



HAL
open science

Streaming Algorithms for Online Selection Problems

José Correa, Paul Dütting, Felix Fischer, Kevin Schewior, Bruno Ziliotto

► **To cite this version:**

José Correa, Paul Dütting, Felix Fischer, Kevin Schewior, Bruno Ziliotto. Streaming Algorithms for Online Selection Problems. Innovations in Theoretical Computer Science, 2021. hal-02920558

HAL Id: hal-02920558

<https://hal.science/hal-02920558>

Submitted on 29 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Streaming Algorithms for Online Selection Problems

José Correa* Paul Dütting† Felix Fischer‡ Kevin Schewior§ Bruno Ziliotto¶

July 12, 2020

Abstract

The model of streaming algorithms is motivated by the increasingly common situation in which the sheer amount of available data limits the ways in which the data can be accessed. Streaming algorithms are typically allowed a single pass over the data and can only store a sublinear fraction of the data at any time. We initiate the study of classic online selection problems in a streaming model where the data stream consists of two parts: historical data points that an algorithm can use to learn something about the input; and data points from which a selection can be made. Both types of data points are i.i.d. draws from an unknown distribution. We consider the two canonical objectives for online selection—maximizing the probability of selecting the maximum and maximizing the expected value of the selection—and provide the first performance guarantees for both these objectives in the streaming model.

1 Introduction

In an online selection problem, items arrive sequentially and have to be accepted or rejected at the time of their arrival and thus with partial information about their value. Online selection has both a substantial theory that goes back more than half a century [e.g., 12, 15] and a wide range of contemporary applications. The latter can for example be found in online retail, where consumers arrive over time and items must be priced appropriately, on online platforms for ride sharing or matching, and in advertising auctions with dynamic reserve prices.

The most fundamental problems in online selection are the secretary problem and the prophet problem. In both problems a decision maker is presented with a sequence of numerical values, and for each value has to make an irreversible decision to accept or reject it. Only one value can be accepted, and a value that has been rejected is lost forever. In the secretary problem the objective is to maximize the probability of selecting the largest value of the entire sequence, while in the prophet problem it is to maximize the ratio between the expected value of the selection and the expected value of the maximum. Both problems have been studied extensively under the assumption that values are drawn independently from a known distribution, and respective optimal performance guarantees of 0.580 [15] and 0.745 [19, 11] are known.

In practice, full knowledge of underlying distributions is of course an unrealistic assumption. Instead, relevant information will often be available in the form of historical data, from which a distribution could potentially be learned. Recent research has thus considered variants of the secretary and prophet problems where values are i.i.d. draws from an unknown distribution and the decision maker has access to additional samples from that distribution [9, 27, 10, 20]. An important additional characteristic shared by many current applications is that the sheer amount of historical data may make it impossible or at least impractical to refer explicitly to all such data, or indeed to store this data in perpetuity. And even in cases where complete historical data sets are a technological possibility, laws such as the General Data Protection Regulation may

*Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile, email: correa@uchile.cl

†Google Research, Brandschenkestrasse 110, CH-8002 Zürich, Switzerland, email: duetting@google.com and Department of Mathematics, London School of Economics, London, UK, email: p.d.duetting@lse.ac.uk

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

§Department Mathematik/Informatik, Universität zu Köln, Cologne, Germany, email: kschewior@gmail.com

¶Centre de Recherche en Mathématiques de la Décision, Université Paris-Dauphine, Paris, France, email: ziliotto@math.cnrs.fr

only allow aggregate data to be stored. These restrictions motivate the study of online selection problems in the streaming model of computation.

In the streaming model, the problem input cannot be accessed in an arbitrary way but instead arrives sequentially as a stream [e.g., 3, 14]. Algorithms are allowed only a small number of passes over the stream and in addition are subject to a space constraint. The amount of available space is typically assumed to be logarithmic both in the length of the stream and the size of the domain of values in the stream. A requirement to pass the stream of data only a small number of times, indeed to pass it only once, is of course very natural in online selection. Restrictions on the amount of space available to store information about past data have to our knowledge not previously been considered.

