Unknown I.I.D. Prophets: Better Bounds, Streaming Algorithms, and a New Impossibility

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

Bruno Ziliotto[¶]

Abstract

A prophet inequality states, for some $\alpha \in [0, 1]$, that the expected value achievable by a gambler who sequentially observes random variables $X_1, \ldots, X_n$ and selects one of them is at least an α fraction of the maximum value in the sequence. We obtain three distinct improvements for a setting that was first studied by Correa et al. (EC, 2019) and is particularly relevant to modern applications in algorithmic pricing. In this setting, the random variables are i.i.d. from an unknown distribution and the gambler has access to an additional βn samples for some $\beta \geq 0$. We first give improved lower bounds on α for a wide range of values of β; specifically, $\alpha \geq (1 + \beta)/e$ when $\beta \leq 1/(e - 1)$, which is tight, and $\alpha \geq 0.648$ when $\beta = 1$, which improves on a bound of around 0.635 due to Correa et al. (SODA, 2020). Adding to their practical appeal, specifically in the context of algorithmic pricing, we then show that the new bounds can be obtained even in a streaming model of computation and thus in situations where the use of relevant data is complicated by the sheer amount of data available. We finally establish that the upper bound of $1/e$ for the case without samples is robust to additional information about the distribution, and applies also to sequences of i.i.d. random variables whose distribution is itself drawn, according to a known distribution, from a finite set of known candidate distributions. This implies a tight prophet inequality for exchangeable sequences of random variables, answering a question of Hill and Kertz (Contemporary Mathematics, 1992), but leaves open the possibility of better guarantees when the number of candidate distributions is small, a setting we believe is of strong interest to applications.

1 Introduction

Prophet inequalities have been studied extensively in optimal stopping theory and have recently seen a surge of interest in theoretical computer science. They are less pessimistic than classic worst-case competitive analysis as a framework for the study of online algorithms, and have major applications in algorithmic mechanism design and algorithmic pricing. These applications include relatively obvious ones in the design and analysis of posted-price mechanisms, but also more indirect ones like the choice of reserve prices in online advertising auctions (see, e.g., the survey of Hajiaghayi and Liaghat [2016]).

[*]Departamento de Ingeniería Industrial, Universidad de Chile, Avenue República 701, Santiago, Chile. Email: correa@uchile.cl. Supported in part by ANID Chile through grant CMM-AFB 170001.

[†]Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK. Email: p.d.duetting@lse.ac.uk. Part of the work was done while the author was at Google Research, Zürich, Switzerland.

[‡]School of Mathematical Sciences, Queen Mary University of London, Mile End Road, London E1 4NS, UK. Email: felix.fischer@qmul.ac.uk. Supported in part by EPSRC grant EP/T015187/1.

[§]Department Mathematik/Informatik, Abteilung Informatik, Universität zu Köln, Weyertal 121, 50931, Köln, Germany. Email: kschewior@gmail.com

[¶]Centre de Recherche en Mathématiques de la Décision, Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, 75016 Paris, France. Email: ziliotto@math.cnrs.fr.

The basic prophet inequality problem—sometimes referred to as the *single-choice* or *single-item* prophet inequality—is the following: a *gambler* observes a sequence of random variables $X_1, \ldots, X_n$ which satisfy certain assumptions, for example that they are drawn independently from possibly distinct known distributions. Upon seeing the realization $X_i = x_i$ the gambler has to decide immediately and irrevocably whether to stop and receive x_i as reward or to continue to the next random variable. Denoting the possibly randomized index at which the gambler stops by τ, the expected reward of the gambler is $\mathbb{E}[X_\tau]$. The goal is to find a stopping rule that is competitive with the expected reward of an all-knowing *prophet* who in particular knows the entire sequence and can simply choose the highest reward in the sequence. The goal thus is to determine the largest $\alpha \in [0, 1]$ such that the inequality $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max_i\{X_1, \ldots, X_n\}]$ holds for all random variables that satisfy the assumptions.

When distributions are known, optimal stopping rules can in principle be found by backward induction. Determining the optimal competitive ratio α may, nevertheless, be an intricate task. Whereas a tight result of $\alpha = 1/2$ for non-identical distributions has been known since the 1980s [Krengel and Sucheston, 1977, 1978; Samuel-Cahn, 1984], a bound of $\alpha \approx 0.745$ for the case of identical distributions was only recently shown to be tight [Hill and Kertz, 1982; Abolhassani et al., 2017; Correa et al., 2017].

An exciting new set of questions arises if we assume that the distributions of the random variables are unknown. In what is arguably the most basic setting of this kind, the gambler faces n i.i.d. draws from an unknown distribution. Since the random variables come from the same distribution, we may hope to be able to learn from earlier draws and apply what we have learned to later ones. Correa et al. [2019a] showed that, perhaps surprisingly, no learning is possible and the optimal solution mirrors the optimal solution to the classic secretary algorithm to achieve a bound of $1/e$; the impossibility persists even with $o(n)$ additional samples from the distribution, whereas a considerably better lower bound can be achieved with βn samples, for $\beta > 0$, and $O(n^2)$ samples are enough to get arbitrarily close to the optimal bound of $\alpha \approx 0.745$ achievable for a known distribution. The latter can actually be achieved already with $O(n)$ samples [Rubinstein et al., 2020]. While improved bounds have been obtained for the case with βn samples [Correa et al., 2020; Kaplan et al., 2020], significant gaps remain between lower and upper bounds across the whole range of values of β. In a setting with non-identical distributions, a single sample from each distribution is enough to match the optimal bound of $\alpha = 1/2$ achievable with full knowledge of the distributions [Rubinstein et al., 2020].

1.1 Our Contribution

We improve the results of Correa et al. [2019a] and the follow-up work in three distinct directions. Firstly, we give a new prophet inequality for i.i.d. random variables from an unknown distribution and access to additional samples. The inequality is tight when the number of samples is small, and otherwise improves on the state of the art. To obtain it we consider a natural class of algorithms and use tools from variational calculus to analyze the optimal algorithm from the class exactly. Secondly, we show that the new prophet inequality can be translated with arbitrarily small loss into a streaming model of computation. While streaming algorithms for optimal stopping have to our knowledge not previously been considered, they become very interesting when we deal with unknown distributions and large amounts of relevant data. The latter are commonly found in modern applications of prophet inequalities to algorithmic pricing. Thirdly, we give an upper bound of $1/e$ for a setting where the random variables are i.i.d. from a distribution F^j which is not itself known but drawn from a known distribution over a finite set of known scenarios $F^1, \ldots, F^m$. The bound is established for finite m and n, and scenarios of finite support, and shows that the impossibility in the absence of additional samples is robust to additional information regarding the distribution. It also implies a tight prophet inequality of $1/e$ for exchangeable sequences of random variables with a known joint distribution, answering a question of Hill and Kertz [1992].

Figure 1: Visualization of the lower bound established in this paper for varying β (solid, red), and comparison with the parametric upper bound (dashed black) and parametric lower bound (solid black) of Correa et al. [2019a] as well as the lower bounds of Kaplan et al. [2020] (blue) and Correa et al. [2020] (violet).

Unknown Distribution. In recent work, Correa et al. [2019a] introduced a prophet inequality problem involving n i.i.d. random variables from an unknown distribution and $k \geq 0$ additional samples. The main motivation for studying this problem arises from modern applications of prophet inequalities, specifically their use in the analysis of posted-price mechanisms and reserve pricing in advertising auctions. While it is common in these applications to model valuations as draws from an underlying distribution, it may not be reasonable to assume that the distribution is known to the auctioneer. The auctioneer may, however, choose to learn the distribution on the fly as opportunities arrive, or may possess some limited historical information in the form of additional samples.

Correa et al. [2019a] showed that in the absence of samples the correct value of α is $1/e$. They also gave parametric upper and lower bounds for the case of βn samples when $\beta > 0$, which for $k = n$ are equal to $1 - 1/e \approx 0.6321$ and $\ln(2) \approx 0.69$. Follow-up work has produced improved lower bounds for $k < n$ [Kaplan et al., 2020] and $k = n$ [Correa et al., 2020], the latter being equal to approximately 0.635.

The guarantee of $1 - 1/e$ for $k = n$ samples, achieved by Correa et al. [2019a] and also by Kaplan et al. [2020], can be obtained by setting a threshold for the acceptance of each arriving value that is equal to the maximum of a uniform random subset of size $n - 1$ of all samples and values seen thus far. The key ingredient in the analysis of this algorithm is a fresh-samples lemma, which states that each of the selected sets of values is distributed like $n-1$ fresh samples. The lemma implies that, conditioned on stopping, the expected value obtained is the same as that of the prophet. The competitive ratio is thus at least, and in fact equal to, the probability of stopping at all. Since conditioned on reaching a random variable we stop with probability $1/n$, this probability is $1 - (1 - 1/n)^n \geq 1 - 1/e$.

We study a natural generalization of the algorithm to what we refer to as *maximum-of-random-subset* (MRS) algorithms. MRS algorithms may use as the threshold in any given step a random subset *of any size* of the samples and values seen so far, and are thus characterized by a function mapping each step to the size of the set used at that step.

Our analysis of MRS algorithms is based on the observation that, with the help of a fresh-samples lemma, the expected value obtained by the algorithm can be written as a linear combination of expected maxima of certain numbers of fresh draws from the distribution. On the other hand, the value obtained by the prophet, to which we are comparing ourselves, is simply

3

the expected maximum of n fresh draws. We can thus determine the worst-case ratio between the two values by writing any maximum of γn draws as the integral $\int_0^\infty F^{\gamma n}(x)\, dx$ and then considering the integrals pointwise for each x. After substituting $a = F^n(x)$, we are left with the minimization of a single rational function in a single variable. As we will see, this technique leads to a tight analysis of MRS algorithms.

To determine the best MRS algorithm using an arbitrary number of samples, we take the supremum over all functions characterizing MRS algorithms of the solution to the aforementioned single-variable optimization problem. An upper bound on the solution of this supremum-infimum (control) problem can be obtained by swapping the supremum and the infimum, solving the inner supremum using the Euler–Lagrange equation, and tackling the remaining optimization problem by multivariable calculus. A lower bound can be obtained by substituting a solution of the form obtained from the Euler–Lagrange equation into the supremum-infimum problem and again solving the remaining optimization problem by multivariable calculus. It turns out that the resulting upper and lower bounds match, which indirectly implies a minimax result. The best guarantee that can be obtained using MRS algorithms is approximately 0.653 and can be realized with around $1.443 \cdot n$ samples.

For the case in which only a bounded number of βn samples is available, we conjecture a certain structure of the optimal MRS algorithm. Given this structure, we obtain a similar control problem as before. Using numerical methods, we obtain lower bounds for different values of β, which we conjecture to be tight up to errors in the numerical approximation. Our lower bounds improve on the state of the art for a broad range of values of β. In particular we establish the existence of an MRS algorithm using n samples that achieves a guarantee of roughly 0.649, which presents a considerable improvement over the previous best guarantee of 0.635 for the exact same setting due to Correa et al. [2020].

An interesting question going forward concerns the types of algorithms that could be used to close the remaining gap between the lower and upper bounds. Two natural candidates are algorithms that may skip an initial fraction of the values, effectively turning them into samples, and algorithms that may consider other (empirical) order statistics of random subsets of variables than just the maximum. We will see that the former class, which includes the classic algorithm for the secretary problem when $\beta = 0$, leads to a tight guarantee of $(1 + \beta)/e$ when $\beta \leq 1/(e - 1) \approx 0.58$.

A summary of our results for i.i.d. random variables from an unknown distribution and a comparison to the results from prior work can be found in Figure 1.

Streaming Prophet Inequalities. As we have mentioned earlier, the study of prophet inequalities for i.i.d. random variables from an unknown distribution is motivated by applications in posted pricing and in reserve pricing in advertising auctions. Such applications are often subject to an additional constraint, imposed by a vast amount of individual transactions that limits the way in which data about these transactions can be accessed and used. In particular, simply storing all past data in its entirety for later use is often impossible.

