeks either max-min optimal mechanisms or performance guarantees that are parameterized by the distance between the assumed and actual priors.

2. Preliminaries

Denote by \mathbb{N} the set of positive integers, and let $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. For $i \in \mathbb{N}$, let $[i] = \{1, \dots, i\}$, and denote by \mathcal{S}_i the set of permutations of $[i]$.

Let $k \in \mathbb{N}_0$ and $n \in \mathbb{N}$. We consider (k, n) -stopping rules that sequentially observe random variables X_1, \dots, X_n and have access to samples S_1, \dots, S_k and for each $i = 1, \dots, n$ decide whether to stop on X_i based on the values of X_1, \dots, X_i and S_1, \dots, S_k . We assume that X_1, \dots, X_n and S_1, \dots, S_k are i.i.d., and respectively denote by f and F the probability density function and cumulative distribution function of their distribution. Formally, a (k, n) -stopping rule \mathbf{r} is a family of functions r_1, \dots, r_n where $r_i : \mathbb{R}_+^{k+i} \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Here, $r_i(s_1, \dots, s_k, x_1, \dots, x_i)$ for $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^i$ is the probability of stopping at X_i conditioned on having received $S_1 = s_1, \dots, S_k = s_k$ as samples and $X_1 = x_1, \dots, X_i = x_i$ as values and not having stopped on any of X_1, \dots, X_{i-1} . The stopping time τ of a (k, n) -stopping rule \mathbf{r} , given S_1, \dots, S_k and X_1, \dots, X_n , is thus the random variable with support $\{1, \dots, n\} \cup \{\infty\}$ such that for all $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$,

$$\Pr[\tau = i \mid S_1 = s_1, \dots, S_k = s_k, X_1 = x_1, \dots, X_n = x_n] = \left(\prod_{j=1}^{i-1} (1 - r_j(s_1, \dots, s_k, x_1, \dots, x_j)) \right) \cdot r_i(s_1, \dots, s_k, x_1, \dots, x_i).$$

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 we will measure its performance relative to the expected maximum $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ of the random variables X_1, \dots, X_n . We will say that a stopping rule achieves an approximation guarantee α , for $\alpha \leq 1$, if, for any distribution, $\mathbb{E}[X_\tau] \geq \alpha \mathbb{E}[\max\{X_1, \dots, X_n\}]$.

For ease of exposition, we will assume continuity of F in proving lower bounds and mainly 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.

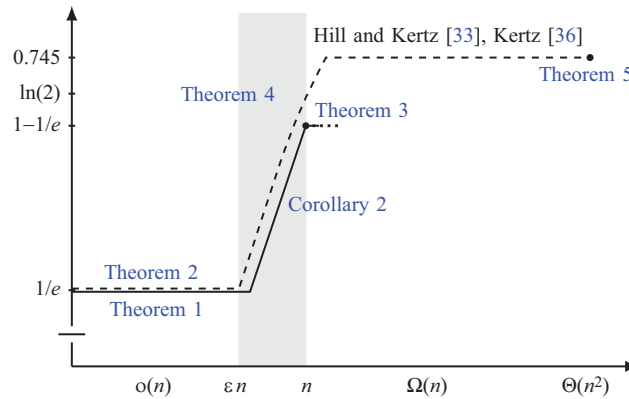
3. Sublinear Number of Samples

We begin by showing that, for $o(n)$ samples, the prophet problem with an unknown distribution behaves like the secretary problem. As we will see in Section 3.1, a straightforward baseline can be obtained from the optimal solution to the secretary problem, which discards a $1/e$ fraction of the values and then accepts the first value that exceeds the maximum of the discarded values. The algorithm does not require any samples, 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. Our main result in this section, which we prove in Section 3.2, shows that the bound of $1/e$ is in fact best possible. This result continues to hold with $o(n)$ samples.

3.1. A $1/e$ -Approximation Without Samples

The following result translates the guarantee of $1/e$ for the secretary problem to a prophet inequality for independent random variables from an unknown distribution.

Figure 1. (Color online) Overview of results. The number of samples is displayed along the horizontal axis, the performance guarantee along the vertical axis. Lower bounds, shown as a solid line and two dots, result from stopping rules with a certain performance guarantee. Upper bounds, shown as dashed lines, correspond to impossibility results that no stopping rule can improve upon. The results for $o(n)$ and $\Theta(n^2)$ samples are tight. With the exception of the upper bound of approximately 0.745, all results are new to this paper.



Theorem 1. Let X_1, X_2, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then there exists a $(0, n)$ -stopping rule with stopping time τ such that

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

The result can be shown in a straightforward way, based on the idea that the realizations of the random variables X_1, \dots, X_n can be obtained by drawing n values from their common distribution and then permuting them uniformly at random. The classic analysis of the secretary problem (Ferguson [28]) implies that for each realization of the n draws, the optimal stopping rule for this problem obtains the maximum value with probability $1/e$. It thus also obtains at least a $1/e$ fraction of the expected value of this maximum. We formalize this idea and prove Theorem 1 in Appendix A.

3.2. A Matching Upper Bound

We next show that it is impossible to improve on the straightforward lower bound of $1/e$.

Theorem 2. Let $\delta > 0$. Then there exists $n_0 \in \mathbb{N}$ such that, for any $n \geq n_0$ and any $(0, n)$ -stopping rule with stopping time τ , there exists a distribution F , not known to the stopping rule, such that when X_1, \dots, X_n are i.i.d. random variables drawn from F ,

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

The main difficulty in showing an impossibility result of this kind is that it applies to the set of all possible $(0, n)$ -stopping rules, which a priori is very rich. Indeed, recall that a $(0, n)$ -stopping rule \mathbf{r} can be any family of functions r_1, \dots, r_n , where $r_i: \mathbb{R}_+^i \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Our main structural insight will be that we can restrict attention to a much simpler class of stopping rules \mathbf{r} that are in a certain sense oblivious to the values of the random variables that they observe. For random variables X_1, \dots, X_n supported on arbitrarily large sets $V \subseteq \mathbb{N}$, under the condition that X_1, \dots, X_i are pairwise distinct and $X_i > \max\{X_1, \dots, X_{i-1}\}$, and up to an arbitrarily small error ε , the probability that \mathbf{r} stops on X_i will not depend on the values of any of the random variables X_1, \dots, X_i . This is made precise by the following definition. Although it is not needed for proving Theorem 2, we will consider the more general case of (k, n) -stopping rules for any $k \in \mathbb{N}_0$. The structural result extends easily to the more general case, and we will use it later to generalize Theorem 2.

Definition 1. Let $\varepsilon > 0$, $k \in \mathbb{N}_0$, and let $V \subseteq \mathbb{N}$. A (k, n) -stopping rule \mathbf{r} is an ε -value-oblivious on V if, for all $i \in [n]$, there exists $q_i \in [0, 1]$ such that, for all pairwise distinct $s_1, \dots, s_k, v_1, \dots, v_i \in V$ with $v_i > \max\{s_1, \dots, s_k, v_1, \dots, v_{i-1}\}$, it holds that $r_i(s_1, \dots, s_k, v_1, \dots, v_i) \in [q_i - \varepsilon, q_i + \varepsilon]$.

Whereas value-obliviousness significantly restricts the expressiveness of a stopping rule, this restriction turns out to be essentially without loss when it comes to the ability of achieving a certain guarantee across all possible

distributions: for any stopping rule and any $\varepsilon > 0$, there exists a stopping rule with the same guarantee that is ε -value-oblivious for some infinite set $V \subseteq \mathbb{N}$. This is made precise by the following lemma, which we prove in Section 3.3.

Lemma 1. *Let $\varepsilon > 0$ and $k \in \mathbb{N}_0$. If there exists a (k, n) -stopping rule with guarantee α , then there exists a (k, n) -stopping rule \mathbf{r} with guarantee α and an infinite set $V \subseteq \mathbb{N}$ such that \mathbf{r} is ε -value-oblivious on V .*

With Lemma 1 at hand, it is not difficult to prove Theorem 2. For any $(0, n)$ -stopping rule and an appropriate value of ε , we identify a $(0, n)$ -stopping rule \mathbf{r} with the same performance guarantee that is ε -value-oblivious on an infinite set $V \subseteq \mathbb{N}$. We then define a distribution F with finite support $S \subseteq V$ such that (i) there is a large gap between the largest and second-largest elements of S , (ii) n independent draws from F are pairwise distinct with probability close to 1, (iii) \mathbf{r} is ε -value-oblivious on S , and (iv) the performance guarantee of \mathbf{r} on the distribution is dominated by the probability of selecting the largest element of S . By (i) and (ii), the prophet problem for the unknown distribution F is then equivalent up to a small error to a secretary problem, and by (iii) and (iv), \mathbf{r} behaves on F essentially like a stopping rule for the secretary problem. A performance guarantee for \mathbf{r} of more than $1/e$ would thus contradict the optimality of this bound for the secretary problem.

Proof of Theorem 2. It suffices to show that the guarantee of any $(0, n)$ -stopping rule is bounded from above by $1/e + o(1)$, where, implicitly, $n \rightarrow \infty$.

To this end, consider an arbitrary $(0, n)$ -stopping rule with guarantee α . Let $\varepsilon = 1/n^2$. By Lemma 1, there then exists a $(0, n)$ -stopping rule \mathbf{r} with guarantee α and an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value-oblivious. Denote by τ the stopping time of \mathbf{r} . 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

$$X_i = \begin{cases} v_1 & \text{with probability } \frac{1}{n^3} \left(1 - \frac{1}{n^2}\right), \\ \vdots & \\ v_{n^3} & \text{with probability } \frac{1}{n^3} \left(1 - \frac{1}{n^2}\right), \\ u & \text{with probability } \frac{1}{n^2}. \end{cases}$$

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

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

On the other hand,

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

where for the first inequality we have applied the law of total expectation to $\mathbb{E}[X_\tau \mid X_{(n)} = u \wedge X_{(n-1)} \neq u]$ to additionally condition on whether $X_\tau = X_{(n)}$, and for the last inequality we have applied the law of total probability to $\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_{(n-1)} \neq u]$ to additionally condition on whether X_1, \dots, X_n are distinct.

Given (1) and (2), to show that $\alpha \leq 1/e + o(1)$, it now suffices to show that

$$\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ distinct}] \leq 1/e + o(1). \tag{3}$$

Note that in the event where $X_{(n)} = u$ and X_1, \dots, X_n are pairwise distinct, the relative ranks of X_1, \dots, X_n are distributed uniformly at random. For a 0-value-oblivious stopping rule $\hat{\mathbf{r}}$ with stopping time $\hat{\tau}$, it thus follows from the well-known optimal solution to the secretary problem (Ferguson [28, section 2]) that

$$\Pr[X_{\hat{\tau}} = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ distinct}] \leq 1/e + o(1). \tag{4}$$

To show this claim for stopping rule \mathbf{r} , which is only ε -value-oblivious with $\varepsilon > 0$, we construct from \mathbf{r} a 0-value-oblivious stopping rule $\hat{\mathbf{r}}$ and show through a coupling argument that the probability that \mathbf{r} stops at $X_{(n)}$ is bounded by the probability that $\hat{\mathbf{r}}$ stops at $X_{(n)}$ plus $n\varepsilon = 1/n$.