1.1 Results

We study the two canonical problems of online selection, the secretary and the prophet problem, in a streaming model and provide the first performance guarantees for both.

For the prophet problem, we consider a specific natural class of algorithms and show that they can be implemented with arbitrarily small loss in the streaming model. Algorithms from the class, which we term MRS algorithms, are characterized by their use of maxima of random subsets of samples and past values as a threshold for acceptance of the current value. They can be defined in terms of a function f such that $f(i)$ is the size of the subset considered at step i .

Given an MRS algorithm with performance guarantee α , we construct a streaming algorithm that uses the same thresholds as that MRS algorithm at $O_\epsilon(1)$ “supporting” values by constructing the random subsets on the fly. It then continues to use each of these thresholds until the next supporting value. By carefully choosing which values are supporting values we guarantee a notion of uniform convergence between the streaming algorithm, which can also be viewed as a different MRS algorithm, and the original MRS algorithm. This turns out to be enough to achieve a guarantee of $\alpha - \sqrt{\epsilon}$.

Two particular MRS algorithms were studied previously by Correa et al. [9] and by Kaplan et al. [20], and our result translates the performance guarantees of these algorithms to the streaming model. We will, however, determine the optimal MRS algorithm and obtain stronger guarantees. To this end we cast the choice of the function f that determines the sizes of the random subsets as a problem of optimal control, and solve the control problem using tools from variational calculus.

We specifically derive the optimal MRS algorithm that is not constrained by the number of samples it can access. This algorithm achieves a performance guarantee of approximately 0.653 and can be realized with around $1.443 \cdot n$ samples, where n is the number of values. We then derive lower bounds on the performance of MRS algorithms with access to $\beta \cdot n$ samples where $\beta < 1.443$. We conjecture these lower bounds to be tight up to errors in the numerical approximation. Finally we analyze a variant of MRS algorithms allowed to skip values unconditionally, and show that the optimal such algorithm is in fact optimal over the class of all algorithms when there are at most $n/(e-1) \approx 0.58 \cdot n$ samples.

Our results for the prophet problem also apply to the model of Correa et al. [9] and Kaplan et al. [20], which is not subject to the restrictions of streaming, and improve on the best known results for this model.

For the secretary problem, we provide an implementation of an algorithm due to Correa et al. [10] in the streaming model. The algorithm uses as thresholds the maxima of a sliding window containing samples and values, and it turns out that only a small number of samples and values need to be stored to determine these thresholds. We thus translate the performance guarantee of the algorithm, which is equal to around 0.452 when the numbers of samples and values are the same, to the streaming model.

All aforementioned results apply to i.i.d. draws from an unknown distribution, which are closely related to the concept of exchangeability.¹ Indeed, De Finetti’s Theorem establishes that infinite sequences of exchangeable random variables are mixtures of i.i.d. random variables. However, for finite sequences, exchangeability defines a larger class. Our final result, given in the appendix, uses a maximin argument to establish a tight bound of $1/e$ for the prophet inequality problem for sequences of exchangeable random variables with a known joint distribution. This resolves a question of Hill and Kertz [18] that had been open for quite some time.

¹A sequence of random variables is exchangeable if its joint distribution is invariant under finite permutations.




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. [9] as well as the lower bounds of Kaplan et al. [20] (blue) and Correa et al. [10] (violet).

1.2 Related Work

Prophets and Secretaries With and Without Samples. The classic work of Dynkin [12] establishes a performance guarantee of $1/e$ for the secretary problem. The result in fact applies to values that are chosen by an adversary, as long as they are presented in random order. Perhaps less known is a beautiful impossibility result which shows that the guarantee of $1/e$ is best possible even if values are i.i.d. draws from an unknown distribution [6, 13].