The canonical model for studying algorithms subject to this kind of constraint is the *streaming model* (see, e.g., the survey of Muthukrishnan [2005]). In the streaming model, an algorithm is allowed only a limited number of sequential passes over the input, often just a single one, and can only store a limited amount of information, logarithmic in the size of the input or constant. In the prophet problem, allowing only a single pass over the sequence of random variables is very natural, but the streaming model additionally limits the amount of memory available to an algorithm to $O(\log(k + n))$ or $O(1)$.

Streaming algorithms for the prophet problem have to our knowledge not previously been studied. This should not be surprising, since stopping rules typically rely on information about underlying distributions like their mean or quantiles, and distributions were until very recently assumed to be known from the outset. When distributions are not known, as in the model of Correa et al. [2019a], any information about them must be inferred from the sequence of random

variables and streaming algorithms suddenly become very desirable. This is true in particular when optimal stopping is used as a tool for modern applications in algorithmic pricing, which often involve very large amounts of relevant data.

MRS algorithms as described above are not streaming algorithms, as they may require access to arbitrary subsets of the random variables seen so far at any point. We will see, however, that any MRS algorithm can be implemented as a streaming algorithm with ε additive loss in the guarantee by storing $O_\varepsilon(1)$ samples. Note that computing a threshold for accepting the next value for *each* such value independently would essentially require remembering all of the seen values. Instead, we recycle the first computed threshold for some number of times steps and recompute it whenever it becomes too bad of an approximation for the threshold the MRS algorithm would actually set. We show that it suffices to (non-adaptively, i.e., independently of the observed values) choose $O_\varepsilon(1)$ time points for recomputing the threshold. The computation of the thresholds is implemented by an on-the-fly construction of the corresponding random subsets which only requires to store a single sample and $O(\log n)$ additional space.

Correa et al. [2019a] showed that $O_\varepsilon(n^2)$ samples are enough to get within ε of the tight bound of around 0.745 for the case of a known distribution. Quite remarkably, the same can already be achieved with just $O_\varepsilon(n)$ samples [Rubinstein et al., 2020]. The algorithms underlying both of these results rely heavily on *empirical quantiles*, and strong communication-complexity lower bounds for quantile estimation [Guha and McGregor, 2009] suggest that they cannot be implemented in the streaming model. This suggests that a gap might exist between the performance achievable respectively by streaming algorithms and algorithms that have random access to past random variables. We leave this interesting question for future work.

Unknown Distribution with Prior. The prophet problem for i.i.d. random variables from an unknown distribution is subject to a strong impossibility result that matches the obvious lower bound of $1/e$, and any improvement requires a significant number of additional samples. It is natural to ask whether an improvement is possible if information about the distribution is available not in the form of samples but more directly.

We explore this question by considering a setting where a distribution F^j is drawn from a set of distributions $F^1, \ldots, F^m$ according to a prior distribution θ, and the gambler is presented with random variables $X_1, \ldots, X_n$ drawn independently from F^j. The gambler knows the prior distribution θ including its support, but not the distribution F^j that has been drawn. This setting models situations where a decision maker is aware of the existence of a number of scenarios and has a good understanding of each scenario, but does not know which scenario is currently unfolding. For example, a merchant may be aware that there are good and bad days for selling a particular item, and may be aware how the valuation potential buyers have for the item is distributed on good days and on bad days, but may not be aware whether it is a good or a bad day.

Given the additional information about the distribution, we may hope to be able to improve on the obvious lower bound of $1/e$. We will see, however, that no such improvement is possible even when n, m, and the distributions $F^1, \ldots, F^m$ are all finite. This shows that the impossibility result for unknown distributions is rather robust. It also implies tightness of a known prophet inequality of $1/e$ due to Elton and Kertz [1991] for exchangeable sequences of random variables with a known joint distribution, answering a question of Hill and Kertz [1992]. Here, a sequence of random variables is called exchangeable if their joint distribution is invariant under permutations of the sequence, which is the case for the sequence of random variables for which we establish the upper bound. The upper bound does not preclude the existence of better lower bounds in settings where n or m is small. We believe such settings to be of significant practical interest and leave their study as an interesting problem for future work.

The proof of the upper bound uses an appropriate minimax argument to provide a reduction to the impossibility result for an unknown distribution. However, the argument itself requires a few new ideas and a modification of the key construction.

The basic idea behind our proof is to interpret the new prophet problem as a two-player zero-sum game, or equivalently a min-max problem, where the first player chooses a prior θ and the second player chooses a stopping rule. The payoff that the first player seeks to minimize and the second player seeks to maximize is the expected reward from the stopping rule minus $1/e$ times the expected maximum in the sequence. If we could reverse the order of minimization and maximization, and thus turn the problem into a max-min problem, we would be looking at a situation where player 2, the maximizer, moves first and chooses a stopping rule without knowing θ, and player 1, the minimizer, gets to choose a difficult θ with knowledge of the stopping rule. This max-min problem is, in fact, more difficult than the prophet problem for i.i.d. random variables from an unknown distribution considered by Correa et al. [2019a].

Indeed, the construction of Correa et al. relies on the infinite version of Ramsey's theorem [Ramsey, 1930], and in particular leads to distributions with infinite support for which the order of minimization and maximization cannot be reversed.

Minimax theorems do exist that can handle finite strategy spaces of player 2 and compact metric strategy spaces for player 1 [e.g., Mertens et al., 2015, Proposition 1.17], and this is the case we would get if the difficult instances for unknown i.i.d. distributions would have a support that is finite and bounded by a number that depends only on n. The argument sketched above could thus be rescued through a variant of the construction of Correa et al. [2019a] with this property. We provide such a construction by using a *finite* version of Ramsey's theorem as given for example by Conlon et al. [2010].

An interesting aspect of our argument for the game theory connoisseur is that the minimax theorem of course requires mixed strategies. This is clearly not a problem for player 2, but for player 1 this involves mixing over stopping rules. The validity of the above argument thus requires that any mixture of stopping rules can be implemented as a stopping rule. We will see that this readily follows from Kuhn's celebrated theorem on behavior strategies in extensive form games [Kuhn, 1953].

1.2 Further Related Work

Prophet Inequalities and Pricing. A comprehensive overview of early work on the classic single-choice prophet inequality can be found in the survey of Hill and Kertz [1992]. Starting from the work of Hajiaghayi et al. [2007], prophet inequalities have seen a surge of interest in the theoretical-computer-science literature.

Two important directions have considered extensions to richer domains where more than one element can be selected [e.g., Kleinberg and Weinberg, 2019; Alaei, 2014; Dütting and Kleinberg, 2015; Feldman et al., 2015; Dütting et al., 2020a; Rubinstein, 2016; Rubinstein and Singla, 2017; Chawla et al., 2019; Gravin and Wang, 2019; Ezra et al., 2020; Dütting et al., 2020b]), and random- or best-order models, both in the single-item setting and in combinatorial domains [e.g., Esfandiari et al., 2015; Abolhassani et al., 2017; Azar et al., 2018; Ehsani et al., 2018; Correa et al., 2019b; Agrawal et al., 2020].

Prophet inequalities with inexact priors were studied by Dütting and Kesselheim [2019], and prophet inequalities for unknown distributions with or without access to samples by Azar et al. [2014]; Babaioff et al. [2017]; Correa et al. [2019a]; Rubinstein et al. [2020]; Correa et al. [2020]; Kaplan et al. [2020].

Exchangeable random variables are generally positively correlated. Apart from the classical work of Rinott and Samuel-Cahn [1987] on negatively correlated random variables, the only work we are aware of to explicitly study correlation in the context of prophet inequalities is that of Immorlica et al. [2020].

Problems in pricing where distributional information is unavailable and must be learned have also been considered in operations research and management science. These problems typically differ significantly from those studied in computer science, and solutions often involve some form of regret minimization. Examples include recent results of Goldenshluger and Zeevi [2017] and those described in a survey of den Boer [2015].

Streaming Algorithms. To the best of our knowledge, we are the first to use the streaming model of computation to study prophet inequality problems. Related work has previously considered streaming algorithms for submodular maximization [e.g., Badanidiyuru et al., 2014; Feldman et al., 2020]. The algorithms in this literature are typically not required to make irrevocable online decisions, but they also don't assume a generative model for the data.

More closely related is work in the streaming literature that concerns the computation of various aggregate statistics of large amounts of stochastic data. For example, most prophet inequalities rely on the mean or median, or on more fine-grained information about the underlying distributions such as quantiles. While the empirical mean and variance of a sequence of values can be updated efficiently West [1979], strong lower bounds exist for the estimation of the median and other quantiles Guha and McGregor [2009].

Even more specifically, our streaming implementation of MRS algorithms shares certain characteristics with reservoir sampling and its variants [e.g., Vitter, 1985; Li, 1994], where the goal is to produce at any point in time a random subset of size k of the values seen so far. Compared to reservoir sampling, however, we require subsets of varying size at certain points in time as well as a better space complexity.

Algorithms from Data. In analyzing MRS algorithm with and without streaming, we optimize over a class of algorithms that has limited information about the problem at hand. Related problems have been considered under the umbrella of application-specific algorithm selection and data-driven algorithm design [e.g., Gupta and Roughgarden, 2017; Ailon et al., 2011; Balcan et al., 2018a,b], in particular in the context of designing revenue-optimal auctions from samples [e.g., Cole and Roughgarden, 2014; Morgenstern and Roughgarden, 2015].

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 $[i] = \{1, \ldots, i\}$.

Sequences of Random Variables. We consider sequences of non-negative and integrable random variables $X_1, \ldots, X_n$ drawn independently from a distribution F. The distribution F will either be unknown, or it will be drawn from a known prior distribution $\theta \in \Theta$, where Θ is the set of finite distributions over real distributions with finite support. A sequence of random variables is called exchangeable if their joint distribution is invariant under permutations of the sequence. The set of exchangeable sequences contains the sequences described above but is more general.

Stopping Rules. A stopping rule observes a sequence of random variables and decides to stop based on the random variables seen so far and possibly some additional information regarding the distribution from which the random variables are drawn. We consider two types of stopping rules, corresponding respectively to the cases where random variables are drawn from an unknown distribution and from a distribution that has itself been drawn according to a prior.

In the case of an unknown distribution, we consider stopping rules that in addition to the random variables observed so far may depend on k additional independent samples $S_1, \ldots, S_k$ from the same distribution as the random variables. Such a stopping rule can be expressed as a family $\mathbf{r}$ of functions $r_1, \ldots, r_n$, where $r_i : \mathbb{R}_+^{k+i} \to [0, 1]$ for all $i = 1, \ldots, n$. Here, for any $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$, $r_i(s_1 \ldots, s_k, x_1, \ldots, x_i)$ is the probability of stopping at X_i when we have observed samples $S_1 = s_1 \ldots, S = s_k$ and values $X_1 = x_1, \ldots, X_i = x_i$ and have not stopped at $X_1, \ldots, X_{i-1}$. The *stopping time* τ of such a stopping rule $\mathbf{r}$ is the random variable with support $\{1, \ldots, n\} \cup \{\infty\}$ such that for all $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$,

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

$$\cdot\, r_i(s_1 \ldots, s_k, x_1, \ldots, x_i).$$

In the case of a distribution drawn from a prior θ, we consider stopping rules that in addition of the random variables observed so far may depend on θ. Such a stopping rule can be expressed as a family $\mathbf{r}$ of functions $r_1, \ldots, r_n$, where $r_i : \mathbb{R}_+^i \times \Theta \to [0,1]$ for all $i = 1, \ldots, n$. Here, for any $\mathbf{x} \in \mathbb{R}_+^n$ and $\theta \in \Theta$, $r_i(x_1, \ldots, x_i, \theta)$ is the probability of stopping at X_i when we have observed the values $X_1 = x_1, \ldots, X_i = x_i$, have not stopped at $X_1, \ldots, X_{i-1}$, and when the prior distribution is θ. The *stopping time* τ of such a stopping rule $\mathbf{r}$ is thus the random variable with support $\{1, \ldots, n\} \cup \{\infty\}$ such that for all $\mathbf{x} \in \mathbb{R}_+^n$ and $\theta \in \Theta$,

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

Prophet Inequalities. For a given stopping rule we will be interested in the expected value $\mathbb{E}\left[X_\tau\right]$ of the variable at which it stops, where we use the convention that $X_\infty = 0$, and will measure its performance relative to the expected maximum $\mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right]$ of the random variables $X_1, \ldots, X_n$. We will say that a stopping rule achieves approximation guarantee α, for $\alpha \le 1$, if for any distribution or prior over distributions, $\mathbb{E}\left[X_\tau\right] \ge \alpha\, \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right]$.