Since \mathbf{r} is ε -value-oblivious on V , $r_i(s_1, \dots, s_i) \in [q_i - \varepsilon, q_i + \varepsilon]$ for all $i \in [n]$, some $q_i \in [0, 1]$, and all distinct $s_1, \dots, s_i \in V$ with $s_i > \max\{s_1, \dots, s_{i-1}\}$. Let $\hat{\mathbf{r}}$ be the stopping rule such that for all $s_1, \dots, s_i \in V$, $\hat{r}_i(s_1, \dots, s_i) = q_i$ if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $\hat{r}_i(s_1, \dots, s_i) = 0$ otherwise. Denote by $\hat{\tau}$ the stopping time of $\hat{\mathbf{r}}$.

Let $s_1, \dots, s_n \in V$ be distinct and assume that $X_1 = s_1, \dots, X_n = s_n$. To compare the performance of \mathbf{r} and $\hat{\mathbf{r}}$, we can view τ and $\hat{\tau}$ as being coupled via n independent draws c_1, \dots, c_n from the uniform distribution on $[0, 1]$. For every $i \in [n]$, and under the condition that $\tau \geq i$, we can assume that $\tau = i$ if and only if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $r_i(s_1, \dots, s_i) > c_i$. Similarly, for every $i \in [n]$, and under the condition that $\hat{\tau} \geq i$, we can assume that $\hat{\tau} = i$ if and only if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $\hat{r}_i(s_1, \dots, s_i) = q_i > c_i$. For $i \in [n]$, let ξ_i be the event that occurs if and only if $s_i > \max\{s_1, \dots, s_{i-1}\}$ and $c_i \in [\min\{r_i(s_1, \dots, s_i), \hat{r}_i(s_1, \dots, s_i)\}, \max\{r_i(s_1, \dots, s_i), \hat{r}_i(s_1, \dots, s_i)\}]$. Then $\Pr[\xi_i] \leq \varepsilon$, while $X_\tau = X_{(n)} \neq X_{\hat{\tau}}$ requires ξ_i to occur for some $i \in [n]$. Thus, by the union bound,

$$\Pr[X_\tau = X_{(n)} \mid X_1 = s_1, \dots, X_n = s_n] \leq \Pr[X_{\hat{\tau}} = X_{(n)} \mid X_1 = s_1, \dots, X_n = s_n] + n\varepsilon.$$

Because this statement holds pointwise for all distinct $s_1, \dots, s_n \in V$,

$$\Pr[X_\tau = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ distinct}] \leq \Pr[X_{\hat{\tau}} = X_{(n)} \mid X_{(n)} = u \wedge X_1, \dots, X_n \text{ distinct}] + n\varepsilon. \quad (5)$$

Substituting (5) into (4) yields (3), which completes the proof. \square

3.3. Proof of Lemma 1

The lemma claims that, for every $\varepsilon > 0$, the existence of a (k, n) -stopping rule with performance guarantee α implies the existence of a (k, n) -stopping rule \mathbf{r} with the same performance guarantee and the existence of an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value oblivious. Since we can interpret a (k, n) -stopping rule as a $(0, n')$ -stopping rule with $n' = k + n$ that never stops on the first k values, it will be sufficient to consider $(0, n)$ -stopping rules with this additional constraint. We prove the lemma through a sequence of steps that successively restrict the expressiveness of the stopping rules that we have to consider. First, we show a restriction to what we call *order-oblivious rules*, which in the decision to stop at random variable X_i , and conditioned on having reached X_i , may take into account the values of random variables X_1, \dots, X_{i-1} but not the order in which they were observed.

Definition 2. A $(0, n)$ -stopping rule \mathbf{r} is order-oblivious if for all $j \in [n]$, all pairwise distinct $v_1, \dots, v_j \in \mathbb{R}_+$, and all permutations $\pi \in \mathcal{S}_{j-1}$, $r_i(v_1, \dots, v_j) = r_i(v_{\pi(1)}, \dots, v_{\pi(j-1)}, v_j)$.

The following result is rather intuitive.

Lemma 2. If there exists a $(0, n)$ -stopping rule \mathbf{r} with guarantee α , then there exists a $(0, n)$ -stopping rule \mathbf{r}' with guarantee α that is order-oblivious and that, for any $i \in [n]$, never selects X_i if \mathbf{r} never selects X_i .

A naive attempt to prove this lemma would be to construct an order-oblivious stopping rule from an arbitrary stopping rule \mathbf{r} by permuting X_1, \dots, X_{i-1} uniformly at random upon observing X_i and accepting X_i if and only if \mathbf{r} would accept it under the random permutation. The resulting stopping rule may, however, have a different guarantee than \mathbf{r} , because the probability that \mathbf{r} arrives at X_i may vary depending on the permutation. Some additional care is therefore required.

Proof of Lemma 2. For $i \in [n]$, let \sim_i be the equivalence relation on \mathbb{R}_+^i such that $(v_1, \dots, v_i) \sim_i (w_1, \dots, w_i)$ if v_1, \dots, v_{i-1} is a permutation of w_1, \dots, w_{i-1} and $v_i = w_i$. Note that a stopping rule \mathbf{r} with stopping time τ is order-oblivious if and only if, for all $i \in [n]$ and $v_1, \dots, v_i, w_1, \dots, w_i \in \mathbb{R}_+$, it holds that $r_i(v_1, \dots, v_i) = r_i(w_1, \dots, w_i)$ whenever $(v_1, \dots, v_i) \sim_i (w_1, \dots, w_i)$. We will refer to the equivalence classes of \sim_i as *states*, and we will say that \mathbf{r} arrives at $s \in \mathbb{R}_+^i / \sim_i$ in the event that $\tau \geq i$ and $X_1 = v_1, \dots, X_{i-1} = v_{i-1}$, where $[v_1, \dots, v_i]_{\sim_i} = s$.

Let \mathbf{r} be an arbitrary stopping rule with stopping time τ , and define a stopping rule \mathbf{r}' with stopping time τ' such that $r'_1(v_1) = r_1(v_1)$ and, for all $i \in \{2, \dots, n\}$ and $v_1, \dots, v_i \in \mathbb{R}_+$ with $\Pr[\mathbf{r} \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] > 0$,

$$r'_i(v_1, \dots, v_i) = \Pr[\tau = i \mid \mathbf{r} \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}].$$

Because $[v_1, \dots, v_i]_{\sim_i}$ is invariant under permutations of the sequence v_1, \dots, v_{i-1} , \mathbf{r}' is indeed order-oblivious. Moreover, for any $i \in [n]$, if \mathbf{r} never selects X_i , then neither does \mathbf{r}' . It remains to be shown that \mathbf{r}' provides guarantee α .

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

$$\Pr[\mathbf{r} \text{ arrives at } s] = \Pr[\mathbf{r}' \text{ arrives at } s]. \tag{6}$$

This holds trivially for $i = 1$, so we assume that it holds for $i = k - 1 \geq 1$ and show then that it holds for $i = k$. For any $v_1, \dots, v_k \in \mathbb{R}_+$ and $s = [v_1, \dots, v_k]_{\sim_k}$, we write $v_{-j} = (v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_{k-1})$ for the sequence of length $k - 2$ in which v_j has been left out and $(v_{-j}, v_j) = (v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_{k-1}, v_j)$ for the sequence of length $(k - 1)$ obtained by appending v_j to v_{-j} . Then,

$$\begin{aligned} \Pr[\mathbf{r} \text{ arrives at } s] &= \sum_{j=1}^{k-1} \Pr[\mathbf{r} \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[\tau \neq i \mid \mathbf{r} \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[X_k = v_k] \\ &= \sum_{j=1}^{k-1} \Pr[\mathbf{r}' \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[\tau' \neq i \mid \mathbf{r}' \text{ arrives at } [v_{-j}, v_j]_{\sim_{k-1}}] \cdot \Pr[X_k = v_k] \\ &= \Pr[\mathbf{r}' \text{ arrives at } s], \end{aligned}$$

where the first and last equalities hold by definition of \sim_{k-1} and the second equality holds by the induction hypothesis and by definition of \mathbf{r}' .

We now claim that

$$\begin{aligned} \mathbb{E}[X_\tau] &= \sum_{i=1}^n \mathbb{E}[X_i \mid \tau = i] \cdot \Pr[\tau = i] \\ &= \sum_{i=1}^n \int_0^\infty \dots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \\ &\quad \cdot \frac{1}{(i-1)!} \cdot \sum_{\pi \in S_{i-1}} \Pr[\tau = i \mid X_1 = v_{\pi(1)}, \dots, X_{i-1} = v_{\pi(i-1)}, X_i = v_i] \, dv_1 \dots dv_i \\ &= \sum_{i=1}^n \int_0^\infty \dots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \cdot \Pr[\tau = i \mid \mathbf{r} \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \\ &\quad \cdot \Pr[\mathbf{r} \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \, dv_1 \dots dv_i \\ &= \sum_{i=1}^n \int_0^\infty \dots \int_0^\infty \prod_{j=1}^i f(v_j) \cdot v_i \cdot r'_i(v_1, \dots, v_i) \cdot \Pr[\mathbf{r}' \text{ arrives at } [v_1, \dots, v_i]_{\sim_i}] \, dv_1 \dots dv_i \\ &= \sum_{i=1}^n \mathbb{E}[X_i \mid \tau' = i] \cdot \Pr[\tau' = i] = \mathbb{E}[X_{\tau'}]. \end{aligned}$$

Indeed, the second equality can be seen to hold by imagining that X_1, \dots, 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 \mathbf{r}' and by (6). This completes the proof. \square

To further restrict the class of stopping rules from order-oblivious to value-oblivious ones, we will now construct, for every order-oblivious rule \mathbf{r} and any $\varepsilon > 0$, an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{r} is ε -value-oblivious. The set V will depend on \mathbf{r} and will be obtained by starting from \mathbb{N} and identifying smaller and smaller subsets on which the behavior of \mathbf{r} is more and more limited. By induction on $i \in [n]$, we will identify a set on which value-obliviousness holds with respect to the i th random variable. We need the following definition.

Definition 3. Consider a $(0, n)$ -stopping rule \mathbf{r} . Let $\varepsilon > 0$, $i \in [n]$, and $V \subseteq \mathbb{N}$. Then \mathbf{r} is (ε, i) -value-oblivious on V if there exists $q \in [0, 1]$ such that, for all pairwise distinct $v_1, \dots, v_i \in V$ with $v_i > \max\{v_1, \dots, v_{i-1}\}$, it holds that $r_i(v_1, \dots, v_i) \in [q - \varepsilon, q + \varepsilon]$.

Note that (ε, i) -value-obliviousness for all $i \in [n]$ is equivalent to ε -value-obliviousness. In establishing (ε, i) -value-obliviousness for a particular value of i , we will appeal to the infinite version of Ramsey’s theorem to show the existence of an appropriate set V .