Recent work by Correa et al. [9] shows that for i.i.d. draws from an unknown distribution, the same tight bound of $1/e$ also applies to the prophet problem. The impossibility result in fact applies to any algorithm with access to $o(n)$ samples. The situation changes with access to $\Omega(n)$ samples, and Correa et al. [9] give an algorithm that achieves improved guarantees with $\beta \cdot n$ samples for $\beta > 0$. They also gave a lower bound for algorithms that have access to $\beta \cdot n$ samples that is equal to $1/e$ for $\beta \leq 1/e$ and then increases linearly to $1 - 1/e \approx 0.632$ at $\beta = 1$; and a parametric upper bound that is equal to $\frac{1+\beta}{e}$ for $\beta \leq 1/(e-1)$ and equal to $-\beta \cdot \ln(\beta/(1+\beta))$ for $\beta \geq 1/(e-1)$. Note that this yields an upper bound of $\ln(2) \approx 0.693$ at $\beta = 1$.

Kaplan et al. [20] improved the lower bound for $\beta < 1$ by proving a bound of $e^{-e^{-\beta}}$ for $\beta \lesssim 0.567$ and $\beta \cdot (1 - \ln(\beta) - e^{-\beta})$ for $0.567 \lesssim \beta \leq 1$. Note that this matches the previously known bound of $1 - 1/e$ at $\beta = 1$. Moreover, except for $\beta = 0$ this lower bound does *not* match the parametric upper bound of [9]. In parallel work, Correa et al. [10] provided an improved bound of ≈ 0.635 for $\beta = 1$. The results of Kaplan et al. [20] and Correa et al. [10] actually apply to slightly more general models than the unknown i.i.d. model, but they do not show whether their algorithms are implementable as streaming algorithms.

Our bounds improve on the state-of-the-art approximation guarantees for all β , and also show that these bounds are achievable in the streaming model. Specifically, Theorem 3 provides a *tight* answer for $\beta \leq 1/(e-1)$ while Theorem 1 and Theorem 2 provide improved bounds for $1/(e-1) < \beta \lesssim 1.443$. This improvement is bigger for larger β . For the particularly interesting case of $\beta = 1$ Theorem 2 shows a lower bound of 0.649. For a visualization of the various bounds see Figure 1.

Correa et al. [9] also showed that with $O(n^2)$ samples it is possible to get arbitrarily close to 0.745, which is the best one can achieve with full knowledge of the underlying distribution. Subsequent work by Rubinfeld et al. [27] improved this result by reducing the number samples that are required for this to $O(n)$. Tight bounds are known also for the case of a *known* distribution, equal to around 0.58 for the secretary problem [15] and around 0.745 for the prophet problem [19, 21, 1, 11]. However, all these algorithms are based on quantiles or empirical quantiles of the distribution, and strong communication-complexity lower bounds for quantile estimation [16] suggest that they cannot be implemented in the streaming model.

An interesting direction for future work would be to prove a formal separation between performance guarantees achievable in the streaming model, and performance guarantees achievable without this requirement.

Relevant Problems in the Streaming Literature. Algorithms for the secretary and prophet problems typically rely on aggregate statistics such as the median or mean of the distribution, or on more fine-grained information such as quantiles. A classic paper by West [29] gives an efficient algorithm for updating the mean and variance. The aforementioned paper by Guha and McGregor [16], on the other hand, gives lower bounds for the estimation of median and quantiles.

Our implementation of MRS algorithms shares certain characteristics with reservoir sampling and its variants [e.g., 28, 23], where the goal is to produce at any point in time a random subset of size k of the values seen so far. We require subsets of varying size at certain points in time, as well as a better space complexity.

Algorithms from Data. 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., 17, 2, 4, 5], in particular in the context of designing revenue-optimal auctions from samples [e.g., 7, 25].

2 Preliminaries

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

Online Selection Problems. We are given k samples S_1, \dots, S_k and n values X_1, \dots, X_n . Both samples and values are obtained as iid draws from an unknown distribution F . The samples are presented up front, the values are revealed one by one. Our goal is to select one of the values immediately and irrevocably when it is revealed.