Streaming Algorithms. In the case of an unknown distribution we will be interested specifically in stopping rules that can be implemented as streaming algorithms. We will assume that the stream consists of the samples $S_1, \ldots, S_k$ followed by the values $X_1, \ldots, X_n$. A streaming algorithm is then allowed a single pass over the sequence, and its space complexity is required to be logarithmic both in the length of the sequence, which will in fact be $O(n)$, and in $\max\{S_1, \ldots, S_k, X_1, \ldots, X_n\}$. Since the streaming algorithms we consider exclusively store values that occur in the stream, we will express their space complexity in terms of n only.

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

3 I.I.D. Random Variables from an Unknown Distribution

We now turn to the setting of unknown i.i.d. random variables, which was first studied by Correa et al. [2019a]. Here a lower bound of $\alpha = 1/e$ can be obtained via the optimal solution to the secretary problem, and this bound cannot be improved upon without access to additional samples from the distribution. A sharp phase transition occurs when there are linearly many samples in the number of variables. In particular, Correa et al. gave an algorithm that with n samples achieves an approximation ratio of $\alpha \ge 1 - 1/e \approx 0.6321$. The same bound was obtained by Kaplan et al. [2020], and improved to 0.635 by Correa et al. [2020]. We improve on this bound with n samples, and also obtain improved bounds for the case of βn samples for variable $\beta > 0$.

We do so by considering a natural class of algorithms that can be thought of as fixing a function $f : [n] \to \mathbb{N}$, and accepting variable X_i if its value exceeds the maximum of $f(i)$ fresh samples from the distribution. In fact, we will not draw fresh samples at each time step but rather choose uniformly at random a subset of size $f(i) \le k + i - 1$ from the set containing the k samples available to the algorithm at the outset and the values of the random variables $X_1, \ldots, X_{i-1}$ observed so far. We will refer to algorithms that follow this general strategy as *maximum-of-random-subset* (MRS) algorithms, and will see in Section 3.1 that they behave as if they would in fact draw fresh samples at each time step. The bound of $1 - 1/e$ of Correa et al. can in fact be achieved with an MRS algorithm where for all i, $f(i) = n - 1$.

We will first consider MRS algorithms that may use an arbitrary number k of samples, and give a tight bound on their approximation guarantee. Interestingly, the optimal such guarantee

8

of $\alpha \approx 0.6534$ is obtained with a bounded number of samples, namely $k \approx 1.4434 \cdot n$. We then give bounds for MRS algorithms that use at most βn samples for $0 \leq \beta \leq 1.4434$, and show specifically that there exists an MRS algorithm using n samples with an approximation guarantee of $\alpha \geq 0.6489 > 0.635 > 1 - 1/e$.

Theorem 3.1. *Consider a sequence of n random variables $X_1, \ldots, X_n$ drawn independently from an unknown distribution. Then, as $n \to \infty$, the best MRS algorithm with an unconstrained number k of samples achieves an approximation guarantee of $\alpha \approx 0.6534$ and requires $k \approx 1.4434 \cdot n$ samples.*

Theorem 3.2. *Consider a sequence of n random variables $X_1, \ldots, X_n$ drawn independently from an unknown distribution. Then there exists an MRS algorithm that uses $k = n$ samples and achieves an approximation ratio of $\alpha \geq 0.6489$ as $n \to \infty$.*

Towards proving these theorems, we first show a structural lemma in Section 3.1 that allows us to analyze MRS algorithms as if they would draw fresh samples at each step. In Section 3.2, we then consider as a warm-up an MRS algorithm using a simple piecewise linear function that improves upon the the approximation guarantee $1 - 1/e$. The remainder of the section is concerned with computing the best MRS algorithm. In Section 3.3, we consider the setting in which an unbounded number of samples is available, leading to the proof of Theorem 3.1. In Section 3.4, we then turn to the setting in which only a limited number of samples is available, proving Theorem 3.2 and additional bounds for other values of k.

3.1 Definition and Structural Lemma

Let $k \in \mathbb{N}$ and consider a function $f : [n] \to \mathbb{N}$ where $f(i) \leq k + i - 1$ for all $i \in [n]$. Then the MRS algorithm based on f proceeds as follows: given that it arrives at random variable X_i, it selects a uniformly random subset $\mathcal{R}_i = \{R_i^1, \ldots, R_i^{f(i)}\}$ of size $f(i)$ from the set $\{S_1, \ldots, S_k, X_1, \ldots, X_{i-1}\}$ of k samples and the first $i - 1$ random variables and sets $\max \mathcal{R}_i$ as threshold for X_i. We have the following lemma.

Lemma 3.3. *Consider some MRS algorithm based on $f : [n] \to \mathbb{N}$ and $i \in [n]$. Conditioned on the fact that the algorithm arrives at step i, the distribution of the set $\{S_1, \ldots, S_k, X_1, \ldots, X_{i-1}\}$ of values seen before step i is identical to the distribution of a set of $k + i - 1$ fresh samples from F.*

Proof. We show the claim by induction on i, and start by observing that it clearly holds for $i = 1$. Now suppose the claim holds for $i = 1, \ldots, i^\star - 1$. Then, conditioned on the fact that the algorithm arrives at step $i^\star - 1$, the set $\mathcal{T} = \{S_1, \ldots, S_k, X_1, \ldots, X_{i^\star - 2}\}$ has the same distribution as the one of a set of $k + i^\star - 2$ fresh samples, so the distribution of the set $\mathcal{T}' = \{S_1, \ldots, S_k, X_1, \ldots, X_{i^\star - 1}\}$ is the same as the one of a set of $k + i^\star - 1$ fresh samples. We will argue that the decision of the algorithm to stop at $X_{i^\star - 1}$ or to continue does not depend on the realization of $\mathcal{T}'$, which implies the claim.

Since F is continuous, we may assume that all the values $S_1, \ldots, S_k, X_1, \ldots, X_{i^\star - 1}$ are distinct, so that each of these values can be identified with a unique rank in $[k + i^\star - 1]$. By definition, the decision of an MRS algorithm to stop or continue only depends on the ranks of the values $R_{i^\star - 1}^1, \ldots, R_{i^\star - 1}^{f(i^\star - 1)}, X_{i^\star - 1}$. Since the distribution of $\mathcal{T}$, from which $R_{i^\star - 1}^1, \ldots, R_{i^\star - 1}^{f(i^\star - 1)}$ are drawn, is that of $k + i^\star - 1$ fresh samples, and since $X_{i^\star - 1}$ *is* a fresh sample, those ranks are $f(i^\star - 1) + 1$ uniform draws without replacement from $[k + i^\star - 1]$ and thus independent of the realization of $\mathcal{T}'$. $\qquad\square$

Note that this lemma immediately implies that a set of size $f(i)$ selected uniformly at random from $\{S_1, \ldots, S_k, X_1, \ldots, X_{i-1}\}$ is also distributed like $f(i)$ fresh samples from F. This insight is what will enable us to give a mathematical expression for the value obtained by an MRS algorithm.

3.2 Warm-Up: Three-Step Functions

As a warm-up, let us convince ourselves that an MRS algorithm with access to $k = n$ samples can improve on the bound of $1 - 1/e$. We will assume for simplicity that n is a multiple of 3, but note that the conclusion holds in general as $n \to \infty$. Let $f : [n] \to \mathbb{N}$ with $f(i) = n-1$ for $i = 1, \ldots, \frac{n}{3}$, $f(i) = 4n/3 - 1$ for $i = \frac{n}{3} + 1, \ldots, \frac{2n}{3}$, and $f(i) = 2n/3 - 1$ for $i = \frac{2n}{3} + 1, \ldots, n$. Observe that the MRS algorithm for f uses only n samples. Consider $X_1, \ldots, X_n$ drawn independently from an arbitrary distribution F, and let τ be the stopping time of the algorithm for these random variables. Then,

$$\mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right] = \int_0^\infty 1 - F(x)^n \, \mathrm{d}x$$

and, by Theorem 3.3,

$$\mathbb{E}[X_\tau] = \underbrace{\sum_{i=1}^{n/3} \left(1 - \frac{1}{n}\right)^{i-1} \frac{1}{n} \int_0^\infty 1 - F(x)^n \, \mathrm{d}x}_{=:T_1} +$$

$$(1 - T_1) \underbrace{\sum_{i=1}^{n/3} \left(1 - \frac{3}{4n}\right)^{i-1} \frac{3}{4n} \int_0^\infty 1 - F(x)^{4n/3} \, \mathrm{d}x}_{=:T_2} +$$

$$(1 - T_1)(1 - T_2) \underbrace{\sum_{i=1}^{n/3} \left(1 - \frac{2}{3n}\right)^{i-1} \frac{2}{3n} \int_0^\infty 1 - F(x)^{2n/3} \, \mathrm{d}x}_{=:T_3}.$$

For $n \to \infty$, $T_1 \to 1 - 1/e^{1/3}$, $T_2 \to 1 - 1/e^{1/4}$, and $T_3 \to 1 - 1/e^{1/2}$, and thus

$$\mathbb{E}[X_\tau] = \left(1 - \frac{1}{e^{1/3}}\right) \int_0^\infty 1 - F(x)^n \, \mathrm{d}x + \frac{1}{e^{1/3}} \left(1 - \frac{1}{e^{1/4}}\right) \int_0^\infty 1 - F(x)^{4n/3} \, \mathrm{d}x +$$

$$\frac{1}{e^{1/3}} \frac{1}{e^{1/4}} \left(1 - \frac{1}{e^{1/2}}\right) \int_0^\infty 1 - F(x)^{2n/3} \, \mathrm{d}x.$$

To show that $\mathbb{E}[X_\tau] \geq \alpha \mathbb{E}[\max\{X_1, \ldots, X_n\}]$, it suffices to show that for all $a \in [0, 1]$,

$$\left(1 - \frac{1}{e^{1/3}}\right)(1 - a) + \frac{1}{e^{1/3}} \left(1 - \frac{1}{e^{1/4}}\right)\left(1 - a^{4/3}\right) + \frac{1}{e^{1/3}} \frac{1}{e^{1/4}} \left(1 - \frac{1}{e^{1/2}}\right)\left(1 - a^{2/3}\right) \geq \alpha(1 - a).$$

By dividing both sides of the inequality by $1 - a$ and minimizing the resulting left-hand side over $a \in [0, 1)$, we see that this is the case for $\alpha \approx 0.6370 > 1 - 1/e$.

As we will see later in Section 3.3.1, requiring the approximation ratio to hold pointwise for each value of the variable of integration rather than just for the sum of integrals is without loss. We will thus be able to obtain tight bounds using this technique.

3.3 Proof of Theorem 3.1

To prove Theorem 3.1 we would like to find the best possible choice of $f : [n] \to \mathbb{N}$, without any restriction on the number of samples it is allowed to use. For a particular choice of f, we will again determine the approximation ratio α of the corresponding MRS algorithm by minimizing the ratio between $\mathbb{E}[X_\tau]$ and $\mathbb{E}[\max\{X_1, \ldots, X_n\}]$. We will again express both in terms of integrals, and turn the minimization over distributions F into a minimization over $a \in [0, 1)$ by requiring the approximation ratio to hold pointwise for each value of the variable of integration. Maximization over f and minimization over a will yield a max-min control problem, which we will solve optimally.