Lemma 3 (Ramsey [40, Theorem A]). *Let $c, d \in \mathbb{N}$, and let H be an infinite complete d -uniform hypergraph whose hyperedges are colored with c colors. Then there exists an infinite complete d -uniform subhypergraph of H that is monochromatic.*

Proof of Lemma 1. Suppose that there exists a (k, n) -stopping rule with guarantee α . By interpreting this rule as a $(0, n')$ -stopping rule with $n' = k + n$, and by Lemma 2, there then exists a $(0, n')$ -stopping rule \mathbf{r} that is

order-oblivious and never stops on X_1, \dots, X_k . 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 $i \in [j]$, \mathbf{r} is (ε, i) -value-oblivious on S^j .

For $j = n'$, this implies that the stopping rule \mathbf{r} is (ε, j) -value-oblivious on S^n for all $j \in \mathbb{N}$ and hence is ε -value-oblivious on S^n . The claim then follows by reinterpreting \mathbf{r} as a (k, n) -stopping rule, which is possible, because it never stops on X_1, \dots, X_k .

It is clear that $S^0 = \mathbb{N}$ clearly satisfies the claim for $j = 0$, and we proceed to show the claim for $j = \ell > 0$ assuming that it holds for $j < \ell$.

Note that it suffices to find an infinite set $S^\ell \subseteq S^{\ell-1}$ such that \mathbf{r} is (ε, ℓ) -value-oblivious on S^ℓ , as the induction hypothesis then implies (ε, i) -value-obliviousness on S^ℓ as a subset of S^i for all $i \in [\ell - 1]$.

Toward the application of Lemma 3, we construct a complete ℓ -uniform hypergraph H with vertex set $S^{\ell-1}$. Consider any set $\{v_1, \dots, v_\ell\} \subseteq S^{\ell-1}$ of cardinality ℓ such that $v_\ell > \max\{v_1, \dots, v_{\ell-1}\}$. Note that there exists a unique $u \in \{1, 2, \dots, \lceil 1/(2\varepsilon) \rceil\}$ such that $r_\ell(v_1, \dots, v_\ell) \in [(2u - 1) \cdot \varepsilon - \varepsilon, (2u - 1) \cdot \varepsilon + \varepsilon)$. Color the hyperedge $\{v_1, \dots, v_\ell\}$ of H with color u .

By Lemma 3 with $c = \lceil 1/2\varepsilon \rceil$ and $d = \ell$, there exists an infinite set of vertices that induces a complete monochromatic subhypergraph of H . Define S^ℓ to be such a set inducing a monochromatic subhypergraph of H with color u . Set $q = (2u - 1)\varepsilon$, and consider distinct $v_1, \dots, v_\ell \in S^\ell$ with $v_\ell > \max\{v_1, \dots, v_{\ell-1}\}$. Since the edge $\{v_1, \dots, v_\ell\}$ in H has color u , $r_\ell(v_{\pi(1)}, \dots, v_{\pi(\ell-1)}, v_\ell) \in [q - \varepsilon, q + \varepsilon)$ for some permutation $\pi \in \mathcal{S}_{\ell-1}$. But since \mathbf{r} is order-oblivious, also $r_\ell(v_1, \dots, v_{\ell-1}, v_\ell) \in [q - \varepsilon, q + \varepsilon)$. So \mathbf{r} is (ε, ℓ) -value-oblivious on S^ℓ . This completes the induction step and the proof. \square

3.4. Extension of the Upper Bound to $o(n)$ Samples

We conclude this section by showing that, even with $o(n)$ samples, the guarantee of $1/e$ is still best possible. To gain some intuition as to why this is true, assume that there existed an $(o(n), n)$ -stopping rule \mathbf{r} with guarantee greater than $1/e$ by some constant. We could then obtain a $(0, n)$ -stopping rule \mathbf{r}' that interprets, for a suitable choice of n' , the first $o(n')$ values as samples and the following n' values as actual values on which it may stop, and then runs \mathbf{r} in this setting. If we choose $n' = (1 - o(1)) \cdot n$, then the expected maxima of n and n' draws from any distribution are identical up to a $(1 - o(1))$ factor. The guarantee of \mathbf{r} would thus carry over to \mathbf{r}' , contradicting Theorem 2.

Corollary 1. *Let $\delta > 0$ and $f : \mathbb{N} \rightarrow \mathbb{N}$ with $f(n) = o(n)$. Then there exists $n_0 \in \mathbb{N}$ such that, for any $n \geq n_0$ and any $(f(n), n)$ -stopping rule with stopping time τ , there exists a distribution F , not known to the stopping rule, such that when X_1, \dots, X_n are i.i.d. random variables drawn from F ,*

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

Proof. For $\delta > 0$, choose $\gamma > 0$ such that $(1 + \gamma)/e \leq 1/e + \delta/2$ and $\gamma/(1 + \gamma) \leq 1/e$. By Theorem 4, there exists an n_1 such that, for all $n \geq n_1$ and every $(\gamma n, n)$ -stopping rule with stopping time τ , there exists a distribution F , not known to the stopping rule, with the following property. When X_1, \dots, X_n are i.i.d. random variables drawn from F , we have

$$\mathbb{E}[X_\tau] \leq \left(\frac{1 + \gamma}{e} + \frac{\delta}{2}\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}] \leq \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where the second inequality follows by our choice of γ .

Now let n_0 be such that $f(n) \leq \gamma n$ for all $n \geq n_0$. As every $(f(n), n)$ -stopping rule can be interpreted as a $(\gamma n, n)$ -stopping rule when $n \geq n_0$, the aforementioned bound for $(\gamma n, n)$ -stopping rules applies to $(f(n), n)$ -stopping rules as well when $n \geq \max\{n_0, n_1\}$. This proves the claim. \square

4. Linear Number of Samples

The previous section has revealed a strong impossibility: even with $o(n)$ samples, it is impossible to improve over the straightforward lower bound of $1/e \approx 0.368$ achieved by the well-known optimal stopping rule for the secretary problem. We proceed to show that there is a sharp phase transition when going from $o(n)$ samples to $\Omega(n)$ samples, by giving an algorithm that uses as few as $n - 1$ samples and improves the lower bound from $1/e$ to $1 - 1/e \approx 0.632$. We also show that the bound of $1 - 1/e$ is in fact tight for two different classes of algorithms that share certain features of our algorithm. This illustrates that our analysis is tight and limits the types of approaches that could conceivably be used to go beyond $1 - 1/e$. We also show a parametric upper bound for

algorithms that use γn samples for any $\gamma \geq 0$. For algorithms that use at most n samples, this bound is equal to $\ln(2) \approx 0.693$ and thus nearly tight.

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 , accepting the first random variable that exceeds the threshold. 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 at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$, since, under this condition, X_t is the maximum of at least 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 detailed analysis of the approach is provided for completeness in Appendix B.

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. Our algorithm improves over the naive approach that obtains a factor of $1/2$ by increasing the probability that we stop at all, while maintaining the property that the expected value that we collect when we do stop is at least $\mathbb{E}[\max\{X_1, \dots, X_n\}]$.

Theorem 3. Let X_1, X_2, \dots, X_n be i.i.d. random variables from an unknown distribution F . Then there exists an $(n-1, n)$ -stopping-rule with stopping time τ such that

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

Note that a guarantee of $1 - 1/e - \varepsilon$ with $O_\varepsilon(n)$ samples follows from a result of Ehsani et al. [22] by observing that $O_\varepsilon(n)$ samples provide a sufficiently good approximation to the $1/e$ -quantile of the distribution of $\max\{X_1, \dots, X_n\}$. Here we take a different route that yields the bound *exactly* and that, more importantly, can be developed further to work with only $n-1$ samples.

Suppose that we were given access to $n(n-1) \in \Theta(n^2)$ samples. Then we could partition the $n(n-1)$ samples into n sets of size $n-1$ each and use the maximum of the i th set as a threshold for the i th 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 i th 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 mimics this approach, but, instead of using $n-1$ fresh samples for each of the n random variables, it constructs $n-1$ fresh-looking samples for each of the n random variables from a *single* set $\{S_1, \dots, S_{n-1}\}$ of $n-1$ samples.

Algorithm 1

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

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

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

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

if $X_t \geq \max S$ **then return** t

else $S \leftarrow$ random subset of size $n-1$ of $\{S_1, \dots, S_{n-1}, X_1, \dots, X_t\}$

end

return $n+1$

The algorithm starts by drawing $n-1$ samples S_1, \dots, S_{n-1} . Then, for each time step t , it compares the current random variable X_t to the maximum $\max S$ of a random subset S of size $n-1$ of the set $\{S_1, \dots, S_{n-1}, X_1, \dots, X_{t-1}\}$ containing the initial samples and the random variables seen previously. If $X_t \geq \max S$, then the algorithm accepts X_t and stops. Otherwise, it continues to the next random variable.

The key ingredient in our analysis is the following lemma, which concerns the distribution of the unordered set of values seen before step t under the condition that the algorithm has reached that step.

Lemma 4. *If Algorithm 1 arrives at step t , then the distribution of the set $\{S_1, \dots, S_{n-1}, X_1, \dots, X_{t-1}\}$ is identical to the distribution of a set of $n + t - 2$ fresh samples from F .*

Proof. We show the claim by induction on t , and we start by observing that it clearly holds for $t = 1$, as $\{S_1, \dots, S_{n-1}\}$ is a set of $n - 1$ fresh samples from F .

Now assume that the claim holds for $t = 1, \dots, t^* - 1$. Then, by the induction hypothesis and under the condition that the algorithm arrives at step $t^* - 1$, the set $\mathcal{T} = \{S_1, \dots, S_{n-1}, X_1, \dots, X_{t^*-2}\}$ has the same distribution as a set of $n + t^* - 3$ fresh samples from F . We now consider the set $\mathcal{T}' = \{S_1, \dots, S_{n-1}, X_1, \dots, X_{t^*-1}\}$, which additionally includes X_{t^*-1} , and claim that, conditioned on arriving at step t^* , this set is distributed like a set of $n + t^* - 2$ fresh samples. Note that the latter is true *before* the decision to stop or not to stop at step $t^* - 1$ is taken. To show the claim, we will argue that the decision of the algorithm to stop or not to stop at step $t^* - 1$ does not depend on the realization of the set \mathcal{T}' .