Stopping Rules. Our objective can be formalized as the search for a *stopping rule* that for each $i \in [n]$ decides whether to select X_i based only on S_1, \dots, S_k and X_1, \dots, X_i . Such a stopping rule can be expressed as a family \mathbf{r} of functions r_1, \dots, r_n , where $r_i : \mathbb{R}_+^{k+i} \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Here, for any $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$, $r_i(s_1, \dots, s_k, x_1, \dots, x_i)$ is the probability of stopping at X_i when we have observed samples $S_1 = s_1, \dots, S_k = s_k$ and values $X_1 = x_1, \dots, X_i = x_i$ and have not stopped at X_1, \dots, X_{i-1} . The *stopping time* τ of such a stopping rule \mathbf{r} is the random variable with support $\{1, \dots, n\} \cup \{\infty\}$ such that for all $\mathbf{s} \in \mathbb{R}_+^k$ and $\mathbf{x} \in \mathbb{R}_+^n$,

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

Prophet and Secretary Objectives. We consider stopping rules for two different objectives. In the *prophet setting* we are interested in the expected value $\mathbb{E}[X_\tau]$ of the variable at which a stopping rule stops, where we use the convention that $X_\infty = 0$, and will measure its performance relative to the expected maximum $\mathbb{E}[\max\{X_1, \dots, X_n\}]$ of the random variables X_1, \dots, X_n . We will say that a stopping rule achieves approximation guarantee α , for $\alpha \leq 1$, if for any distribution $\mathbb{E}[X_\tau] \geq \alpha \mathbb{E}[\max\{X_1, \dots, X_n\}]$. In the *secretary setting* we are interested in maximizing the probability $\alpha = \Pr[X_\tau = \max_i X_i]$ with which the stopping rule stops on a maximum of the sequence of random variables.

Streaming Model. We are interested specifically in stopping rules that can be implemented as streaming algorithms. We assume that the stream consists of the samples S_1, \dots, S_k followed by the values X_1, \dots, X_n . An algorithm is 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 $\max\{S_1, \dots, S_k, X_1, \dots, X_n\}$.

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 Prophet Objective

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

3.1 Exact Analysis of MRS Algorithms

Theorem 1. *Consider a sequence of n random variables X_1, \dots, X_n drawn independently from an unknown distribution. Then, as $n \rightarrow \infty$, 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.*

3.1.1 Structural Lemma

Lemma 1. *Consider some MRS algorithm based on $f : [n] \rightarrow \mathbb{N}$ and $i \in [n]$. Conditioned on the fact that the algorithm arrives at step i , the distribution of the set $\{S_1, \dots, S_k, X_1, \dots, X_{i-1}\}$ 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 $i = 1, \dots, i^* - 1$. Then, conditioned on the fact that the algorithm arrives at step $i^* - 1$, the set $\mathcal{T} = \{S_1, \dots, S_k, X_1, \dots, X_{i^*-2}\}$ has the same distribution as the one of a set of $k + i^* - 2$ fresh samples, so the distribution of the set $\mathcal{T}' = \{S_1, \dots, S_k, X_1, \dots, X_{i^*-1}\}$ is the same as the one of a set of $k + i^* - 1$ fresh samples. We will argue that the decision of the algorithm to stop at X_{i^*-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 $S_1, \dots, S_k, X_1, \dots, X_{i^*-1}$ are distinct, so that each of these values can be identified with a unique rank in $[k + i^* - 1]$. By definition, the decision of an MRS algorithm to stop or continue only depends on the ranks of the values $R_{i^*-1}^1, \dots, R_{i^*-1}^{f(i^*-1)}, X_{i^*-1}$. Since the distribution of \mathcal{T} , from which $R_{i^*-1}^1, \dots, R_{i^*-1}^{f(i^*-1)}$ are drawn, is that of $k + i^* - 1$ fresh samples, and since X_{i^*-1} is a fresh sample, those ranks are $f(i^* - 1) + 1$ uniform draws without replacement from $[k + i^* - 1]$ and thus independent of the realization of \mathcal{T}' . \square