3.3.1 Formulation as a Control Problem

Fix $n \in \mathbb{N}$, and consider an MRS algorithm given by the function $f : [n] \to \mathbb{N}$. We can construct a continuous function $g : [0,1] \to \mathbb{R}_+$ from f by setting $g(i/n) = f(i)/n$ for all $i \in [n]$ and linearly interpolating between these values. Similarly, if we were only given a continuous function g in the first place, we could obtain f from g by setting $f(i) = \lceil g(i/n) \cdot n \rceil$ for all $i \in [n]$. In what follows we will compute the optimal such function g and thereby the optimal MRS algorithm for all values of n. To do so, consider a sequence $X_1, \ldots, X_n$ of random variables drawn i.i.d. from a distribution F, and denote the stopping time of the MRS algorithm on this sequence by τ. Then

$$\mathbb{E}\left[X_\tau\right] = \sum_{i=1}^{n} \Pr\left[A \text{ arrives at } X_i\right] \cdot \Pr\left[A \text{ accepts } X_i \mid A \text{ arrives at } X_i\right] \cdot \mathbb{E}\left[X_i \mid A \text{ accepts } X_i\right]$$

$$= \sum_{i=1}^{n} \prod_{j=1}^{i-1} \left(1 - \frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right) \cdot \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \cdot \int_0^\infty \left(1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)\right) \mathrm{d}x$$

$$= \sum_{i=1}^{n} \exp\left(\sum_{j=1}^{i-1} \ln\left(1 - \frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right)\right) \cdot \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \cdot \int_0^\infty \left(1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)\right) \mathrm{d}x$$

$$= \sum_{i=1}^{n} \exp\left(-\sum_{j=1}^{i-1} \left(\frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1} - O\left(\frac{1}{n^2}\right)\right)\right) \cdot \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \cdot \int_0^\infty \left(1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)\right) \mathrm{d}x$$

$$= e^{-O(\frac{1}{n})} \sum_{i=1}^{n} \exp\left(-\sum_{j=1}^{i-1} \left(\frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right)\right) \cdot \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \cdot \int_0^\infty \left(1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)\right) \mathrm{d}x,$$

where for the fourth equality we have used that the Laurent series of $\ln(1 - \frac{1}{x}) - (-\frac{1}{x})$ at $x = \infty$ is $\sum_{i=2}^{\infty} -\frac{x^{-i}}{i} = -O(\frac{1}{x^2})$ (and assumed w.l.o.g. that g is bounded away from 0). Thus, for $n \to \infty$,

$$\mathbb{E}\left[X_\tau\right] = \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, \mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \int_0^\infty \left(1 - F^{g(y) \cdot n}(x)\right) \mathrm{d}x \, \mathrm{d}y$$

$$= \int_0^\infty \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, \mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \left(1 - F^{g(y) \cdot n}(x)\right) \mathrm{d}y \, \mathrm{d}x,$$

where we exchange the order of integration in the second step using Fubini's theorem, which may be applied, because the integrand is clearly positive.

Our goal is to find the maximum value $\alpha \in \mathbb{R}_+$ for which $\mathbb{E}\left[X_\tau\right] \geq \alpha \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right]$ or, equivalently,

$$\int_0^\infty \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, \mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \left(1 - F^{g(y) \cdot n}(x)\right) \mathrm{d}y \, \mathrm{d}x \geq \int_0^\infty \alpha \cdot \left(1 - F^n(x)\right) \mathrm{d}x. \quad (1)$$

A sufficient condition for the latter is that for all $a \in [0,1]$,

$$\int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, \mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \left(1 - a^{g(y)}\right) \mathrm{d}y \geq \alpha \cdot (1 - a), \quad (2)$$

and this condition is in fact also necessary. Indeed, if (2) is violated for some α and a, then (1) is violated for α and the cumulative distribution function F of a random variable that has value 0 with probability a and value 1 with probability $(1 - a)$. This choice of F makes the integrand on the right-hand side of (1) greater than the integrand on the left-hand side for all x for which the integrands are non-zero, i.e., for all $x < 1$, thus violating (1).

To determine the approximation ratio of the MRS algorithm A we can thus focus on finding the maximum value α such that (2) is satisfied for all a. Since (2) is trivially satisfied for $a = 1$, we are interested in the optimum value of the control problem

$$
\begin{aligned}
\mathcal{P} &= \sup_{g:[0,1]\to\mathbb{R}_+} \inf_{a\in[0,1)} \left\{ \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)}\, \mathrm{d}z \right) \cdot \frac{1 - a^{g(y)}}{g(y)\cdot(1-a)}\, \mathrm{d}y \right\} \\
&= \sup_{\substack{h:[0,1]\to\mathbb{R}_+,\\ h(0)=0}} \inf_{a\in[0,1)} \left\{ \int_0^1 e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1-a}\, \mathrm{d}y \right\},
\end{aligned} \tag{3}
$$

where the second equality can be seen to hold by choosing $h : [0,1] \to \mathbb{R}_+$ such that $h(y) = \int_0^y \frac{1}{g(z)}\, \mathrm{d}z$ for all $y \in [0,1]$, which implies that $g(y) = \frac{1}{h'(y)}$.

3.3.2 Solving the Control Problem

We solve the control problem $\mathcal{P}$ by giving matching upper and lower bounds. For the upper bound we swap supremum and infimum and apply the Euler–Lagrange equation to the supremum, which is now the inner problem, to write any optimal function h in terms of a and a single parameter μ. We then guess the value of a at which the infimum is attained and solve the remaining supremum over μ. For the lower bound we replace h by its parametric form, guess the values of the parameters at which the supremum is attained, and solve the remaining infimum over a. In both cases we obtain the same value of approximately 0.6534. Inspection of the optimal function h reveals that it is non-increasing, which implies that $g(0) \cdot n = \frac{1}{h'(0)} n \approx 1.4434 \cdot n$ samples are sufficient to implement the optimal MRS algorithm.

It is worth pointing out that the change of the order of supremum and infimum and the substitution of a particular form of h are potentially lossy but turn out to be without loss. This means that a minimax theorem holds for $\mathcal{P}$, and that a universal worst-case distribution F exists that applies to all MRS algorithms.

Upper Bound. By the max-min inequality,

$$
\mathcal{P} \le \inf_{a\in[0,1)} \sup_{\substack{h:[0,1]\to\mathbb{R}_+,\\ h(0)=0}} \left\{ \int_0^1 e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1-a}\, \mathrm{d}y \right\}. \tag{4}
$$

Now the inner problem can be written as

$$
\sup_{\substack{h:[0,1]\to\mathbb{R}_+,\\ h(0)=0}} \int_0^1 L(y, h(y), h'(y))\, \mathrm{d}y,
$$

where

$$
L(y, h(y), h'(y)) = e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1-a}.
$$

A necessary condition for optimality of h is the Euler–Lagrange equation

$$
\frac{\partial}{\partial h} L(y, h(y), h'(y)) - \frac{\mathrm{d}}{\mathrm{d}y} \frac{\partial}{\partial h'} L(y, h(y), h'(y)) = 0, \tag{5}
$$

where

$$
\frac{\partial}{\partial h} L(y, h(y), h'(y)) = -e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1-a} \tag{6}
$$

12

and

$$\frac{\mathrm{d}}{\mathrm{d}y}\frac{\partial}{\partial h'}L(y,h(y),h'(y)) = \frac{\mathrm{d}}{\mathrm{d}y}\left(e^{-h(y)}\cdot\frac{1-a^{\frac{1}{h'(y)}}}{1-a} + e^{-h(y)}\cdot h'(y)\cdot\frac{\ln a\cdot a^{\frac{1}{h'(y)}}}{(1-a)\cdot(h'(y))^2}\right)$$

$$= -e^{-h(y)}\cdot h'(y)\cdot\frac{1-a^{\frac{1}{h'(y)}}}{1-a} + e^{-h(y)}\cdot\frac{\ln a\cdot a^{\frac{1}{h'(y)}}\cdot h''(y)}{(1-a)\cdot(h'(y))^2}$$

$$+ e^{-h(y)}\cdot(h''(y)-(h'(y))^2)\cdot\frac{\ln a\cdot a^{\frac{1}{h'(y)}}}{(1-a)\cdot(h'(y))^2}$$

$$+ e^{-h(y)}\cdot h'(y)\cdot\left(-\frac{(\ln(a))^2 a^{\frac{1}{h'(y)}}h''(y)}{(1-a)(h'(y))^4} - \frac{2\ln(a)a^{\frac{1}{h'(y)}}h''(y)}{(1-a)(h'(y))^3}\right). \quad (7)$$

Substitution of (6) and (7) into (5) and simplification yields that

$$-e^{-h(y)}\frac{\ln(a)a^{\frac{1}{h'(y)}}}{1-a} - e^{-h(y)}\frac{(\ln(a))^2 a^{\frac{1}{h'(y)}}h''(y)}{(1-a)(h'(y))^3} = 0.$$

Since $e^x > 0$ for all x and $1-a > 0$ for $a \in [0,1)$, an equivalent requirement is that

$$-\frac{h''(y)}{(h'(y))^3} = \frac{1}{\ln(a)}$$

with the boundary condition $h(0) = 0$.

Solving this second-order nonlinear ordinary differential equation yields two classes of parametric solutions

$$h_1(y) = \sqrt{\kappa-\mu y} - \sqrt{\kappa}, \quad h_2'(y) = -\frac{\mu}{2\sqrt{\kappa-\mu y}}, \quad \mu = -2\ln(a)\geq 0, \quad \kappa\geq\mu, \quad \text{and}$$

$$h_2(y) = \sqrt{\kappa} - \sqrt{\kappa-\mu y}, \quad h_2'(y) = \frac{\mu}{2\sqrt{\kappa-\mu y}}, \quad \mu = -2\ln(a)\geq 0, \quad \kappa\geq\mu,$$

where only the latter guarantees that $g(y) = 1/h'(y) \geq 0$.

Let $\bar{\mu} \approx 1.9202$ be the unique value such that

$$1 - \frac{e^{\sqrt{\bar{\mu}}}}{\sqrt{\bar{\mu}}} + \frac{e^{\frac{\bar{\mu}}{2}}}{\sqrt{\bar{\mu}}} = 0,$$

and $\bar{a} = e^{-\frac{\bar{\mu}}{2}} \approx 0.3829$.