Fix any realization $\{y_1, \dots, y_{n+t^*-2}\}$ of \mathcal{T}' , and assume that the algorithm has arrived at step $t^* - 1$. Since F is continuous, we may assume that the values y_1, \dots, y_{n+t^*-2} are pairwise distinct and, without loss of generality, that $y_1 < y_2 < \dots < y_{n+t^*-2}$. Since \mathcal{T} is distributed like a set of fresh samples and X_{t^*-1} is a fresh sample, it must be the case that X_{t^*-1} is distributed uniformly over $\{y_1, \dots, y_{n+t^*-2}\}$ and the elements of \mathcal{T} are equal to the remaining values in $\{y_1, \dots, y_{n+t^*-2}\}$. The algorithm now stops at step $t^* - 1$ if $X_{t^*-1} = y_{n+t^*-2}$, and this happens with probability $1/n$. \square

Proof of Theorem 3. The value $\mathbb{E}[X_\tau]$ obtained by Algorithm 1 can be written by summing over all possible stopping times $t = 1, \dots, n$ the product of the probability of stopping at X_t and the expectation of X_t upon stopping. Writing S^t for the random subset used in step t , we thus have

$$\mathbb{E}[X_\tau] = \sum_{t=1}^n \left(\Pr[X_t \geq \max\{S^t\} \wedge X_j < \max\{S^j\} \text{ for } j < t] \cdot \mathbb{E}[X_t \mid X_t \geq \max\{S^t\} \wedge X_j < \max\{S^j\} \text{ for } j < t] \right). \quad (7)$$

For any $t \in \{1, \dots, n\}$,

$$\begin{aligned} & \Pr[X_t \geq \max\{S^t\} \wedge X_j < \max\{S^j\} \text{ for } j < t] \\ &= \Pr[X_t \geq \max\{S^t\} \mid X_j < \max\{S^j\} \text{ for } j < t] \cdot \prod_{\ell < t} \Pr[X_\ell < \max\{S^\ell\} \mid X_j < \max\{S^j\} \text{ for } j < \ell] \\ &= \frac{1}{n} \left(1 - \frac{1}{n} \right)^{t-1}, \end{aligned} \quad (8)$$

where the first equality can be obtained by repeated application of the definition of conditional probabilities and the second equality follows from Lemma 4.

Denoting by $\{T_1, \dots, T_{n-1}\}$ a set of fresh samples from F , we claim that

$$\begin{aligned} \mathbb{E}[X_t \mid X_t \geq \max\{S^t\} \wedge X_j < \max\{S^j\} \text{ for } j < t] &= \mathbb{E}[X_t \mid X_t \geq \max\{T_1, \dots, T_{n-1}\} \wedge X_j < \max\{S^j\} \text{ for } j < t] \\ &= \mathbb{E}[X_t \mid X_t \geq \max\{T_1, \dots, T_{n-1}\}] \\ &= \mathbb{E}[\max\{X_1, \dots, X_n\}]. \end{aligned} \quad (9)$$

Indeed, the first equality holds because, under the condition that $X_j < \max\{S^j\}$ for $j < t$ and by Lemma 4, S^t is distributed like $\{T_1, \dots, T_{n-1}\}$. The second equality holds because X_t itself is independent of whether $X_j < \max\{S^j\}$ for $j < t$, and the third equality holds because X_t is distributed like a fresh sample.

By substituting (8) and (9) into (7), we obtain

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

as claimed. \square

4.3. Going Beyond $1 - 1/e$

We proceed to show an upper bound of $1 - 1/e$ for two different classes of algorithms that share certain characteristics of Algorithm 1. This shows that our analysis of Algorithm 1 is tight and limits the class of algorithms that could conceivably improve on the guarantee of $1 - 1/e$.

Algorithms in the first class, upon reaching the i th random variable, stop at this random variable with a probability that is independent of i . This is the case for Algorithm 1, which, by Lemma 4, stops with probability $1/n$ upon reaching a particular random variable. The upper bound we obtain applies even in the case where the distribution F is known and to stopping rules that, like Algorithm 1, use dependent thresholds.

Proposition 1. *Let $\delta > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that, for any stopping time τ for which $\Pr[\tau = i \mid \tau > i - 1]$ is independent of i ,*

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

Proof. For $n \geq 3$ and $i \in [n]$, let

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

We begin by bounding $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ from below. For every $\varepsilon > 0$, there exists $m \in \mathbb{N}$ such that, for all $n \geq m$,

$$\begin{aligned} \mathbb{E}[\max\{X_1, \dots, X_n\}] &= \Pr\left[X_i = \frac{\sqrt{n}}{e-2} \text{ for some } i\right] \cdot \frac{\sqrt{n}}{e-2} + \\ &\quad \Pr\left[X_i < \frac{\sqrt{n}}{e-2} \text{ for all } i\right] \cdot \Pr\left[X_i = 1 \text{ for some } i \mid X_i < \frac{\sqrt{n}}{e-2} \text{ for all } i\right] \cdot 1 \\ &= \left(1 - \left(1 - \frac{1}{n^{3/2}}\right)^n\right) \cdot \frac{\sqrt{n}}{e-2} + \left(1 - \frac{1}{n^{3/2}}\right)^n \cdot \left(1 - \left(\frac{1 - \frac{1}{\sqrt{n}} - \frac{1}{n^{3/2}}}{1 - \frac{1}{n^{3/2}}}\right)^n\right) \cdot 1 \\ &\geq \frac{1}{e-2} + 1 - \varepsilon = \frac{e-1}{e-2} - \varepsilon. \end{aligned}$$

In bounding $\mathbb{E}[X_\tau]X_\tau$ from above, we can restrict attention to stopping rules that always accept a value of $\frac{\sqrt{n}}{e-2}$ and never accept a value of 0. Given the property that $\Pr[\tau = i \mid \tau > i - 1]$ is independent of i , any such stopping rule is characterized by the probability with which it accepts a value of 1. Denoting this probability, which may depend on n , by q_n , and the corresponding stopping time by τ_{q_n} ,

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

We now distinguish three cases depending on the limit behavior of $q_n\sqrt{n}$.

If $\limsup_{n \rightarrow \infty} q_n\sqrt{n} = \infty$, then, for every $\varepsilon > 0$ and $m \in \mathbb{N}$, there exists $n \geq m$ such that

$$\begin{aligned} \mathbb{E}[X_{\tau_{q_n}}] &= \left(1 - \left(1 - \frac{q_n\sqrt{n} + 1/\sqrt{n}}{n}\right)^n\right) \frac{q_n\sqrt{n} + 1/(e-2)}{q_n\sqrt{n} + 1/\sqrt{n}} \\ &\leq \frac{q_n\sqrt{n} + 1/(e-2)}{q_n\sqrt{n} + 1/\sqrt{n}} \\ &\leq 1 + \varepsilon. \end{aligned}$$

Indeed, the first inequality holds because $(1 - (1 - \frac{1}{n}(q_n\sqrt{n} + 1/\sqrt{n}))^n) \leq 1$, and the second inequality holds because $\limsup_{n \rightarrow \infty} q_n\sqrt{n} = \infty$.

If $\liminf_{n \rightarrow \infty} q_n \sqrt{n} = 0$, then, for every $\varepsilon > 0$ and every $m \in \mathbb{N}$, there exists $n \geq m$ such that

$$\begin{aligned} \mathbb{E}[X_{\tau_{q_n}}] &= \left(1 - \left(1 - \frac{q_n \sqrt{n} + 1/\sqrt{n}}{n}\right)^n\right) \frac{q_n \sqrt{n} + 1/(e-2)}{q_n \sqrt{n} + 1/\sqrt{n}} \\ &\leq (q_n \sqrt{n} + 1/\sqrt{n}) \frac{q_n \sqrt{n} + 1/(e-2)}{q_n \sqrt{n} + 1/\sqrt{n}} \\ &\leq \frac{1}{e-2} + \varepsilon. \end{aligned}$$

For the first inequality, we have used that, for $y = q_n \sqrt{n} + 1/\sqrt{n}$ and by Bernoulli's inequality, $1 - (1 - y/n)^n \leq y$, provided that $y/n \leq 1$, which is satisfied because $q_n \leq 1$ and $n \geq 3$. The second inequality holds because $\liminf_{n \rightarrow \infty} q_n \sqrt{n} = 0$.

Finally, if $\limsup_{n \rightarrow \infty} q_n \sqrt{n} < \infty$ and $\liminf_{n \rightarrow \infty} q_n \sqrt{n} > 0$, then there exist constants c_1, c_2 with $0 < c_1 \leq c_2$ and infinitely many values of n such that $q_n \sqrt{n} + 1/\sqrt{n} \in [c_1, c_2]$. Then, for every $\varepsilon > 0$ and $m \in \mathbb{N}$, and for $\varepsilon' = \varepsilon \cdot (c_2 + 1/(e-2))/c_1$, there exists n with $n \geq m$ and $q_n \sqrt{n} + 1/\sqrt{n} \in [c_1, c_2]$ such that

$$\begin{aligned} \mathbb{E}[X_{\tau_{q_n}}] &= \left(1 - \left(1 - \frac{q_n \sqrt{n} + 1/\sqrt{n}}{n}\right)^n\right) \frac{q_n \sqrt{n} + 1/(e-2)}{q_n \sqrt{n} + 1/\sqrt{n}} \\ &\leq (1 - e^{-(q_n \sqrt{n} + 1/\sqrt{n})} + \varepsilon') \frac{q_n \sqrt{n} + 1/(e-2)}{q_n \sqrt{n} + 1/\sqrt{n}} \\ &\leq (1 - e^{-(q_n \sqrt{n} + 1/\sqrt{n})}) \frac{q_n \sqrt{n} + 1/(e-2)}{q_n \sqrt{n} + 1/\sqrt{n}} + \varepsilon \\ &\leq \max_{x \geq 0} \left((1 - e^{-x}) \frac{x + 1/(e-2)}{x} \right) + \varepsilon \\ &\leq \frac{(e-1)^2}{(e-2)e} + \varepsilon. \end{aligned}$$

For the first inequality, we have used that, for every $\varepsilon > 0$, there exists $m \in \mathbb{N}$ such that, for all $n \geq m$ and all $x \in [c_1, c_2]$, it holds that $1 - (1 - x/n)^n \leq 1 - e^{-x} + \varepsilon$. The last inequality holds because the maximum of $(1 - e^{-x}) \cdot (x + 1/(e-2))/x$ is attained at $x = 1$, where it is equal to $(e-1)^2/((e-2)e)$.

For every stopping time τ such that $\Pr[\tau = i \mid \tau > i - 1]$ is independent of i , and for every $\varepsilon \geq 0$ and $m \in \mathbb{N}$, there thus exists $n \geq m$ such that

$$\mathbb{E}[X_\tau] \leq \max \left\{ 1 + \varepsilon, \frac{1}{e-2} + \varepsilon, \frac{(e-1)^2}{(e-2)e} + \varepsilon \right\} = \frac{(e-1)^2}{(e-2)e} + \varepsilon.$$

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ such that, for every $\varepsilon \geq 0$,

$$f(\varepsilon) = \frac{\frac{(e-1)^2}{(e-2)e} + \varepsilon}{\frac{e-1}{e-2} - \varepsilon}.$$

Then, for every $\varepsilon > 0$, there exists $n \in \mathbb{N}$ such that

$$\mathbb{E}[X_\tau] / \mathbb{E}[\max\{X_1, \dots, X_n\}] \leq f(\varepsilon).$$

Since f is continuous and $\lim_{\varepsilon \rightarrow 0} f(\varepsilon) = 1 - \frac{1}{e}$, there exists, for every $\delta > 0$, a value $\varepsilon > 0$ such that $f(\varepsilon) \leq 1 - \frac{1}{e} + \delta$, and thus $n \in \mathbb{N}$ such that

$$\mathbb{E}[X_\tau] / \mathbb{E}[\max\{X_1, \dots, X_n\}] \leq 1 - \frac{1}{e} + \delta,$$

as claimed. \square

Algorithms in the second class have access to $n - 1$ samples S_1, \dots, S_{n-1} from the underlying distribution and satisfy the following two natural conditions: (i) if the value of the first random variable X_1 is greater than all $n - 1$ samples, then they stop; and (ii) conditioned on reaching X_i , their probability of stopping at X_i is nondecreasing in i . It is again easily verified that Algorithm 1 belongs to this class.