3.1.2 Formulation as a Control Problem

Fix $n \in \mathbb{N}$, and consider an MRS algorithm given by the function $f : [n] \rightarrow \mathbb{N}$. We can construct a continuous function $g : [0, 1] \rightarrow \mathbb{R}_+$ 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 the MRS algorithm A defined by g , a sequence X_1, \dots, 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

$$\begin{aligned} \mathbb{E}[X_\tau] &= \sum_{i=1}^n \Pr[A \text{ arrives at } X_i] \cdot \Pr[A \text{ accepts } X_i \mid A \text{ arrives at } X_i] \cdot \mathbb{E}[X_i \mid A \text{ accepts } X_i] \\ &= \sum_{i=1}^n \prod_{j=1}^{i-1} \left(1 - \frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right) \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \int_0^\infty (1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)) \, dx \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n \exp\left(\sum_{j=1}^{i-1} \ln\left(1 - \frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right)\right) \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \int_0^\infty (1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)) \, dx \\
&= \sum_{i=1}^n \exp\left(-\sum_{j=1}^{i-1} \left(\frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1} - O\left(\frac{1}{n^2}\right)\right)\right) \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \int_0^\infty (1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)) \, dx \\
&= e^{-O(\frac{1}{n})} \sum_{i=1}^n \exp\left(-\sum_{j=1}^{i-1} \frac{1}{\lceil g(\frac{j}{n}) \cdot n \rceil + 1}\right) \frac{1}{\lceil g(\frac{i}{n}) \cdot n \rceil + 1} \int_0^\infty (1 - F^{\lceil g(\frac{i}{n}) \cdot n \rceil + 1}(x)) \, dx,
\end{aligned}$$

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 $\sum_{i=2}^\infty -\frac{x^{-i}}{i} = -O(\frac{1}{x^2})$. Thus, for $n \rightarrow \infty$,

$$\begin{aligned}
\mathbb{E}[X_\tau] &= \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, dz\right) \cdot \frac{1}{g(y)} \cdot \int_0^\infty (1 - F^{g(y) \cdot n}(x)) \, dx \, dy \\
&= \int_0^\infty \int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} \, dz\right) \cdot \frac{1}{g(y)} \cdot (1 - F^{g(y) \cdot n}(x)) \, dy \, dx,
\end{aligned}$$

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 $\mathbb{E}[X_\tau] \geq \alpha \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}]$ or, equivalently,

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

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

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

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

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

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

3.1.3 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.

Upper Bound. By the max-min inequality,

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

Now the inner problem can be written as

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

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

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

where

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

and

$$\begin{aligned} \frac{d}{dy} \frac{\partial}{\partial h'} L(y, h(y), h'(y)) &= -e^{-h(y)} \cdot h'(y) \cdot \frac{1 - a^{\frac{1}{h'(y)}}}{1 - a} + e^{-h(y)} \cdot \frac{\ln a \cdot a^{\frac{1}{h'(y)}} \cdot h''(y)}{(1 - a) \cdot (h'(y))^2} \\ &+ e^{-h(y)} \cdot (h''(y) - (h'(y))^2) \cdot \frac{\ln a \cdot a^{\frac{1}{h'(y)}}}{(1 - a) \cdot (h'(y))^2} \\ &+ e^{-h(y)} \cdot h'(y) \cdot \left(-\frac{(\ln(a))^2 a^{\frac{1}{h'(y)}} h''(y)}{(1 - a)(h'(y))^4} - \frac{2 \ln(a) a^{\frac{1}{h'(y)}} h''(y)}{(1 - a)(h'(y))^3} \right). \end{aligned} \quad (7)$$

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

$$\begin{aligned} h_1(y) &= \sqrt{\kappa - \mu y} - \sqrt{\kappa}, & h'_2(y) &= -\frac{\mu}{2\sqrt{\kappa - \mu y}}, & \mu &= -2 \ln(a) \geq 0, & \kappa &\geq \mu, & \text{and} \\ h_2(y) &= \sqrt{\kappa} - \sqrt{\kappa - \mu y}, & h'_2(y) &= \frac{\mu}{2\sqrt{\kappa - \mu y}}, & \mu &= -2 \ln(a) \geq 0, & \kappa &\geq \mu, \end{aligned}$$