By setting $h = h_2$ and $a = \bar{a}$ in (4), and showing that the remaining supremum over κ is attained for $\kappa = \bar{\mu}$, we conclude that

$$\mathcal{P} \leq \frac{e^{-\sqrt{\bar{\mu}}}(1-e^{\sqrt{\bar{\mu}}}+\sqrt{\bar{\mu}})}{e^{-\frac{\bar{\mu}}{2}}-1} \approx 0.6534.$$

Lower Bound. By restricting the supremum in (3) to functions of the form $h(y) = \sqrt{\mu} - \sqrt{\mu\cdot(1-y)}$ for some $\mu \in \mathbb{R}_+$, which satisfy the boundary condition that $h(0) = 0$, we see that

$$\mathcal{P} \geq \sup_{\mu\in\mathbb{R}_+}\inf_{a\in[0,1)}\left\{\frac{e^{-\sqrt{\mu}}}{1-a}\cdot\int_0^1\frac{e^{\sqrt{\mu\cdot(1-y)}}\cdot\mu\cdot\left(1-a^{\frac{2}{\mu}\sqrt{\mu\cdot(1-y)}}\right)}{2\cdot\sqrt{\mu\cdot(1-y)}}\,\mathrm{d}y\right\}$$

$$= \sup_{\mu\in\mathbb{R}_+}\inf_{a\in[0,1)}\left\{\frac{e^{-\sqrt{\mu}}}{1-a}\cdot\left[\frac{e^{\sqrt{\mu\cdot(1-y)}\cdot(1+\frac{2\ln a}{\mu})}}{1+\frac{2\ln a}{\mu}} - e^{\sqrt{\mu\cdot(1-y)}}\right]_0^1\right\}$$

13

$$= \sup_{\mu \in \mathbb{R}_+} \inf_{a \in [0,1)} \left\{ \frac{e^{-\sqrt{\mu}}}{1-a} \cdot \left(\frac{1}{1 + \frac{2\ln a}{\mu}} - 1 - \frac{e^{\sqrt{\mu} \cdot (1 + \frac{2\ln a}{\mu})}}{1 + \frac{2\ln a}{\mu}} + e^{\sqrt{\mu}} \right) \right\}$$

$$= \sup_{\mu \in \mathbb{R}_+} \inf_{\substack{b \in \mathbb{R}_+; \\ b \notin \{0,1\}}} \left\{ \frac{e^{-\sqrt{\mu}}}{1 - e^{-\frac{\mu b}{2}}} \cdot \left(\frac{1}{1-b} - 1 - \frac{e^{\sqrt{\mu} \cdot (1-b)}}{1-b} + e^{\sqrt{\mu}} \right) \right\},$$

where the last equality can be seen to hold by setting $b = -\frac{2\ln a}{\mu}$ and $a = e^{-\frac{\mu b}{2}}$.

By setting $\mu = \bar{\mu}$ in the last expression and showing that the remaining infimum over b is attained for $b \to 1$, we conclude that

$$\mathcal{P} \geq \frac{e^{-\sqrt{\bar{\mu}}}(1 - e^{\sqrt{\bar{\mu}}} + \sqrt{\bar{\mu}})}{e^{-\frac{\bar{\mu}}{2}} - 1},$$

which equals the upper bound.

The resulting optimal choice of g, given by $g(y) = 1/h'(y) = 2\sqrt{\mu - \bar{\mu}y}/\bar{\mu}$, is non-increasing in y and thus has a maximum value of $g(0) = 2/\sqrt{\bar{\mu}} \approx 1.4434$. This means that the optimal MRS algorithm can be implemented with slightly fewer than $3n/2$ samples.

3.4 Proof of Theorem 3.2

We finally consider MRS algorithms that have access to βn samples for some $\beta < 2/\sqrt{\bar{\mu}}$. This imposes the constraint that $g(y) \leq \beta + y$ for all $y \in [0,1]$, and since the optimal MRS algorithm for the unconstrained case uses more than βn samples the constraint must bind for some non-empty subset of $[0,1]$. To obtain a lower bound on the performance of the best MRS algorithm we may in fact assume that the constraint binds on $[0,t]$ for some $t \in [0,1]$, such that $g(y) = \beta + y$ and $h(y) = \int_0^y 1/g(z) \, dz = \ln(\beta + y) - \ln(\beta)$ for all $y \in [0,t]$. Proceeding as in Section 3.3.2, we can write the performance of the best MRS algorithm from the restricted class as a control problem

$$\mathcal{Q} = \sup_{\substack{t \in [0,1], \\ h:[t,1] \to \mathbb{R}_+, \\ h(t) = \ln(\beta + t)}} \inf_{a \in [0,1)} \left\{ \frac{\int_0^t \frac{\beta}{(\beta+y)^2} \cdot \left(1 - a^{\beta+y}\right) \, dy + \int_t^1 e^{-h(y)} \cdot h'(y) \cdot \left(1 - a^{\frac{1}{h'(y)}}\right) \, dy}{1-a} \right\}$$

$$= \sup_{\substack{t \in [0,1], \\ h:[t,1] \to \mathbb{R}_+, \\ h(t) = \ln(\beta + t)}} \inf_{a \in [0,1)} \left\{ \int_0^t \frac{\beta(1 - a^{\beta+y})}{(\beta+y)^2(1-a)} \, dy + \int_t^1 e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1-a} \, dy \right\}.$$

Note that the objective is now a sum of two integrals. The first integral is constant with respect to h. The second integral has the same integrand as the integral in problem $\mathcal{P}$ from Section 3.3.2, but it begins at t rather than 0 and involves a function h that is subject to a different boundary condition, $h(t) = \ln(\beta + t) - \ln(\beta)$ instead of $h(0) = 0$. As our application of the Euler–Lagrange equation in Section 3.3.2 relied neither on the limits of integration nor on the boundary condition we obtain the same differential equation as before, $-h''(y)/(h'(y))^3 = 1/\ln(a)$, but subject to the new boundary condition that $h(t) = \ln(\beta + t) - \ln(\beta)$.

Since $g(y) = 1/h'(y)$ for $y \in (0,1)$ and thus

$$g(y) \cdot g'(y) = \frac{((g(y))^2)'}{2} = \frac{1}{2} \left(\frac{1}{(h'(y))^2} \right)' = -\frac{h''(y)}{(h'(y))^3}$$

for $y \geq t$, we can alternatively solve the first-order non-linear differential equation $g(t) \cdot g'(t) = 1/\ln(a)$. From the requirement that $g(y) \geq 0$ for all y we conclude that

$$g(y) = \sqrt{2} \cdot \sqrt{\frac{1}{\ln(a)} \cdot y + \kappa}$$

14

β	$\alpha \geq$	$t \approx$	$a \approx$	β	$\alpha \geq$	$t \approx$	$a \approx$
1.4	0.653368	0.025503	0.383230	$\vdots$	$\vdots$	$\vdots$	$\vdots$
1.3	0.653280	0.087540	0.387562				
1.2	0.652853	0.155180	0.398509	0.7	0.616281	0.720814	0.758359
1.1	0.651654	0.230674	0.419390	0.6	0.588379	0.949784	0.959047
1.0	0.648957	0.317590	0.455588	0.5	0.549306	1.000000	1.000000
0.9	0.643563	0.421611	0.515673	0.4	0.501105	1.000000	1.000000
0.8	0.633580	0.551596	0.612066	0.3	0.439901	1.000000	1.000000
$\vdots$	$\vdots$	$\vdots$	$\vdots$	0.2	0.358351	1.000000	1.000000
				0.1	0.239789	1.000000	1.000000

Table 1: Lower bounds on the performance of the optimal MRS algorithm with access to βn samples for varying values of β. The bounds arise as the minimum over a of a function in t, and the values of t and a corresponding to each bound are given alongside it.

for some $\kappa \geq -1/\ln(a)$, and by choosing κ to satisfy the boundary condition that $g(t) = \beta + t$ we obtain

$$g(y) = \sqrt{2} \cdot \sqrt{\frac{1}{\ln(a)} \cdot y + \frac{1}{2}\left(-\frac{2}{\ln(a)} \cdot t + t^2 + 2\beta t + \beta^2\right)}.$$

In analogy to Section 3.3.2 we may derive a lower bound on the value of $\mathcal{Q}$ by considering the parametric class of functions

$$g(y) = \sqrt{2} \cdot \sqrt{cy + \frac{1}{2}\left(-2ct + t^2 + 2\beta t + \beta^2\right)},$$

where $c \leq 0$, and we may in fact choose $c = (\beta + t)^2/(2(t-1))$ such that $g(1) = 0$ as before. Then

$$g(y) = (\beta + t)\sqrt{\frac{y-1}{t-1}}$$

and

$$h(y) = \ln(\beta + t) - \ln(\beta) + \frac{2(y-1)}{\sqrt{\frac{(y-1)(\beta+t)^2}{t-1}}} - \frac{2(t-1)}{\sqrt{(\beta+t)^2}}$$

We can now substitute h into $\mathcal{Q}$ and solve the integrals to obtain a simpler control problem with a supremum over t and an infimum over a. While we cannot solve this problem exactly, we may conjecture in analogy to problem $\mathcal{P}$ that for the optimal choice of t the infimum over a is attained for $a \to e^{2(t-1)/(\beta+1)^2}$. We can then determine the value of t for which the conjectured infimum is smallest, which turns out to be unique, and obtain a lower bound on $\mathcal{Q}$ and thus on the approximation guarantee of the best MRS algorithm by substituting this value of t into the simplified control problem and solving the remaining minimization problem over a.

Table 1 shows a selection of bounds obtained in this way for different values of β, along with the choice of t that leads to each bound and the corresponding optimal choice of a. The lower bounds are also shown graphically in Figure 1. For $\beta = 1$ in particular we obtain a lower bound of $\alpha \geq 0.6489$.

We conjecture that the bounds shown in Table 1 are in fact tight up to errors in the numerical approximation. More specifically, we believe that it is without loss of generality to assume that $g(y) = \beta + y$ for all $y < t$ and some $t \in [0, 1]$, $g(t) = \beta + t$, and $g(1) = 0$.

3.5 A Tight Bound for At Most $n/(e-1)$ Samples

Comparison to the upper bound of Correa et al. [2019a] reveals that the bound in the previous section is in fact tight when $\beta = 1/(e-1)$. The optimal MRS algorithm in this case stops at the first value that exceeds all samples, a behavior that should remind us of the optimal algorithm for the case $\beta = 0$, which is identical to the optimal algorithm for the secretary problem. Indeed, $\beta = 1/(e-1)$ implies that $\beta/(\beta+1) = 1/e$, so the only difference to the case $\beta = 0$ is that the algorithm does not need to skip any values because it is given the correct number of samples for free.

For the intermediate case where $0 \le \beta \le 1/(e-1)$, we may now conjecture that we should skip values until the combined number of samples and skipped values amounts to a (rounded) $1/e$ fraction of the combined number of samples and values, and stop at the first value thereafter that exceeds all samples and values seen so far. It turns out that this algorithm matches the upper bound of Correa et al. [2019a], and is thus optimal, when $0 \le \beta \le 1/(e-1)$ and $n \to \infty$. Note that the limit is needed in our analysis to account for possible losses incurred when rounding.

Theorem 3.4. *Let $\beta \le 1/(e-1)$, $n \in \mathbb{N}$, $k = \lfloor \beta n \rfloor$, and $m = \lfloor (\frac{1+\beta}{e} - \beta)n \rfloor$. Consider a sequence of i.i.d. random variables $S_1, \ldots, S_k, X_1, \ldots, X_n$. Let τ be the stopping time of the algorithm that stops at X_j if (i) it has not stopped previously, (ii) $j > m$, and (iii) $X_j > \max\{S_1, \ldots, S_k, X_1, \ldots, X_{j-1}\}$. Then, as $n \to \infty$, $\mathbb{E}[X_\tau] \ge \frac{1+\beta}{e} \cdot \mathbb{E}[\max\{X_1, \ldots, X_n\}]$.*

Proof. To simplify the exposition, we will assume that $\frac{1+\beta}{e} - \beta$ and βn are integers and drop the symbols $\lfloor$ and $\rfloor$. Let $\delta = \frac{1+\beta}{e} - \beta$ and $x \in \mathbb{R}_+$. We will show that for large n and uniformly in F,

$$\Pr(X_\tau \ge x) \ge \left(\frac{1+\beta}{e} + o(1) \right) \cdot \Pr(\max\{X_1, \ldots, X_n\} \ge x). \tag{8}$$

The theorem then follows by integrating over x.