Proposition 2. Let $\delta > 0$. Then there exists $n \in \mathbb{N}$ and a distribution F such that, for any $(n - 1, n)$ -stopping rule with stopping time τ that satisfies conditions (i) and (ii),

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

Proof. Let $\varepsilon > 0$. Let Y_i be distributed uniformly on $[0, \varepsilon]$, and let $X_i = 1 + Y_i$ with probability $1/n^2$ and $X_i = Y_i$ otherwise. Consider a stopping rule \mathbf{r} with stopping time τ that satisfies conditions (i) and (ii). Then, $\Pr[\tau = 1] \geq 1/n$ and $\Pr[\tau = i | \tau \geq i] \geq \Pr[\tau = 1] \geq 1/n$. Moreover, it is easy to see by induction that $\Pr[\tau \geq i] \leq (1 - 1/n)^{i-1}$. Indeed, $\Pr[\tau \geq 1] = 1$; for $i > 1$, and assuming that the claim holds for $i - 1$,

$$\Pr[\tau \geq i] = \Pr[\tau \geq i - 1] \cdot \Pr[\tau \neq i - 1 | \tau \geq i - 1] \leq (1 - 1/n)^{i-2} \cdot (1 - 1/n) = (1 - 1/n)^i.$$

It follows that

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

On the other hand,

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] \geq 1 - \left(1 - \frac{1}{n^2}\right)^n,$$

and, therefore,

$$\frac{\mathbb{E}[X_\tau]}{\mathbb{E}[\max\{X_1, \dots, X_n\}]} \leq \frac{\frac{1}{n} \left(1 - \left(1 - \frac{1}{n}\right)^n\right) + \varepsilon}{1 - \left(1 - \frac{1}{n^2}\right)^n}.$$

The right-hand side tends to $1 - 1/e$ as $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$, so, for every $\delta > 0$, there exists $n \in \mathbb{N}$ such that $\mathbb{E}[X_\tau] \leq (1 - 1/e + \delta) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}]$. \square

4.4. A Parametric Lower Bound

The lower bound of Theorem 3 can be generalized to a situation with γn samples when $\gamma \in [0, 1]$. The idea is to interpret some of the values X_1, \dots, X_n as samples, so that the number of remaining values equals the number of samples and Algorithm 1 can be applied.

Corollary 2. Let X_1, X_2, \dots, X_n be i.i.d. random variables from an unknown distribution F . Let $\gamma \in [0, 1]$ such that $\gamma n + n$ is an even number. Then there exists an $(\gamma n, n)$ -stopping-rule with stopping time τ such that

$$\mathbb{E}[X_\tau] \geq \frac{1 + \gamma}{2} \cdot \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}].$$

Proof. Let $n' = \frac{1 + \gamma}{2}n$, and note that $n' \in \mathbb{N}$. Define $S'_i = S_i$ for all $i \in [\gamma n]$, $S'_{\gamma n + i} = X_i$ for all $i \in [n - \gamma n]$, and $X'_i = X_{n' - \gamma n + i}$ for all $i \in [n']$. Note that $X'_{n'} = X_n$, so this assignment is well defined. We use Algorithm 1 with stopping time τ on $X'_1, \dots, X'_{n'}$ with samples $S'_1, \dots, S'_{n'-1}$. Then, by applying Theorem 3, we get

$$\mathbb{E}[X'_\tau] \geq \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X'_1, \dots, X'_{n'}\}] \geq \frac{1 + \gamma}{2} \cdot \left(1 - \frac{1}{e}\right) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

as claimed. \square

4.5. A Parametric Upper Bound

Whereas an improvement over the bound of $1 - 1/e \approx 0.632$ remains conceivable via more complicated stopping rules, such an improvement cannot go beyond $\ln(2) \approx 0.693$. This is a consequence of the following generalization of Theorem 2, which provides a parametric upper bound for stopping rules that have access to γn samples for some $\gamma \geq 0$. We prove this result by generalizing the proof of Theorem 2 and bounding the performance of the algorithm by bounding the probability that it accepts the maximum of the entire sequence of $(1 + \gamma)n$ values.

Theorem 4. Let $\delta > 0, \gamma \in \mathbb{R}_+$. Then there exists $n_0 \in \mathbb{N}$ such that, for any $n \geq n_0$ and any $(\gamma n, n)$ -stopping rule with stopping time τ , there exists a distribution F , not known to the stopping rule, with the following property. When X_1, \dots, X_n are i.i.d. random variables drawn from F ,

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

where

$$b(\gamma) = \begin{cases} \frac{1+\gamma}{e} & \text{if } \frac{1}{e} \geq \frac{\gamma}{1+\gamma}, \\ -\gamma \cdot \log \frac{\gamma}{1+\gamma} & \text{otherwise.} \end{cases}$$

Proof. We will restrict attention to $n \in \mathbb{N}$ such that $\gamma n \in \mathbb{N}$. It suffices to show that the guarantee of any $(\gamma n, n)$ -stopping rule is bounded from above by $b(\gamma) + o(1)$, where implicitly $n \rightarrow \infty$. To this end, consider an arbitrary $(\gamma n, n)$ -stopping rule with guarantee α . Let $\varepsilon = 1/n^2$. By Lemma 1, there then exists an infinite set $V \subseteq \mathbb{N}$ on which \mathbf{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

$$X_i = \begin{cases} v_1 & \text{with probability } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right), \\ \vdots \\ v_{n^3} & \text{with probability } \frac{1}{n^3} \cdot \left(1 - \frac{1}{n^2}\right), \\ u & \text{with probability } \frac{1}{n^2}. \end{cases}$$

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

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] \geq \Pr[X_{(n)} = u] \cdot u = \frac{1 - o(1)}{n} \cdot u. \tag{10}$$

For $i \in [(1 + \gamma)n]$, let

$$R_i = \begin{cases} S_i & \text{if } i \leq \gamma n, \\ X_{i-\gamma n} & \text{otherwise,} \end{cases}$$

and let $\tau' = \tau + \gamma n$ be the stopping time of \mathbf{r} viewed as a $(0, (1 + \gamma)n)$ -stopping rule on $R_1, \dots, R_{(1+\gamma)n}$. Then, analogously to the proof of Theorem 2,

$$\begin{aligned} \mathbb{E}[X_\tau] &= \Pr[R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} \neq u] \cdot \mathbb{E}[R_{\tau'} \mid R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} \neq u] \\ &\quad + \Pr[R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} = u] \cdot \mathbb{E}[R_{\tau'} \mid R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} = u] \\ &\quad + \Pr[R_{((1+\gamma)n)} \neq u] \cdot \mathbb{E}[R_{\tau'} \mid R_{((1+\gamma)n)} \neq u] \\ &\leq \frac{1+\gamma}{n} \cdot \left(\Pr[R_{\tau'} = R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} \neq u] \cdot u \right. \\ &\quad \left. + \Pr[R_{\tau'} \neq R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} \neq u] \cdot O(n^{-3}) \cdot u \right) \\ &\quad + O(n^{-2}) \cdot u + 1 \cdot O(n^{-3}) \cdot u \\ &\leq \frac{1+\gamma}{n} \cdot \Pr[R_{\tau'} = R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_{((1+\gamma)n-1)} \neq u] \cdot u + o\left(\frac{1}{n}\right) \cdot u \\ &\leq \frac{1+\gamma}{n} \cdot \Pr[R_{\tau'} = R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_1, \dots, R_{(1+\gamma)n} \text{ are distinct}] \cdot u + o\left(\frac{1}{n}\right) \cdot u. \end{aligned} \tag{11}$$

Given (10) and (11), to show $\alpha \leq b(\gamma) + o(1)$, it suffices to show that

$$\Pr[R_{\tau'} = R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_1, \dots, R_{(1+\gamma)n} \text{ are distinct}] \leq b(\gamma)/(1 + \gamma) + o(1). \tag{12}$$


Note that, in the event where $R_{((1+\gamma)n)} = u$ and $R_1, \dots, R_{(1+\gamma)n}$ are distinct, the relative ranks of $R_1, \dots, R_{(1+\gamma)n}$ are distributed uniformly at random. The optimal stopping rule for accepting the value with the largest relative rank is known to set, for some $x \in [0, 1]$, $q_i = 0$ for all $i < x \cdot (1 + \gamma) \cdot n$ and $q_i = 1$ for all $i \geq x \cdot (1 + \gamma) \cdot n$ (Gilbert and Mosteller [29]). Then, for any $(0, (1 + \gamma) \cdot n)$ -stopping rule $\hat{\mathbf{r}}$ with stopping time $\hat{\tau}$ that does not accept any of the values $X_1, \dots, X_{\gamma n}$,

$$\Pr[X_{\hat{\tau}} = R_{((1+\gamma)n)} \mid R_{((1+\gamma)n)} = u \wedge R_1, \dots, R_{(1+\gamma)n} \text{ are distinct}] = -x \cdot \log x + o(1),$$

which, subject to $x \geq \gamma/(1 + \gamma)$, is maximized for

$$x = \max\left\{\frac{1}{e}, \frac{\gamma}{1 + \gamma}\right\}.$$

Figure 2. Visualization of the parametric lower bound (solid) and the parametric upper bound (dashed).



Thus,

$$\Pr[X_{\hat{\tau}} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_1, \dots, R_{(1+\gamma)\cdot n} \text{ are distinct}] \leq b(\gamma)/(1 + \gamma) + o(1), \tag{13}$$

where

$$b(\gamma) = \begin{cases} \frac{1 + \gamma}{e} & \text{if } \frac{1}{e} \geq \frac{\gamma}{1 + \gamma}, \\ -\gamma \cdot \log \frac{\gamma}{1 + \gamma} & \text{otherwise.} \end{cases}$$

Analogously to the proof of Theorem 2, \mathbf{r} can be coupled with a $(0, (1 + \gamma) \cdot n)$ -stopping rule $\hat{\mathbf{r}}$ such that

$$\begin{aligned} & \Pr[X_{\hat{\tau}'} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_1, \dots, R_{(1+\gamma)\cdot n} \text{ are distinct}] \\ & \leq \Pr[X_{\hat{\tau}} = R_{((1+\gamma)\cdot n)} \mid R_{((1+\gamma)\cdot n)} = u \wedge R_1, \dots, R_{(1+\gamma)\cdot n} \text{ are distinct}] + n\varepsilon. \end{aligned} \tag{14}$$

Then, substituting (14) into (13) yields (12), completing the proof. \square

A visualization of the upper bound of Theorem 4 and the lower bound of Corollary 2 is shown in Figure 2. Note that the function b defined in Theorem 4 is continuous and that $b(1) = \ln(2) \approx 0.693$. Moreover, b meets the tight bound of approximately 0.745 due to Correa et al. [15], which implies an upper bound in the setting where the distribution is unknown, at $\gamma \approx 1.32$.