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

$$\begin{aligned} \mathcal{P} &\geq \sup_{\mu \in \mathbb{R}_+} \inf_{a \in [0,1]} \left\{ \frac{e^{-\sqrt{\mu}}}{1-a} \cdot \int_0^1 \frac{e^{\sqrt{\mu \cdot (1-y)}} \cdot \mu \cdot \left(1 - a^{\frac{2}{\mu}} \sqrt{\mu \cdot (1-y)}\right)}{2 \cdot \sqrt{\mu \cdot (1-y)}} dy \right\} \\ &= \sup_{\mu \in \mathbb{R}_+} \inf_{a \in [0,1]} \left\{ \frac{e^{-\sqrt{\mu}}}{1-a} \cdot \left[\frac{e^{\sqrt{\mu \cdot (1-y)} \cdot \left(1 + \frac{2 \ln a}{\mu}\right)}}{1 + \frac{2 \ln a}{\mu}} - e^{\sqrt{\mu \cdot (1-y)}} \right]_0^1 \right\} \\ &= \sup_{\mu \in \mathbb{R}_+} \inf_{a \in [0,1]} \left\{ \frac{e^{-\sqrt{\mu}}}{1-a} \cdot \left(\frac{1}{1 + \frac{2 \ln a}{\mu}} - 1 - \frac{e^{\sqrt{\mu} \cdot \left(1 + \frac{2 \ln a}{\mu}\right)}}{1 + \frac{2 \ln a}{\mu}} + e^{\sqrt{\mu}} \right) \right\} \\ &= \sup_{\mu \in \mathbb{R}_+} \inf_{\substack{b \in \mathbb{R}_+; \\ b \notin \{0,1\}}} \left\{ \frac{e^{-\sqrt{\mu}}}{1 - e^{-\frac{\mu b}{2}}} \cdot \left(\frac{1}{1-b} - 1 - \frac{e^{\sqrt{\mu} \cdot (1-b)}}{1-b} + e^{\sqrt{\mu}} \right) \right\}, \end{aligned}$$

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 $b \rightarrow 1$, 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 $g(y) = 1/h'(y) = 2\sqrt{\bar{\mu} - \bar{\mu}y}/\bar{\mu}$, 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.

3.2 MRS Algorithms with Constraints on the Number of Samples

Theorem 2. Consider a sequence of n random variables X_1, \dots, X_n drawn independently from an unknown distribution. A lower bound on the performance of MRS algorithms follows by considering MRS algorithms that use $g(y) = (\beta + t)\sqrt{\frac{y-1}{t-1}}$ for some parameter $t \in [0, 1]$ and minimizing (2) over $a \in [0, 1]$. For $\beta = n$ this yields an approximation ratio of $\alpha \geq 0.6489$ as $n \rightarrow \infty$.

Consider an MRS algorithm that has access to βn samples for some $\beta < 2/\sqrt{\bar{\mu}} \approx 1.4434$. 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 3.1.3, we can write the performance of the best MRS algorithm from the

restricted class as a control problem

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

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 3.1.3, 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 3.1.3 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

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

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 3.1.3 we may derive a lower bound on the value of \mathcal{Q} by considering the parametric class of functions

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

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 $a \rightarrow e^{2(t-1)/(\beta+1)^2}$. 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.