We have

$$\Pr(X_\tau \ge x) = \sum_{i=\delta n+1}^{n} \Pr(\{X_i \ge x\} \cap \{\tau = i\}),$$

and for $i \in \{\delta n+1, \ldots, n\}$, $\Pr(\{X_i \ge x\} \cap \{\tau = i\})$ is equal to

$$\Pr([X_i \ge \max\{x, S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{i-1}\}]$$
$$\cap [\max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{\delta n}\} \ge \max\{X_{\delta n+1}, \ldots, X_{i-1}\}])$$
$$= \Pr(X_i \ge \max\{x, S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{i-1}\})$$
$$\cdot \Pr(\max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{\delta n}\} \ge \max\{X_{\delta n+1}, \ldots, X_{i-1}\}).$$

We have

$$\Pr(\max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{\delta n}\}) \ge \max\{X_{\delta n+1}, \ldots, X_{i-1}\}) = \frac{\beta + \delta}{\beta + \frac{i-1}{n}},$$

and $\Pr(X_i \ge \max\{x, S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{i-1}\})$ is equal to

$$\Pr(X_i \ge x | X_i \ge \max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{i-1}\}) \cdot \Pr(X_i \ge \max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{i-1}\})$$
$$= \Pr(\max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_i\} \ge x) \cdot \Pr(X_i = \max\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_i\})$$
$$= \frac{1 - F^{\beta n+i}(x)}{\beta n + i}.$$

Thus,

$$\Pr(X_\tau \ge x) = (\beta + \delta) \frac{1}{n} \sum_{i=\delta n}^{n} \frac{1 - F^{\beta n+i}(x)}{(\beta + \frac{i}{n})(\beta + \frac{i-1}{n})} \ge (\beta + \delta) \frac{1}{n} \sum_{i=\delta n}^{n} \frac{1 - F^{\beta n+i}(x)}{(\beta + \frac{i}{n})^2}.$$

16

Set $a = F^n(x)$, and notice that $\Pr(\max\{X_1, \ldots, X_n\} \geq x) = 1 - a^n$. Hence, to prove (8), it is enough to prove that for all $a \in [0, 1]$ and large n,

$$(\beta + \delta)\frac{1}{n}\sum_{i=\delta n}^{n} \frac{1 - a^{\beta + \frac{i}{n}}}{(\beta + \frac{i}{n})^2} \geq \left(\frac{(1+\beta)}{e} + o(1)\right)(1-a), \tag{9}$$

where the $o(1)$ is independent of a. Let $g(t) = \frac{1 - a^{\beta + t}}{(\beta + t)^2}$. There exists $C > 0$ such that for all $t \in [0, 1]$ and $a \in [0, 1]$, $|g'(t)| \leq C(1 - a)$. By property of the Riemann integral, it follows that for all $a \in [0, 1]$ and $n \geq 1$,

$$\left|(\beta + \delta)\frac{1}{n}\sum_{i=\delta n}^{n} \frac{1 - a^{\beta + \frac{i}{n}}}{(\beta + \frac{i}{n})^2} - (\beta + \delta)\int_{\delta}^{1} \frac{1 - a^{\beta + t}}{(\beta + t)^2}\, dt\right| \leq \frac{C(1 - a)}{n}.$$

Since $\beta + \delta = \frac{1+\beta}{e}$, to prove (9), it is thus enough to prove that for all $a \in [0, 1]$,

$$\int_{\delta}^{1} \frac{1 - a^{\beta + t}}{(\beta + t)^2}\, dt \geq 1 - a.$$

The above inequality clearly holds for $a = 1$, and thus by the change of variables $t' = t + \beta$, we want to prove that for all $a \in [0, 1)$,

$$\int_{\frac{1+\beta}{e}}^{1+\beta} \frac{1 - a^t}{(1 - a)t^2}\, dt \geq 1.$$

It is enough to prove that the above integral is decreasing in a. Indeed, its limit as a goes to 1 is 1. Define

$$H(a) = \int_{\frac{1+\beta}{e}}^{1+\beta} \frac{1 - a^t}{(1 - a)t^2}\, dt.$$

We have

$$H'(a) = \int_{\frac{1+\beta}{e}}^{1+\beta} \frac{-(1 - a)ta^{t-1} + (1 - a^t)}{(1 - a)^2 t^2}\, dt.$$

Thus, we want to prove that the function I defined by

$$I(a) = \int_{\frac{1+\beta}{e}}^{1+\beta} \frac{-(1 - a)ta^{t-1} + (1 - a^t)}{t^2}\, dt$$

is negative. Notice that

$$I(1) = 0,$$

thus it is enough to prove that I is increasing, which means that I' is positive. We have

$$I'(a) = -\int_{\frac{1+\beta}{e}}^{1+\beta} \frac{a^{t-2}(1 - a)(t - 1)}{t}\, dt.$$

Let

$$J(a) = \int_{\frac{1+\beta}{e}}^{1+\beta} \frac{a^{t-1}(t - 1)}{t}\, dt.$$

Thus, we want to prove that J is negative. For all $a \in [0, 1)$ and $t \in [\frac{1+\beta}{e}, 1 + \beta]$, we have $\frac{a^{t-1}(t-1)}{t} \leq \frac{(t-1)}{t}$, and thus

$$J(a) \leq \int_{\frac{1+\beta}{e}}^{1+\beta} \frac{t - 1}{t}\, dt = \left(1 - \frac{1}{e}\right)(1 + \beta) - 1 \leq \left(1 - \frac{1}{e}\right)\left(1 + \frac{1}{e - 1}\right) - 1 = 0.$$

This finishes the proof. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

4 Streaming Prophet Inequalities

We will now argue that with arbitrarily small additive loss ε in the guarantee, our MRS algorithms can be implemented as streaming algorithms. To this end, consider $x_0 \in [0, 1]$, $\bar{y} \in \mathbb{R}_+$, and the continuous function $g : [0, 1] \to [0, \bar{y}]$ based on which the MRS algorithm is defined. The property that we need our functions g to satisfy (and our algorithms from Section 3 do satisfy) is that the graph of g has $O_\varepsilon(1)$ intersection points with the (infinitely many) horizontal lines at height $0, \varepsilon, 2\varepsilon, \ldots$.

Theorem 4.1. *Let $\varepsilon > 0$. Assume there exists an MRS algorithm with guarantee α for the unknown-distribution setting with $O(n)$ samples. Further assume that the MRS algorithm is based on continuous function g with $|\{x \in [0, 1] : \exists q \in \mathbb{N} : g(x) = q \cdot \varepsilon\}| = O_\varepsilon(1)$. Then there exists a streaming algorithm using $O_\varepsilon(\log n)$ space and achieving a guarantee of $\alpha - \varepsilon$ in the same setting.*

The algorithm divides the x-range $[0, 1]$ into strips of width ε and the y-range $[0, \bar{y}]$ into strips of width ε. This creates $\gamma \leq \lceil 1/\varepsilon \rceil + O_\varepsilon(1)$ intersection points with g. Let $0 = x_1 \leq x_2 \leq \cdots \leq x_\gamma = 1$ be the corresponding x-coordinates of these intersection points.

For all $i = 1, \ldots, \gamma - 1$, the algorithm uses a single threshold that is distributed like the maximum of $\lceil g(x_i) \cdot n \rceil$ fresh samples for all steps in $[x_i \cdot n, x_{i+1} \cdot n)$. We observe that the emerging algorithm can be viewed as an MRS algorithm again. Towards this, let $\tilde{g} : [0, 1] \to [0, \bar{y} + \varepsilon]$ be the function that is equal to $g(x_i)$ at x_i and then grows linearly with slope 1 until (and not including) x_{i+1}. We can (essentially) view the new algorithm as the MRS algorithm based on $\tilde{g}$.

Lemma 4.2. *Let $j \in [x_i \cdot n, x_{i+1} \cdot n) \cap \mathbb{Z}$ for some $i \in \{1, \ldots, \gamma - 1\}$. Conditioned on arriving in step j, the above algorithm sets a threshold for X_j that is distributed like the maximum of $\tilde{g}(j/n) \cdot n \pm O(1)$ fresh samples.*

Proof. Let j_0 be the first integer in $[x_i \cdot n, x_{i+1} \cdot n)$. Denote the subset of values selected uniformly at random from $\{S_1, \ldots, S_{\beta n}, X_1, \ldots, X_{j_0-1}\}$ by $Y_S = \{S'_1, \ldots, S'_\ell\}$ where $\ell = \lceil g(j_0/n) \cdot n \rceil = \tilde{g}(j_0/n) \cdot n \pm O(1)$ by continuity of g. By Theorem 3.3, this set and the set $Y = \{S'_1, \ldots, S'_\ell, X_{j_0}, \ldots, X_{j-1}\}$ are distributed like sets of ℓ and $|Y| = \ell + (j - j_0) - 1$, respectively, fresh samples. Note that $|Y| = \tilde{g}(j/n) \cdot n \pm O(1)$. It suffices to show that, conditioned on arriving in step j_0 and any such set Y, (i) the probability of arriving in step j is independent of Y, and (ii), if the algorithm arrives in step j, the threshold it sets in step j is $\max Y$.

Towards showing (i) and (ii), again condition on arriving in step j_0 and any set Y. Note that the algorithm arrives in step j if and only if $\max Y = \max Y_S$. This implies that, throughout steps $j_0, \ldots, j$, the algorithm sets $\max Y$ as threshold, showing (ii). Finally notice that, since both Y_S and $Y \setminus Y_S$ are sets of fresh draws from F, $\max Y = \max Y_S$ happens with probability independent of Y, showing (i). $\qquad\square$

Further note that our construction ensures that $g(x) \leq \tilde{g}(x) \leq g(x) + 2\varepsilon$ for all $x \in [0, 1]$. See Figure 2 for a visualization of the construction.

Lemma 4.3. *Suppose that the MRS algorithm defined by g achieves approximation ratio α via (2) and that $|\tilde{g}(x) - g(x)| \leq 2\varepsilon$ for all $x \in [0, 1]$, then the MRS algorithm defined by $\tilde{g}$ achieves approximation ratio $\alpha - O(\sqrt{\varepsilon})$.*

Proof. As shown previously, g satisfies (2): For all $a \in [0, 1]$,

$$\int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, \mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \left(1 - a^{g(y)}\right) \mathrm{d}y \geq \alpha \cdot (1 - a). \tag{10}$$

Moreover, to prove our claim, it is enough to show that the above equation holds, replacing g by $\tilde{g}$ and α by $\alpha - O(\sqrt{\varepsilon})$. Note that, as $\varepsilon \to 0$ and uniformly in a,

$$(1-a)^{-1} \int_{1-\varepsilon}^{1} \exp\left(-\int_0^y \frac{1}{g(z)}\,\mathrm{d}z\right) \cdot \frac{1}{g(y)} \cdot \left(1 - a^{g(y)}\right)\,\mathrm{d}y = O(\sqrt{\varepsilon}). \qquad (11)$$

Second, for $y \in [0, 1-\varepsilon]$, we have $g(y) \geq \tilde{g}(y) - 2\varepsilon \geq \tilde{g}(y) \times (1 - 2\varepsilon/\tilde{g}(y)) \geq \tilde{g}(y)(1 - 2\varepsilon/g(1-\varepsilon)) = \tilde{g}(y)(1 - O(\sqrt{\varepsilon}))$, and thus $1/\tilde{g}(y) \geq 1/g(y) - O(\sqrt{\varepsilon})$. Hence, by (10) and (11), as ε tends to 0 and uniformly in a,

$$(1-a)^{-1} \int_0^{1-\varepsilon} \exp\left(-\int_0^y \frac{1}{\tilde{g}(z)}\,\mathrm{d}z\right) \cdot \frac{1}{\tilde{g}(y)} \cdot \left(1 - a^{\tilde{g}(y)}\right)\,\mathrm{d}y \geq \alpha - O(\sqrt{\varepsilon}).$$

To obtain the above inequality, we have used in addition the fact that the left-hand side term and right-hand side term in the integrand of equation (10) increase when one replaces g by $\tilde{g}$. This completes the proof. $\qquad \square$

To implement our approach as a streaming algorithm, for each $i = 0, \ldots, \gamma$, we construct the maximum of the corresponding random subset on the fly: We count how many random positions are left to consider and include the current position with probability proportional to that count.