5. Superlinear Number of Samples

Our final result is that we can, in fact, match the optimal guarantee achievable by a stopping rule that knows the distribution, up to any $\varepsilon > 0$, if we have access to $O_\varepsilon(n^2)$ samples.

Theorem 5. *For every $\varepsilon > 0$, there exists an $n_\varepsilon \in \mathbb{N}$ such that the following holds for all $n \geq n_\varepsilon$. Let X_1, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then there exists an algorithm for choosing a stopping time $\tilde{\tau}$ that uses $O(n^2)$ samples from the same distribution with*

$$\mathbb{E}[X_{\tilde{\tau}}] \geq (\beta^{-1} - \varepsilon) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where $\beta^{-1} \approx 0.745$ is the guarantee shown by Correa et al. [15].

As our algorithm is related to that of Correa et al. [15], we first recall how that algorithm works. It computes a decreasing sequence x_1, x_2, \dots, x_n . As $n \rightarrow \infty$, it can be shown that x_i approaches $y(i/n)^{1/(n-1)}$ pointwise for all $i \in [n]$ (Kertz [36, theorem C]), where y is the unique solution to the following ordinary differential equation:

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

where $\beta \approx 1.3414 \approx 1/0.745$. This solution turns out to be decreasing and convex. Then, conditional on reaching random variable X_i , it chooses a quantile $q_i \in [1 - x_{i-1}, 1 - x_i]$ according to the probability density function

$$f_i(q) = \frac{(n-1)(1-q)^{n-2}}{\alpha_i} \quad \text{where} \quad \alpha_i = \int_{1-x_{i-1}}^{1-x_i} (n-1)(1-r)^{n-2} dr,$$

and sets $F^{-1}(1-q)$ as the threshold for accepting X_i .

Now let τ be the stopping time implied by the algorithm, and, for $q \in [0, 1]$, define $R(q) = \mathbb{E}[X | X \geq F^{-1}(1-q)]$ to be the expected value of random variable X conditioned on X exceeding threshold $F^{-1}(1-q)$. It can then be shown that

$$\begin{aligned} \mathbb{E}[X_\tau] &= \sum_{i=1}^n \Pr[\tau \geq i] \int_{1-x_{i-1}}^{1-x_i} f_i(q)R(q)q \, dq \\ &= \sum_{i=1}^n \rho \cdot \alpha_i \int_{1-x_{i-1}}^{1-x_i} f_i(q)R(q)q \, dq = \frac{\rho}{n} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}], \end{aligned} \tag{15}$$

where the i th term in the sum can be viewed as the contribution of X_i to the expectation $\mathbb{E}[X_\tau | X_\tau \geq i]$, and $\rho/n \geq \beta^{-1}$.

To simplify the presentation of our result, we first show that setting deterministic thresholds is also sufficient to achieve a guarantee of β^{-1} . In particular, for all $i \in [n]$, define \bar{q}_i such that

$$\bar{q}_i = \int_{1-x_{i-1}}^{1-x_i} f_i(q)q \, dq.$$

Let $\bar{\tau}$ be the stopping time of the algorithm that sets deterministic thresholds $F^{-1}(1-\bar{q}_i)$. The following lemma shows that we can also consider this algorithm.

Lemma 5. *Let X_1, \dots, X_n be i.i.d. random variables drawn from a known distribution F . Then, for all $i \in [n]$, we have*

$$R(\bar{q}_i) \cdot \bar{q}_i \geq \int_{1-x_{i-1}}^{1-x_i} f_i(q)R(q)q \, dq.$$

Proof. The left-hand side is $\mathbb{E}[X | X > T_d] \cdot \Pr[X > T_d]$, where T_d is the deterministic threshold $F^{-1}(1-\bar{q}_i)$ that corresponds to \bar{q}_i . The right-hand side is $\mathbb{E}[X | X > T_r] \cdot \Pr[X > T_r]$, where T_r is the randomized threshold that arises from first drawing $q \in [1 - x_{i-1}, 1 - x_i]$ with probability $f_i(q)$ and then setting threshold $F^{-1}(1-q)$.

We have chosen \bar{q}_i so that $\Pr[X > T_d] = \Pr[X > T_r]$; it thus suffices to show that $\mathbb{E}[X | X > T_d] \geq \mathbb{E}[X | X > T_r]$.

Let us prove the stronger statement that, for all t , $\Pr[X > t | X > T_d] \geq \Pr[X > t | X > T_r]$. Indeed, if $t > T_d$, then the claimed inequality becomes

$$\frac{\Pr[X > t]}{\Pr[X > T_d]} \geq \frac{\Pr[X > \max\{t, T_r\}]}{\Pr[X > T_r]},$$

which holds because $\Pr[X > T_d] = \Pr[X > T_r]$ and $\Pr[X > t] \geq \Pr[X > \max\{t, T_r\}]$. On the other hand, if $t \leq T_d$, then $\Pr[X > t | X > T_d] = 1$, and the claimed inequality applies as well. \square

By applying Lemma 5 to (15), we obtain

$$\mathbb{E}[X_\tau]X_{\bar{\tau}} = \sum_{i=1}^n \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i = \sum_{i=1}^n \rho \cdot \alpha_i \cdot R(\bar{q}_i) \cdot \bar{q}_i \geq \frac{\rho}{n} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}], \tag{16}$$

where we used that, conditioned on reaching step i , both τ and $\bar{\tau}$ accept X_i with the same probability, and so $\Pr[\bar{\tau} \geq i] = \Pr[\tau \geq i] = \rho \cdot \alpha_i$.

The algorithm that achieves the bound claimed in Theorem 5 starts by skipping some random variables until the acceptance probability \bar{q}_i of algorithm $\bar{\tau}$ becomes sufficiently large, say, δ/n , where $0 < \delta < 1$ is some constant. Such a step exists for sufficiently large n , because if all acceptance probabilities $\bar{q}_1, \dots, \bar{q}_n$ were at most δ/n , then would be at most $1 - 1/e \approx 0.632$ in the limit for $n \rightarrow \infty$, contradicting $\Pr[\bar{\tau} \leq n] \rightarrow \beta^{-1} \approx 0.745$ (Correa et al. [15]).

So, assume that ℓ is the first such step with $\bar{q}_\ell \geq \delta/n$. From then on, it uses the empirical distribution function of the samples to estimate the quantiles $\bar{q}_\ell, \dots, \bar{q}_n$ 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{q}_i , where \tilde{q}_i is its estimate of \bar{q}_i . More formally, when the original algorithm chooses threshold $\bar{T}_i = F^{-1}(1-\bar{q}_i)$, so that $1 - F(\bar{T}_i) = \bar{q}_i$, our algorithm will choose $\tilde{T}_i = \tilde{F}^{-1}(1-\tilde{q}_i)$, where \tilde{F} is the empirical distribution function, and $\tilde{q}_i = 1 - \tilde{F}(\tilde{T}_i)$. Denote the stopping time of this algorithm by $\tilde{\tau}_\delta$.

The reason why we skip the first few random variables is because the initial acceptance probability of the optimal algorithm that knows the distribution is of the order of $1/n^2$; therefore, with n^2 samples, we cannot get a reliable estimate of the corresponding quantile.

We will lower-bound the performance of our algorithm $\tilde{\tau}_\delta$ in terms of the performance of algorithm $\bar{\tau}$ through an intermediate algorithm, whose stopping time we denote by $\bar{\tau}_\delta$, that also skips the first few random variables but then uses the actual quantiles $\bar{q}_\ell, \dots, \bar{q}_n$.

Lemma 6. For every $\varepsilon > 0$, there exists an $n_\varepsilon \in \mathbb{N}$ such that the following holds for all $n \geq n_\varepsilon$. Let X_1, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then, for any δ such that $0 < \delta < 1/2$,

$$\mathbb{E}[X_{\tilde{\tau}_\delta}] \geq (1 - 2\delta) \cdot \mathbb{E}[X_{\bar{\tau}}].$$

Proof of Lemma 6. Note that if $\ell = 1$, then there is nothing to show as $\bar{\tau}_\delta$ and $\bar{\tau}$ are identical. Otherwise, $\ell \geq 2$, $\bar{q}_\ell \geq \delta/n$, and $\bar{q}_{\ell-1} \leq \delta/n$, which implies that $1 - x_{\ell-2} < \delta/n$.

The expected value achieved by the algorithm that skips the first few random variables until the acceptance probability becomes δ/n and then uses the actual quantiles $\bar{q}_\ell, \dots, \bar{q}_n$ is

$$\mathbb{E}[X_{\bar{\tau}_\delta}] = \sum_{i=\ell}^n \Pr[\bar{\tau}_\delta \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i \geq \sum_{i=\ell}^n \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i$$

while

$$\mathbb{E}[X_{\bar{\tau}}] = \sum_{i=1}^{\ell-1} \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i + \sum_{i=\ell}^n \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i.$$

Now observe that

$$\begin{aligned} \sum_{i=1}^{\ell-1} \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_i) \cdot \bar{q}_i &\leq \left(\sum_{i=1}^{\ell-2} \Pr[\bar{\tau} \geq i] + \Pr[\bar{\tau} \geq \ell - 1] \right) \cdot R(\bar{q}_\ell) \cdot \bar{q}_\ell \\ &\leq 2 \cdot \sum_{i=1}^{\ell-2} \Pr[\bar{\tau} \geq i] \cdot R(\bar{q}_\ell) \cdot \bar{q}_\ell = 2 \cdot \rho \cdot \left(\sum_{i=1}^{\ell-2} \alpha_i \right) \cdot R(\bar{q}_{\ell-1}) \cdot \bar{q}_{\ell-1} \\ &= 2 \cdot \rho \cdot \left(\int_0^{1-x_{\ell-2}} (n-1)(1-q)^{n-2} dq \right) \cdot R(\bar{q}_\ell) \cdot \bar{q}_\ell \\ &= 2 \cdot \rho \cdot (1 - x_{\ell-2}^{n-1}) \cdot R(\bar{q}_\ell) \cdot \bar{q}_\ell \leq 2 \cdot \rho \cdot (1 - (1 - \delta/n)^{n-1}) \cdot R(\bar{q}_\ell) \cdot \bar{q}_\ell \\ &\leq 2 \cdot \rho \cdot (1 - e^{-\delta}) \cdot R(\bar{q}_\ell) \bar{q}_\ell \leq 2 \cdot \rho \cdot \delta \cdot R(\bar{q}_\ell) \bar{q}_\ell, \end{aligned}$$

where for the first inequality we used that $R(q)q = \int_0^q F^{-1}(1-r)dr$ is monotone, for the second inequality we used that $\Pr[\bar{\tau} \geq \ell - 1] \leq \Pr[\bar{\tau} \geq \ell - 2]$, and for the third inequality we used that $x_{\ell-2} \geq 1 - \delta/n$.