β	$\alpha \geq$	$t \approx$	$a \approx$	β	$\alpha \geq$	$t \approx$	$a \approx$
1.4	0.653368	0.02550	0.38325	\vdots	\vdots	\vdots	\vdots
1.3	0.653280	0.08754	0.38756	0.7	0.616281	0.72082	0.75831
1.2	0.652853	0.15519	0.39845	0.6	0.588379	0.94979	0.95905
1.1	0.651654	0.23067	0.41942	0.5	0.549306	0.99997	1.00000
1.0	0.648957	0.31759	0.45559	0.4	0.501105	1.00000	1.00000
0.9	0.643563	0.42162	0.51560	0.3	0.439901	1.00000	1.00000
0.8	0.633580	0.55160	0.61203	0.2	0.358351	1.00000	1.00000
\vdots	\vdots	\vdots	\vdots	0.1	0.239789	1.00000	1.00000

Table 1: Lower bounds on the performance (w.r.t. prophet objective) of the optimal MRS algorithm with access to $\beta \cdot n$ 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 alongside it.

3.3 An Optimal Algorithm for Up to Around $0.58 \cdot n$ Samples

Theorem 3. *Let $\beta \leq \frac{1}{e-1}$. Consider a sequence of n random variables X_1, \dots, X_n drawn independently from an unknown distribution, and assume that the decision-maker has access to $\lfloor \beta n \rfloor$ samples beforehand. Consider the following algorithm: Until stage $\lfloor (\frac{1+\beta}{e} - \beta)n \rfloor$, do not pick anything. From stage $\lfloor (\frac{1+\beta}{e} - \beta)n \rfloor + 1$ onward, pick anything that is higher than the maximum of the observed variables (samples and proposed variables). Then, as $n \rightarrow +\infty$, this algorithm achieves an approximation ratio of $\alpha = \frac{1+\beta}{e}$, which matches the upper bound in [9].*

Proof. To simplify the exposition, we will pretend that $\frac{1+\beta}{e} - \beta$ and βn are integers, and thus drop the symbols \lfloor and \rfloor . Let $\delta := \frac{1+\beta}{e} - \beta$. Let $(S_1, \dots, S_{\beta n})$ be the set of sample variables, and (X_1, \dots, X_n) be the set of proposed variables, and F their cumulative distribution. Let τ be the stopping time of the algorithm, and $x \in \mathbb{R}_+$. We are going to prove that for large n and uniformly in F ,

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

Then, integrating over x yields the theorem.

We have

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

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

$$\begin{aligned} & \Pr(\{X_i \geq \max \{x, S_1, \dots, S_{\beta n}, X_1, \dots, X_{i-1}\}\} \\ & \quad \cap \{\max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_{\delta n}\} \geq \max \{X_{\delta n+1}, \dots, X_{i-1}\}\}) \\ &= \Pr(X_i \geq \max \{x, S_1, \dots, S_{\beta n}, X_1, \dots, X_{i-1}\}) \\ & \quad \cdot \Pr(\max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_{\delta n}\} \geq \max \{X_{\delta n+1}, \dots, X_{i-1}\}) \end{aligned}$$

We have

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

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

$$\begin{aligned} & \Pr(X_i \geq x | X_i \geq \max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_{i-1}\}) \cdot \Pr(X_i \geq \max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_{i-1}\}) \\ &= \Pr(\max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_i\} \geq x) \cdot \Pr(X_i = \max \{S_1, \dots, S_{\beta n}, X_1, \dots, X_i\}) \end{aligned}$$

$$= \frac{1 - F^{\beta n + i}(x)}{\beta n + i}.$$

Thus,

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

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

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

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. \square

3.4 Implementation as Streaming Algorithms

We next show how to implement MRS algorithms as streaming algorithms. To this end, consider $x_0 \in [0, 1]$, $\bar{y} \in \mathbb{R}_+$, and any continuous decreasing function $g : [0, 1] \rightarrow [0, \bar{y}]$. 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 $\gamma \leq \lceil 1/\epsilon \rceil + 2 \cdot \lceil \bar{y}/\epsilon \rceil + 2$ intersection points with g . Let $0 = x_1 \leq x_2 \leq \dots \leq x_\gamma = 1$ be the corresponding x -coordinates of these intersection points.