Lemma 4.4. *For each x_i the threshold corresponding to $g(x_i)$ can be computed with a single pass over the data and $O(\log n)$ space.*

Proof. Consider the first j such that $j/n \geq x_i$ and let $q = \lceil g(x_i) \cdot n \rceil$. We will construct a 0/1-vector of length $k + j - 1$ with exactly q many 1's on the fly such that the positions where the bit vector is 1 correspond to a subset of size q chosen uniformly at random without replacement from $\{S_1, \ldots, S_k, X_1, \ldots, X_{j-1}\}$. We can then compute the threshold in an online fashion by remembering the maximum T of all values where we have set the bit to 1.

We do this as follows: We remember the number s of 1's that we still need and the number of positions t still to come. Initially, $s = q$ and $t = k + j - 1$. Then for $\ell = 1$ to $k + j - 1$ we toss a biased coin that comes up 1 with probability s/t and is 0 otherwise. If it comes up 1 we update $s = s - 1$ and $t = t - 1$, otherwise we keep s and just set $t = t - 1$.

It now suffices to show that this process always yields a 0/1-vector of length $k + j - 1$ with exactly q many 1's, and that all such bit vectors are equally likely. The former follows from the fact that the probability of seeing another 1 is set to zero once there are already q many 1's and that once the number of remaining positions equals the number of 1's that are still needed the probability of seeing a 1 is set to one for all remaining steps.

It remains to show that all 0/1-vectors of length $k + j - 1$ with q many 1's are equally likely, i.e., that the likelihood of seeing any such vector vector is $1/\binom{k+j-1}{q}$. Indeed, consider an arbitrary such vector z. Let $E = \{e_1, e_2, \ldots, e_q\} \subseteq [k+j-1]$ with $e_1 < e_2 < \cdots < e_q$ be the indices ℓ where $z_\ell = 1$ and let $N = \{n_1, \ldots, n_{k+j-1-q}\}$ with $n_1 < \cdots < n_{k+j-1-q}$ be the indices ℓ where $z_\ell = 0$. Then,

$$\Pr[z] = \prod_{\ell=1}^{q} \frac{q - \ell + 1}{k + j - e_\ell} \cdot \prod_{\ell=1}^{k+j-1-q} \frac{k + j - q - \ell}{k + j - n_\ell} = \frac{1}{\binom{k+j-1}{q}}.$$

The space complexity is $O(\log(n))$ because all the algorithm needs to store is the threshold, the remaining number of positions, and the number of ones that are still required. $\qquad \square$

Theorem 4.1 then follows by combining the above lemmata.

Figure 2: Visualization of the g function from Section 3.3 (bold, black), the grid with $\varepsilon = 0.2$ (dotted, black), and the resulting $\tilde{g}$ function (green, bold).

5 Unknown Distribution With Prior

Let θ be a distribution over a set of real distributions $\{F^1, \ldots, F^m\}$, where $m \geq 1$. Consider the following learning problem, which we will call *unknown i.i.d. problem with prior θ*:

- Nature draws F according to θ and $(X_1, \ldots, X_n)$ i.i.d. according to F;

- the gambler knows θ but not F, and has to choose a stopping rule with stopping time τ in order to maximize $\mathbb{E}(X_\tau)$.

The key difference with the unknown i.i.d. problem without samples studied by Correa et al. [2019a] is that the gambler *knows* θ. Thus, it is plausible that a constant better than $1/e$ can be guaranteed in this model. Our main result in this section shows that $1/e$ remains best possible.

Theorem 5.1. *For any $\delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$, there exists a finite number of finitely supported distributions $F^1, \ldots, F^m$ and a distribution θ over $\{F^1, \ldots, F^m\}$ such that for any stopping rule with stopping time τ,*

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

An important consequence of Theorem 5.1 is that it answers a question of Hill and Kertz [1992]. Indeed, in their survey on the state of the art in prophet inequalities, Hill and Kertz stated five groups of open questions. The second group, Question 2 on Page 203, concerns exchangeable random variables. Regarding the maximum value obtainable by a stopping rule τ and the expected maximum value in the sequence they specifically asked for the largest universal constant γ such that, if $X_1, X_2, \ldots, X_n$ are exchangeable random variables taking values in $[0, 1]$ then

$$\mathbb{E}[X_\tau] \geq \gamma \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right]. \tag{12}$$

We resolve this question by showing that the universal constant is $\gamma = 1/e \approx 0.368$. Prior to this work it was only known that the tight approximation ratio for $n = 2$ is $1/2$, that the ratio is weakly decreasing in n, and that it is at least $1/e$ [Elton and Kertz, 1991].

Note that the assumption that $X_1, \ldots, X_n$ take value in $[0, 1]$ is actually without loss of generality. Indeed, given non-negative random variables $X_1, \ldots, X_n$ and $M \geq 1$, one can define $X_i' = (X_i/M)1_{X_i \leq M}$ for all $i \in [n]$. The variables $X_1', \ldots, X_n'$ then are exchangeable and take values in $[0, 1]$, and for M large enough, the best approximation guarantees achievable for both instances are almost the same. We will therefore omit the assumption when stating our results.

Corollary 5.2. *For any $\delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$, there exists a sequence of exchangeable random variables $X_1, X_2, \ldots, X_n$ such that for any stopping rule with stopping time τ,*

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

In particular, the largest constant γ in (12) is $1/e$.

Proof. Let $\delta > 0$, $n_0 \in \mathbb{N}$, $n \geq n_0$, $m \geq 1$, $F^1, \ldots, F^m$ and θ be given by Theorem 5.1. Let F be a random variable distributed according to θ, and $(X_1, \ldots, X_n)$ a sequence of random variables such that its distribution conditional to F is the one of i.i.d. random variables from F. The random variables $(X_1, \ldots, X_n)$ are exchangeable, and by Theorem 5.1, no stopping rule can guarantee a constant better than $1/e + \delta$. $\qquad\square$

The proof of Theorem 5.1 extends on a technique developed recently by Correa et al. [2019a] to show a lower bound of $1/e$ for unknown i.i.d. random variables with finite support and without any additional samples. We modify the central element of this technique, which relies on the infinite version of Ramsey's theorem, to instead use the finite version (Section 5.1). We then deduce the desired impossibility using a minimax argument (Section 5.2).

5.1 Hard Finite Instances for Unknown Distribution

Let us return to the setting considered by Correa et al. [2019a], where random variables are drawn independently from the same unknown distribution and we do not have access to any additional samples and there is no prior over the unknown distribution. For $p \in \mathbb{N}$, denote by B_p the set of probability measures on $[p]$. We prove the following result.

Proposition 5.3. *For all $\delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$, there exists $p \in \mathbb{N}$ such that for any stopping rule $\mathbf{r}$ with associated stopping time τ, there exists $b \in B_p$ such that when $X_1, \ldots, X_n$ are i.i.d. random variables drawn from b,*

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

What distinguishes this result from Theorem 3.2 of Correa et al. is that the distribution b has finite support and the cardinality of its support is independent of the stopping rule. This property will be crucial when applying the minimax theorem in Section 5.2.

The result can be obtained by modifying the central construction in the proof of Correa et al. while keeping much of the structure of that proof intact. To make it easier to compare the two results we will follow the original structure, and begin by recalling the definition of oblivious stopping rules.

Definition 1. *Let $\varepsilon > 0$ and $V \subset \mathbb{N}$.*

- *A stopping rule $\mathbf{r}$ is (ε, i)-value-oblivious on V if, there exists a $q_i \in [0, 1]$ such that, for all pairwise distinct $v_1, \ldots, v_i \in V$ with $v_i > \max\{v_1, \ldots, v_{i-1}\}$, it holds that $r_i(v_1, \ldots, v_i) \in [q_i - \varepsilon, q_i + \varepsilon)$.*

- *A stopping rule $\mathbf{r}$ is ε-value-oblivious on V if, for all $i \in [n]$, it is (ε, i)-value-oblivious on V.*

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

The cornerstone of our proof is the following lemma.

Lemma 5.4. *Let $\varepsilon > 0$. For any $n \in \mathbb{N}$, there exists $p \in \mathbb{N}$ such that if there exists a stopping rule with guarantee α, then there exists a stopping rule $\mathbf{r}$ with guarantee α such that $\mathbf{r}$ is ε-value-oblivious on V, for some finite set $V \subset [p]$ with cardinality $n^3 + 1$.*

The difference to Lemma 3.4 of Correa et al. is that the set V is finite, and in addition is uniformly bounded by an integer p that depends only on n. Consequently, instead of the infinite version of Ramsey's theorem used by Correa et al., we need the following finite version given for example by Conlon et al. [2010].

Lemma 5.5. *There exists a function $R : \mathbb{N}^3 \to \mathbb{N}$ such that for all $n \geq 1$, for all complete m-hypergraph with c colors and order larger than $R(m, n, c)$, there exists a sub-hypergraph of order n that is monochromatic.*

Proof of Theorem 5.4. Fix $\varepsilon > 0$ and set $c = \lfloor (2\varepsilon)^{-1} \rfloor$. Define an integer sequence $(p_i)_{0 \le i \le n}$ by induction as $p_n = n^3 + 1$ and $p_{i-1} = R(i, p_i, c)$. Consider a stopping rule $\mathbf{r}$ with guarantee α. By Lemma 3.6 of Correa et al. it is without loss of generality to assume that $\mathbf{r}$ is order-oblivious. We show by induction on $j \in \{0, 1, \ldots, n\}$ that there exists a set $S^j \subset S^0$ such that $|S^j| = p_j$ and for all $i \in [j]$, $\mathbf{r}$ is (ε, i)-value-oblivious on S^j.

The set $S_0 = \left\{ n^{3s} : s = 0, \ldots, p_0 \right\}$ satisfies the induction hypothesis for $j = 0$. We proceed to show it for $j > 0$. First, observe that we only need to find a set $S^j \subset S^{j-1}$ such that $|S_j| = p_j$ and $\mathbf{r}$ is (ε, j)-value oblivious on S^j, because it follows from the induction hypothesis that for all $i \in [j-1]$, $\mathbf{r}$ is (ε, i)-value-oblivious on S^i and thus on the subset $S^j \subset S^i$.

Toward the application of Theorem 5.5, we construct a complete j-hypergraph H with vertex set S^{j-1}. Consider any set $v_1, \ldots, v_j \subset S^{j-1}$ of cardinality j such that $v_j > max(v_1, \ldots, v_{j-1})$. Note that there exists a unique $u \in \{1, 2, \ldots, c\}$ such that $r_j(v_1, \ldots, v_j) \in [(2u-1)\varepsilon - \varepsilon, (2u-1)\varepsilon + \varepsilon)$, and color the hyperedge $\{v_1, \ldots, v_j\}$ of H with color u. By Theorem 5.5, there exists a finite set S^j of vertices with cardinality p_j that induces a complete monochromatic sub-hypergraph of H. Let u be the color of this sub-hypergraph, set $q = (2u-1)\varepsilon$, and consider distinct $v_1, \ldots, v_j \in S^j$ with $v_j > max(v_1, \ldots, v_{j-1})$. Since the edge $\{v_1, \ldots, v_j\}$ in H has color u, $r_j(v_{\pi(1)}, \ldots, v_{\pi(j-1)}, v_j) \in [q - \varepsilon, q + \varepsilon)$ for some permutation π of S^{j-1}. But since $\mathbf{r}$ is order-oblivious, also $r_j(v_1, \ldots, v_{j-1}, v_j) \in [q - \varepsilon, q + \varepsilon)$. So $\mathbf{r}$ is (ε, j)-value oblivious on S^j. This completes the induction step. $\qquad\square$

With Theorem 5.4 at hand we are now ready to prove Theorem 5.3.