On the other hand,

$$\begin{aligned} \sum_{i=\ell}^n \Pr[\bar{\tau} \geq i] R(\bar{q}_i) \bar{q}_i &\geq \sum_{i=\ell}^n \Pr[\bar{\tau} \geq i] R(\bar{q}_\ell) \bar{q}_\ell \\ &= \rho \left(\sum_{i=\ell}^n \alpha_i \right) R(\bar{q}_\ell) \bar{q}_\ell \\ &= \rho \cdot \int_{1-x_{\ell-1}}^1 (n-1)(1-q)^{n-2} dq \cdot R(\bar{q}_\ell) \bar{q}_\ell \\ &= \rho \cdot \left(1 - \int_0^{1-x_{\ell-1}} (n-1)(1-q)^{n-2} dq \right) \cdot R(\bar{q}_\ell) \bar{q}_\ell \\ &= \rho \left(1 - \sum_{i=1}^{\ell-1} \alpha_i \right) \cdot R(\bar{q}_\ell) \bar{q}_\ell \\ &= \left(\rho - \sum_{i=1}^{\ell-1} \Pr[\bar{\tau} \geq i] \right) \cdot R(\bar{q}_\ell) \bar{q}_\ell \\ &\geq \rho \cdot (1 - 2\delta) \cdot R(\bar{q}_\ell) \bar{q}_\ell, \end{aligned}$$

where we again used the monotonicity of $R(q)q$ for the first inequality and the upper bound for $\sum_{i=1}^{\ell-1} \Pr[\bar{\tau} \geq i]$ derived for the second inequality.

This shows that the ratio between the two terms is at most $2\delta/(1 - 2\delta)$, which in turn implies that

$$\mathbb{E}[X_{\bar{\tau}}] \leq \left(1 + \frac{2\delta}{1 - 2\delta}\right) \cdot \mathbb{E}[X_{\bar{\tau}_\delta}],$$

and, after rearranging, shows the claim. \square

Lemma 7. For every $\varepsilon' > 0$ and $\delta \in (0, 1/2)$, there exist $\varepsilon'' > \gamma > 0$, and $n_{\varepsilon'} \in \mathbb{N}$ such that the following holds for all $n \geq n_{\varepsilon'}$. Let X_1, \dots, X_n be i.i.d. random variables drawn from an unknown distribution F . Then with $k \geq n^2 \ln(2/\varepsilon'')/(2\gamma^2)$ samples, it holds that

$$\mathbb{E}[X_{\bar{\tau}_\delta}] \geq (1 - \varepsilon') \cdot \mathbb{E}[X_{\bar{\tau}_\delta}].$$

To prove this lemma, we make use of the following auxiliary lemma, which can be proven using the Dvoretzky-Kiefer-Wolfowitz inequality (Dvoretzky et al., lemma 2 [21]).

Lemma 8. Fix $\varepsilon'' > 0$ and $\gamma > 0$. Then, with $k \geq n^2 \ln(2/\varepsilon'')/(2\gamma^2)$ samples,

$$\Pr\left[\max_i |\bar{q}_i - \tilde{q}_i| > \frac{\gamma}{n}\right] < \varepsilon''.$$

Proof. Let F denote the true underlying distribution, and let \tilde{F} denote the empirical cumulative density function from k samples. With $k \geq n^2 \ln(2/\varepsilon'')/(2\gamma^2)$ samples, the Dvoretzky-Kiefer-Wolfowitz inequality from [21] states that

$$\Pr\left[\sup_x |\tilde{F}(x) - F(x)| > \frac{\gamma}{n}\right] \leq 2 \cdot e^{-2k(\gamma/n)^2} \leq \varepsilon''.$$

So with probability at least $1 - \varepsilon''$, we have that, for all pairs \bar{q}_i and $\tilde{q}_i = 1 - F(\tilde{T}_i)$,

$$\begin{aligned} \bar{q}_i = 1 - F(\tilde{T}_i) &\leq 1 - \tilde{F}(\tilde{T}_i) + \frac{\gamma}{n} = 1 - (1 - \bar{q}_i) + \frac{\gamma}{n} = \bar{q}_i + \frac{\gamma}{n}, \quad \text{and} \\ \tilde{q}_i = 1 - F(\tilde{T}_i) &\geq 1 - \tilde{F}(\tilde{T}_i) - \frac{\gamma}{n} = 1 - (1 - \bar{q}_i) - \frac{\gamma}{n} = \bar{q}_i - \frac{\gamma}{n}, \end{aligned}$$

as claimed. \square

Proof of Lemma 7. Given $\varepsilon' > 0$ and $\delta \in (0, 1/2)$, choose $\varepsilon'', \varepsilon''', \gamma > 0$ such that $\gamma \leq \delta, \gamma \leq (1 - \beta^{-1}) \cdot \varepsilon''' \approx 0.255 \cdot \varepsilon'''$, and $(1 - \varepsilon'') \cdot (1 - \varepsilon''') \cdot (1 - \gamma/\delta) \geq 1 - \varepsilon'$.

Toward relating $\mathbb{E}[X_{\bar{\tau}_\delta}]$ and $\mathbb{E}[X_{\bar{\tau}_\delta}]$, denote by ω the event that $\max_i |\bar{q}_i - \tilde{q}_i| \leq \gamma/n$. Note that we can lower-bound the expected value obtained by our algorithm by only considering the case that ω occurs and then summing over all steps:

$$\mathbb{E}[X_{\bar{\tau}_\delta}] \geq \Pr[\omega] \cdot \sum_{i=\ell+1}^n \Pr[\bar{\tau}_\delta \geq i | \omega] \cdot \Pr[X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i) | \omega \wedge \bar{\tau}_\delta \geq i] \cdot \mathbb{E}[X_i | \omega \wedge \bar{\tau}_\delta \geq i \wedge X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i)]. \quad (17)$$

Fix some $i \in [n]$ with $i \geq \ell + 1$. We first bound $\Pr[\bar{\tau}_\delta \geq i | \omega]$ with respect to $\Pr[\bar{\tau}_\delta \geq i | \omega] = \Pr[\bar{\tau}_\delta \geq i]$. First note that, for any $j \leq \ell$, it holds that $\Pr[\bar{\tau}_\delta = j] = \Pr[\bar{\tau}_\delta = j] = 0$. Furthermore, note that, for all $\ell + 1 \leq j < i$ and conditioned on ω , we have $|\bar{q}_j - \tilde{q}_j| \leq \gamma/n$. Thus, conditioned on ω , the probability that precisely one of $\bar{\tau}_\delta, \tilde{\tau}_\delta$ is j is bounded by γ/n . By the union bound, we therefore have

$$\Pr[\bar{\tau}_\delta \geq i | \omega] \geq \Pr[\bar{\tau}_\delta \geq i] - (i - \ell - 1) \cdot \gamma/n \geq \Pr[\bar{\tau}_\delta \geq i] - i \cdot \gamma/n.$$

Now note that $\Pr[\bar{\tau}_\delta \geq i] \geq \Pr[\bar{\tau}_\delta > n]$; that is, the probability that $\bar{\tau}_\delta$ stops at step i or later is lower-bounded by the probability that $\bar{\tau}_\delta$ does not stop at all. Moreover, $\Pr[\bar{\tau}_\delta > n] \geq \Pr[\tau > n]$; that is, the probability that $\bar{\tau}_\delta$ does not stop at all is at least the probability that τ does not stop at all. Since $\Pr[\tau > n] = 1 - \beta^{-1} \approx 0.255$, if we choose γ such that $\gamma \leq (1 - \beta^{-1}) \cdot \varepsilon'''$, then

$$\Pr[\bar{\tau}_\delta \geq i | \omega] \geq (1 - \varepsilon''') \cdot \Pr[\bar{\tau}_\delta \geq i]. \quad (18)$$

Next, we bound $\Pr[X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i) \mid \omega \wedge \tilde{\tau}_\delta \geq i] \cdot \mathbb{E}[X_i \mid \omega \wedge \tilde{\tau}_\delta \geq i \wedge X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i)]$, that is, the value extracted from X_i conditioned on ω and arriving in step i . We obtain

$$\begin{aligned} & \Pr[X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i) \mid \omega \wedge \tilde{\tau}_\delta \geq i] \cdot \mathbb{E}[X_i \mid \omega \wedge \tilde{\tau}_\delta \geq i \wedge X_i \geq \tilde{F}^{-1}(1 - \bar{q}_i)] \\ &= \mathbb{E}\left[\int_{\bar{q}_i}^1 F^{-1}(q) \, dq \mid \omega \wedge \tilde{\tau}_\delta \geq i\right] \\ &= \mathbb{E}\left[\int_{\bar{q}_i}^1 F^{-1}(q) \, dq - \int_{\bar{q}_i}^{\tilde{q}_i} F^{-1}(q) \, dq \mid \omega \wedge \tilde{\tau}_\delta \geq i\right] \\ &\geq \mathbb{E}\left[\left(1 - \frac{\gamma}{\delta}\right) \cdot \int_{\bar{q}_i}^1 F^{-1}(q) \, dq \mid \omega \wedge \tilde{\tau}_\delta \geq i\right] \\ &= \left(1 - \frac{\gamma}{\delta}\right) \cdot \int_{\bar{q}_i}^1 F^{-1}(q) \, dq \\ &= \left(1 - \frac{\gamma}{\delta}\right) \cdot \Pr[X_i \geq F^{-1}(1 - \bar{q}_i) \mid \bar{\tau}_\delta \geq i] \cdot \mathbb{E}[X_i \mid \bar{\tau}_\delta \geq i \wedge X_i \geq F^{-1}(1 - \bar{q}_i)], \end{aligned} \tag{19}$$

where in the third-to-last step we used that $F^{-1}(q)$ is monotonically increasing in q . Note that the inequality holds independently of how \bar{q}_i and \tilde{q}_i are ordered.

We now substitute the two bounds given in (18) and (19) into (17) and apply Lemma 8 to obtain

$$\mathbb{E}[X_{\tilde{\tau}_\delta}] \geq \Pr[\omega] \cdot (1 - \varepsilon''') \cdot \left(1 - \frac{\gamma}{\delta}\right) \cdot \mathbb{E}[X_{\bar{\tau}_\delta}] \geq (1 - \varepsilon'') \cdot (1 - \varepsilon''') \cdot \left(1 - \frac{\gamma}{\delta}\right) \cdot \mathbb{E}[X_{\bar{\tau}_\delta}] \geq (1 - \varepsilon') \cdot \mathbb{E}[X_{\bar{\tau}_\delta}],$$

as claimed. \square

We are now ready to prove the theorem.

Proof of Theorem 5. First choose $\varepsilon' > 0$ and $\delta \in (0, 1/2)$ such that $(1 - \varepsilon')(1 - 2\delta) \geq 1 - \varepsilon$. Then, by combining Lemma 6 with Lemma 7, we obtain that there exist $\varepsilon'', \gamma > 0$ such that, for sufficiently large n , with $k \geq n^2 \ln(2/\varepsilon'')/(2\gamma^2) = O(n^2)$ samples,

$$\mathbb{E}[X_{\tilde{\tau}_\delta}] \geq (1 - \varepsilon')(1 - \delta) \cdot \mathbb{E}[X_{\bar{\tau}}] \geq (1 - \varepsilon) \cdot \mathbb{E}[X_{\bar{\tau}}] \geq (\beta^{-1} - \varepsilon) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}],$$

where the last step follows from (16). \square

Acknowledgments

Discussions with Fábio Botler and valuable feedback from the anonymous referees are gratefully acknowledged. An extended abstract announcing the main results appeared in the *Proceedings of the 20th ACM Conference on Economics and Computation*. Part of the work was done while the second author was at Google Research in Zürich, Switzerland, and while the fourth author was at École Normale Supérieure in Paris, France, and Technische Universität München, Germany.

Appendix A. 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 their maximum as a threshold for the remaining ones. Since X_1, X_2, \dots, 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, \dots, X_n are drawn by f ,

$$\begin{aligned} \mathbb{E}[X_\tau] &= \int_0^\infty \dots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \mathbb{E}_\pi[v_{\pi(\tau)}] \, dv_1 \dots dv_n \\ &\geq \int_0^\infty \dots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \Pr_\pi[v_{\pi(\tau)} = \max\{v_1, \dots, v_n\}] \cdot \max\{v_1, \dots, v_n\} \, dv_1 \dots dv_n \\ &\geq \frac{1}{e} \cdot \int_0^\infty \dots \int_0^\infty \prod_{i=1}^n f(v_i) \cdot \max\{v_1, \dots, v_n\} \, dv_1 \dots dv_n \\ &= \frac{1}{e} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}], \end{aligned}$$

where the second inequality holds because the values v_1, \dots, v_n have been randomly ordered and τ is thus guaranteed to select $\max\{v_1, \dots, v_n\}$ with probability at least $1/e$ for any realization (Ferguson [28]). This proves the claim.

Appendix B. A 1/2-Approximation with $n - 1$ Samples

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 B.1. *Let X_1, X_2, \dots, X_n be i.i.d. draws from an unknown distribution F . Then there exists a $(n - 1, n)$ -stopping-rule with stopping time τ such that*

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

To prove Theorem B.1, we will analyze Algorithm B.1 a slight variation of Algorithm 1 that 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 B.1

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

$\tau \leftarrow n$

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

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

if $X_t \geq \max\{S_1, \dots, S_{n-1}\}$ **then**

$\tau \leftarrow t$

Break

end

end

return τ

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

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

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

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

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

$$\mathbb{E}[X_i | \tau = i] = \mathbb{E}[\max\text{ of } n - 1 + i \text{ i.i.d. RVs}] \geq \mathbb{E}[\max\{X_1, \dots, X_n\}]. \tag{B.3}$$

Substituting (B.2) and (B.3) into (B.1) completes the proof. \square

As we have argued in the proof of Theorem B.1, the probability that Algorithm B.1 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 $\mathbb{E}[\max\{X_1, \dots, X_n\}]$.

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

Proposition B.1. *Let X_1, \dots, X_n be drawn independently from the distribution with $F(x) = 1 - e^{-x}$. Then, for the stopping time τ determined by Algorithm B.1,*

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

Proof. It is a well-known fact that the maximum of n independent, exponentially distributed random variables X_1, \dots, X_n is equal to the n th harmonic number, that is, that

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

As we have argued in the proof of Theorem B.1, the expected value obtained by Algorithm B.1 can be written as

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

Tedious calculations allow us to express the expected value via the digamma function and the Euler-Mascheroni constant γ as

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

which can be used to show that

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

This proves the claim. \square

References

- [1] Abolhassani M, Ehsani S, Esfandiari H, Hajiaghayi MT, Kleinberg RD, Lucier B (2017) Beating $1 - 1/e$ for ordered prophets. *Proc. 49th Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 61–71.
- [2] Alaei S (2014) Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43(2):930–972.
- [3] Alijani R, Banerjee S, Gollapudi S, Munagala K, Wang K (2020) Predict and match: Prophet inequalities with uncertain supply. *Proc. ACM Measurement Anal. Comput. Systems* 4(1):1–23.
- [4] Anari N, Niazadeh R, Saberi A, Sharneli A (2019) Nearly optimal pricing algorithms for production constrained and laminar Bayesian selection. *Proc. 20th ACM Conf. Econom. Comput.* (ACM, New York), 91–92.
- [5] Azar Y, Chiplunkar A, Kaplan H (2018) Prophet secretary: Surpassing the $1 - 1/e$ barrier. *Proc. 19th ACM Conf. Econom. Comput.* (ACM, New York), 303–318.
- [6] Azar PD, Kleinberg R, Weinberg SM (2014) Prophet inequalities with limited information. *Proc. 25th Annual ACM-SIAM Sympos. Discrete Algorithms* (ACM, New York), 1358–1377.
- [7] Babaioff M, Blumrosen L, Dughmi S, Singer Y (2017) Posting prices with unknown distributions. *ACM Trans. Econom. Comput.* 5(2):13.
- [8] Berezovskiy BA, Gnedin AV (1984) *The Best Choice Problem* (in Russian) (Nauka, Moscow).
- [9] Bergemann D, Schlag K (2008) Pricing without priors. *J. Eur. Econom. Assoc.* 6(2–3):560–569.
- [10] Bergemann D, Schlag K (2011) Robust monopoly pricing. *J. Econom. Theory* 146(6):2527–2543.
- [11] Chawla S, Miller JB, Teng Y (2019) Posted pricing for online resource allocation: Intervals and paths. *Proc. 30th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1962–1981.
- [12] Chawla S, Hartline JD, Malec DL, Sivan B (2010) Multi-parameter mechanism design and sequential posted pricing. *Proc. 42nd Annual ACM Sympos. Theory Comput.* (ACM, New York), 311–320.
- [13] Correa J, Saona R, Ziliotto B (2019) Prophet secretary through blind strategies. *Proc. 30th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1946–1961.
- [14] Correa J, Cristi A, Epstein B, Soto JA (2020) The two-sided game of googol and sample-based prophet inequalities. *Proc. 31st Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 2066–2081.
- [15] Correa J, Foncea P, Hoeksma R, Oosterwijk T, Vredeveld T (2017) Posted price mechanisms for a random stream of customers. *Proc. 18th ACM Conf. Econom. Comput.* (ACM, New York), 169–186.
- [16] den Boer AV (2015) Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys Oper. Res. Management Sci.* 20(1):1–18.
- [17] Dütting P, Kesselheim T (2019) Posted pricing and prophet inequalities with inaccurate priors. *Proc. ACM Conf. Econom. Comput.* (ACM, New York), 111–129.
- [18] Dütting P, Kleinberg RD (2015) Polymatroid prophet inequalities. Bansal N, Finocchi I, eds. *Proc. 23rd Eur. Sympos. Algorithms* (Springer, Berlin), 437–449.
- [19] Dütting P, Kesselheim T, Lucier B (2020) An $O(\log \log m)$ prophet inequality for subadditive combinatorial auctions. *Proc. 61st Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 306–317.
- [20] Dütting P, Feldman M, Kesselheim T, Lucier B (2020) Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.* 49(3):540–582.
- [21] Dvoretzky A, Kiefer J, Wolfowitz J (1956) Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Ann. Math. Statist.* 27(3):642–669.
- [22] Ehsani S, Hajiaghayi MT, Kesselheim T, Singla S (2018) Prophet secretary for combinatorial auctions and matroids. *Proc. 29th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 700–714.
- [23] Esfandiari H, Hajiaghayi MT, Liaghat V, Monemizadeh M (2015) Prophet secretary. Bansal N, Finocchi I, eds. *Proc. 23rd Eur. Sympos. Algorithms* (Springer, Berlin), 496–508.
- [24] Ezra T, Feldman M, Gravin N, Tang ZG (2020) Online stochastic max-weight matching: Prophet inequality for vertex and edge arrival models. *Proc. 21st ACM Conf. Econom. Comput.* (ACM, New York), 769–787.

- [25] Feldman M, Gravin N, Lucier B (2015) Combinatorial auctions via posted prices. *Proc. 26th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 123–135.
- [26] Feldman M, Svensson O, Zenklusen R (2016) Online contention resolution schemes. *Proc. 27th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1014–1033.
- [27] Feng Y, Niazadeh R, Saberi A (2019) Linear programming based online policies for real-time assortment of reusable resources. Preprint, submitted July 17, <https://dx.doi.org/10.2139/ssrn.3421227>.
- [28] Ferguson TS (1989) Who solved the secretary problem? *Statist. Sci.* 4(3):282–289.
- [29] Gilbert JP, Mosteller F (1966) Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.* 61(313):35–73.
- [30] Goldenshluger A, Zeevi A (2021) Optimal stopping of a random sequence with unknown distribution. *Math. Oper. Res.*, ePub ahead of print October 25, <https://doi.org/10.1287/moor.2020.1109>.
- [31] Gravin N, Wang H (2019) Prophet inequality for bipartite matching: Merits of being simple and non adaptive. *Proc. 20th ACM Conf. Econ. Comput.* (ACM, New York), 93–109.
- [32] Hajiaghayi MT, Kleinberg RD, Sandholm TW (2007) Automated mechanism design and prophet inequalities. Cohn A, ed. *Proc. 22nd AAAI Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 58–65.
- [33] Hill TP, Kertz RP (1982) Comparisons of stop rule and supremum expectations of i.i.d. random variables. *Ann. Probab.* 10(2):336–345.
- [34] Hill TP, Kertz RP (1992) A survey of prophet inequalities in optimal stopping theory. *Contemporary Math.* 125:191–207.
- [35] Kaplan H, Naori D, Raz D (2020) Competitive analysis with a sample and the secretary problem. *Proc. 31st Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 2082–2095.
- [36] Kertz RP (1986) Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *J. Multivariate Anal.* 19(1):88–112.
- [37] Kleinberg RD, Weinberg SM (2012) Matroid prophet inequalities. *Proc. 44th Annual ACM Sympos. Theory Comput.* (ACM, New York), 123–136.
- [38] Krengel U, Sucheston L (1977) Semiamarts and finite values. *Bull. Amer. Math. Soc. (N.S.)*. 83(4):745–747.
- [39] Krengel U, Sucheston L (1978) On semiamarts, amarts, and processes with finite value. *Probability on Banach Spaces* (Dekker, New York), 197–266.
- [40] Ramsey FP (1930) On a problem of formal logic. *Proc. London Math. Soc. (3)* 30(1):264–286.
- [41] Rubinstein A (2016) Beyond matroids: Secretary problem and prophet inequality with general constraints. *Proc. 48th Annual ACM Sympos. Theory Comput.* (ACM, New York), 324–332.
- [42] Rubinstein A, Singla S (2017) Combinatorial prophet inequalities. *Proc. 28th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1671–1687.
- [43] Rubinstein A, Wang JZ, Weinberg SM (2020) Optimal single-choice prophet inequalities from samples. *Proc. 11th Sympos. Innovations Theoret. Comput. Sci., Seattle*, January 12–14, 1–60.
- [44] Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.* 12(4):1213–1216.