Proof of Theorem 5.3. Let $\delta > 0$ and $n \in \mathbb{N}$. We proceed by contradiction, and consider a stopping rule $\mathbf{r}$ with performance guarantee $1/e + \delta$. Set $\varepsilon = 1/n^2$. By Theorem 5.4, there exists a stopping rule $\mathbf{r}$ with performance guarantee $1/e + \delta$ and a set $V \subset [p]$ with cardinality $n^3 + 1$ on which $\mathbf{r}$ is ε-value-oblivious. Let u be the maximum of V, and write $V = \{v_1, \ldots, v_{n^3}, u\}$. Denote by τ the stopping time of $\mathbf{r}$. By construction, we have $u \ge n^3 \max\{v_1, \ldots, v_{n^3}\}$. The rest of the proof proceeds in the same way as the proof of Theorem 3.2 of Correa et al., and we give an informal summary for completeness. For each $i \in [n]$, let

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

For this particular instance, the performance of any stopping rule corresponds approximately to the probability of picking u. Let us, therefore, investigate the probability that $\mathbf{r}$ picks u. First note that with probability almost one, $X_1, \ldots, X_n$ are distinct. Moreover, because $\mathbf{r}$ is ε-value-oblivious on V, it can be changed with an error of $\varepsilon = n^{-2}$ into a stopping rule that considers only the relative ranks of the values it has seen before making its decision. As there are only n stages, the error is insignificant.

The problem thus reduces to the classic secretary problem, for which it is known that no stopping rule can guarantee a probability of picking the maximum that is higher than $1/e + o(1)$ as n goes to infinity [Ferguson, 1989]. However, the stopping rule constructed from $\mathbf{r}$ considers only relative ranks and selects the maximum with probability at least $1/e + \delta - o(1)$, which is a contradiction. $\qquad\square$

5.2 From Unknown to Unknown with Prior: A Minimax Argument

We now use a minimax argument to convert Theorem 5.3, which concerns the unknown i.i.d. case, into an upper bound (impossibility result) for the case of *unknown i.i.d. with prior*. Fix $\delta > 0$ and let n_0 and p be as in Theorem 5.3. Fix $n \ge n_0$. We need the following definitions.

Definition 2.

- *A deterministic stopping rule is a sequence $a = (a_1, ..., a_n)$ such that $a_i : [p]^{i-1} \to \{0,1\}$. Denote the set of such rules by A.*

- *A mixed stopping rule is a distribution over A. Denote the set of such rules by $\mathcal{P}(A)$.*

- *A behavior stopping rule is a sequence $\mathbf{r} = (r_1, ..., r_n)$ such that $r_i : [p]^{i-1} \to [0,1]$. Denote the set of such rules by C.*

Note that we have defined each class of stopping rule to consider only values in $[p]$ as inputs because we will consider random variables supported on $[p]$. As before, considering inputs in $\mathbb{R}$ would not have a significant effect on the performance guarantee. Note further that behavior stopping rules correspond to stopping rules as defined in Section 2. The minimax argument will require us to consider mixed stopping rules. We will apply Kuhn's theorem [Kuhn, 1953] to prove that mixed stopping rules and behavior stopping rules provide the same performance guarantee.

According to our purpose any of these stopping rules will have the two interpretations, (i) as a stopping rule in the unknown i.i.d. problem, and (ii) as a stopping rule in the i.i.d. problem with some prior distribution θ, where in the latter case we have omitted the dependence of a on θ.

Recall that B_p is the set of probability distributions over $[p]$. Because p has been fixed and for ease of exposition, we will henceforth write B instead of B_p. Let $\mathcal{P}(B)$ be the set of probability distributions over B. For given $a \in A$ and $b \in B$, define

$$g(a,b) = \mathbb{E}\left[X_\tau\right] - \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right],$$

where τ is the stopping time of the stopping rule a in the i.i.d. problem where $X_1, X_2, ..., X_n$ are drawn from b. We can extend g linearly to $\mathcal{P}(A) \times \mathcal{P}(B)$ by letting

$$g(x,y) = \int_A \int_B g(a,b)x(\mathrm{d}a)y(\mathrm{d}b).$$

Let

$$V^- = \max_{x \in \mathcal{P}(A)} \min_{b \in B} g(x,b) \quad \text{and} \quad V^+ = \min_{y \in \mathcal{P}(B)} \max_{a \in A} g(a,y).$$

Note that in the above expressions, by linearity of g with respect to x and y, $\min_{b \in B} g(x,b) = \min_{y \in \mathcal{P}(B)} g(x,y)$ and $\max_{a \in A} g(a,y) = \max_{x \in \mathcal{P}(A)} g(x,y)$. The key point is that V^- is related to the universal constant in the unknown i.i.d. problem, while V^+ is related to the universal constant in the unknown i.i.d. problem with prior. Indeed, the following proposition holds:

Proposition 5.6. *If $V^+ < 0$, then there exists $y \in \mathcal{P}(B)$ a finitely supported distribution over B such that, in the unknown i.i.d. problem with prior y, for all behavior stopping rule with stopping time τ,*

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

Proof. Assume that $V^+ < 0$. By [Mertens et al., 2015, Proposition I.1.9], there exists $y \in \mathcal{P}(B)$ with finite support such that $\max_{a \in A} g(a,y) \leq 0$. This means that no deterministic stopping rule provides a better guarantee than $(1/e + \delta)$ in the unknown i.i.d. problem with prior distribution y. Then no distribution over deterministic stopping rule can provide a better guarantee by linearity of expectation, neither can a behavior stopping rule by Kuhn's theorem [Kuhn, 1953], which proves the result. $\square$

To complete the proof of Theorem 5.1 it is thus enough to show that $V^+ < 0$. To this aim, we first prove that $V^+ = V^-$, and then that $V^- < 0$.

Proposition 5.7. *We have $V^+ = V^-$.*

Proof. The set A is finite, the set B compact metric. Moreover, for all $a \in A$, the mapping $b \to g(a, b)$ is continuous. By the minimax theorem [Mertens et al., 2015, Proposition I.1.9] it follows that the mixed extension of the normal-form game (A, B, g) has a value, i.e., that $V^+ = V^-$. $\qquad\square$

To conclude the proof of Theorem 5.1 it remains to show that $V^- < 0$. First note that by Kuhn's theorem [Kuhn, 1953],

$$V^- = \max_{r \in C} \min_{b \in B} \mathbb{E}\left[X_\tau\right] - \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right].$$

By definition of p, which we have chosen as in Theorem 5.3, for each stopping rule in C there exists $b \in B$ such that

$$\mathbb{E}\left[X_\tau\right] - \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}\left[\max\{X_1, \ldots, X_n\}\right] < 0.$$

It follows that $V^- < 0$, which proves Theorem 5.1.

References

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

S. Agrawal, J. Sethuraman, and X. Zhang. On optimal ordering in the optimal stopping problem. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 187–188, 2020.

N. Ailon, B. Chazelle, K. L. Clarkson, D. Liu, W. Mulzer, and C. Seshadhri. Self-improving algorithms. *SIAM Journal on Computing*, 40(2):350–375, 2011.

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

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

Y. Azar, A. Chiplunkar, and H. Kaplan. Prophet secretary: Surpassing the $1 - 1/e$ barrier. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pages 303–318, 2018.

M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. Posting prices with unknown distributions. *ACM Transactions on Economics and Computation*, 5(2):13:1–13:20, 2017.

A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause. Streaming submodular maximization: massive data summarization on the fly. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 671–680, 2014.

M. Balcan, T. Dick, and E. Vitercik. Dispersion for data-driven algorithm design, online learning, and private optimization. In M. Thorup, editor, *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science*, pages 603–614, 2018a.

M. Balcan, T. Dick, and C. White. Data-driven clustering via parameterized Lloyd's families. In *Proceedings of the 2018 Conference on Neural Information Processing Systems*, pages 10664–10674, 2018b.

S. Chawla, J. B. Miller, and Y. Teng. Pricing for online resource allocation: Intervals and paths. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1962–1981, 2019.

R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In D. B. Shmoys, editor, *Proceedings of the 46th Symposium on Theory of Computing*, pages 243–252, 2014.

D. Conlon, J. Fox, and B. Sudakov. Hypergraph Ramsey numbers. *Journal of the American Mathematical Society*, 23(1):247–266, 2010.

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

J. Correa, P. Dütting, F. Fischer, and K. Schewior. Prophet inequalities for i.i.d. random variables from an unknown distribution. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pages 3–17, 2019a.

J. Correa, R. Saona, and B. Ziliotto. Prophet secretary through blind strategies. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1946–1961, 2019b.

J. Correa, A. Cristi, B. Epstein, and J. A. Soto. The two-sided game of googol and sample-based prophet inequalities. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2066–2081. SIAM, 2020.

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

P. Dütting and T. Kesselheim. Posted pricing and prophet inequalities with inaccurate priors. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pages 111–129, 2019.

P. Dütting and R. D. Kleinberg. Polymatroid prophet inequalities. In *Proceedings of the 23rd European Symposium on Algorithms*, pages 437–449, 2015.

P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49:540–582, 2020a.

P. Dütting, T. Kesselheim, and B. Lucier. An O(log log m) prophet inequality for subadditive combinatorial auctions. In *Proceedings of the 61st IEEE Symposium on Foundations of Computer Science*, 2020b. forthcoming.

S. Ehsani, M. T. Hajiaghayi, T. Kesselheim, and S. Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–714, 2018.

J. Elton and R. Kertz. Comparison of stop rule and maximum expectations for finite sequences of exchangeable random variables. *Stochastic Analysis and Applications*, 9:1–23, 1991.

H. Esfandiari, M. T. Hajiaghayi, V. Liaghat, and M. Monemizadeh. Prophet secretary. In *Proceedings of the 23rd European Symposium on Algorithms*, pages 496–508, 2015.

T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang. Online stochastic max-weight matching: Prophet inequality for vertex and edge arrival models. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 769–787, 2020.

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

M. Feldman, A. Norouzi-Fard, O. Svensson, and R. Zenklusen. The one-way communication complexity of submodular maximization with applications to streaming and robustness. In *Proccedings of the 52nd Annual ACM Symposium on Theory of Computing*, pages 1363–1374, 2020.

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

A. Goldenshluger and A. Zeevi. Optimal stopping of a random sequence with unknown distribution. Working paper, 2017.

N. Gravin and H. Wang. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pages 93–109, 2019.

S. Guha and A. McGregor. Stream order and order statistics: Quantile estimation in random-order streams. *SIAM Journal on Computing*, 38(5):2044–2059, 2009.

R. Gupta and T. Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.

M. Hajiaghayi and V. Liaghat. Prophet inequality and online auctions. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*. Springer, 2016.

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

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

T. P. Hill and R. P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *Annals of Probability*, 10(2):336–345, 1982.

N. Immorlica, S. Singla, and B. Waggoner. Prophet inequalities with linear correlations and augmentations. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 159–185, 2020.

H. Kaplan, D. Naori, and D. Raz. Competitive analysis with a sample and the secretary problem. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2082–2095, 2020.

R. D. Kleinberg and S. M. Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, 2019.

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

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

H. W. Kuhn. Extensive games and the problem of information. In *Annals of Mathematics Studies*, volume 28, pages 193–216. Princeton University Press, 1953.

K.-H. Li. Reservoir-sampling algorithms of time complexity o(n(1+log(n/n))). *ACM Transactions on Mathematical Software*, 20:481–493, 1994.

J.-F. Mertens, S. Sorin, and S. Zamir. *Repeated Games*. Cambridge University Press, 2015.

J. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 2015 Conference on Neural Information Processing Systems*, pages 136–144, 2015.

S. Muthukrishnan. Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science*, 1(2):117–236, 2005.

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

Y. Rinott and E. Samuel-Cahn. Comparisons of optimal stopping values and prophet inequalities for negatively dependent random variables. *The Annals of Statistics*, 15(4):1482–1490, 12 1987.

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

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

A. Rubinstein, J. Z. Wang, and S. M. Weinberg. Optimal single-choice prophet inequalities from samples. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference*, pages 60:1–60:10, 2020.

E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12:1213–1216, 1984.

J. S. Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 11(1):37–57, 1985.

D. H. D. West. Updating mean and variance estimates: An improved method. *Communications of the ACM*, 22(9):532–535, 1979.