For all $i = 1, \dots, \gamma - 1$, 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 $\tilde{g} : [0, 1] \rightarrow [0, \bar{y} + \epsilon]$ 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 2. *Let $j \in [x_i \cdot n, x_{i+1} \cdot n) \cap \mathbb{Z}$ for some $i \in \{1, \dots, \gamma - 1\}$. 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 $\{S_1, \dots, S_{\beta n}, X_1, \dots, X_{j_0-1}\}$ by $Y_S := \{S'_1, \dots, S'_\ell\}$ 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 Lemma 1, this set and the set $Y := \{S'_1, \dots, S'_\ell, X_{j_0}, \dots, X_{j-1}\}$ 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, \dots, 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). \square

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

Lemma 3. *Suppose that the MRS algorithm defined by g achieves approximation ratio α via inequality (2) and that $|\tilde{g}(x) - g(x)| \leq 2\epsilon$ for all $x \in [0, 1]$, then the MRS algorithm defined by \tilde{g} achieves approximation ratio $\alpha - O(\sqrt{\epsilon})$.*

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

$$\int_0^1 \exp\left(-\int_0^y \frac{1}{g(z)} dz\right) \cdot \frac{1}{g(y)} \cdot (1 - a^{g(y)}) dy \geq \alpha \cdot (1 - a). \quad (10)$$

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

$$(1 - a)^{-1} \int_{1-\epsilon}^1 \exp\left(-\int_0^y \frac{1}{g(z)} dz\right) \cdot \frac{1}{g(y)} \cdot (1 - a^{g(y)}) dy = O(\sqrt{\epsilon}). \quad (11)$$

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

$$(1 - a)^{-1} \int_0^{1-\epsilon} \exp\left(-\int_0^y \frac{1}{\tilde{g}(z)} dz\right) \cdot \frac{1}{\tilde{g}(y)} \cdot (1 - a^{\tilde{g}(y)}) dy \geq \alpha - O(\sqrt{\epsilon}).$$

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 concludes the proof. \square

To implement our approach as a streaming algorithm, for each $i = 0, \dots, \gamma$, 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.




Figure 2: Visualization of the g function from Section 3.1 (bold, black), the grid with $\epsilon = 0.2$ (dotted, black), and the resulting \tilde{g} function (green, bold).

Lemma 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 $\{S_1, \dots, S_k, X_1, \dots, X_{j-1}\}$. 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 is $1/\binom{k+j-1}{q}$. Indeed, consider an arbitrary such vector z . Let $E = \{e_1, e_2, \dots, e_q\} \subseteq [k + j - 1]$ with $e_1 < e_2 < \dots < e_q$ be the indices ℓ where $z_\ell = 1$ and let $N = \{n_1, \dots, n_{k+j-1-q}\}$ with $n_1 < \dots < n_{k+j-1-q}$ 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. \square

4 Secretary Objective

To get a streaming algorithm for the secretary objective with n samples, we consider the sliding-window approach [10] that achieves a guarantee of $\ln(2) - \ln^2(2)/2 \approx 0.453$ but uses a linear amount of space: This algorithm accepts X_j iff it exceeds $\max\{S_j, \dots, S_n, X_1, \dots, X_{j-1}\}$. To implement this algorithm in the streaming model with ϵ loss in the guarantee, we remember the $O_\epsilon(1)$ largest order statistics of all values and their positions. We then bound the loss incurred from situations where we cannot compute the sliding-window algorithm's threshold from memory by ϵ .

Theorem 4. *For $\beta = 1$ and any $\epsilon > 0$, there exists an algorithm with guarantee $\ln(2) - \ln^2(2)/2 - \epsilon$ using a single pass over the data and $O(\log n)$ space.*

Proof. Assume $\epsilon < 1$. We describe our streaming algorithm. As we pass through the stream $S_1, \dots, S_n, X_1, \dots, X_n$, we remember the $\ell \in O_\epsilon(1)$ (with ℓ yet to be specified) largest order statistics of the values seen up until then as well as their positions in the stream. This can be trivially done in space $O_\epsilon(\log n)$.

To decide whether or not to accept X_j for $j \in \{1, \dots, n\}$, we compute the maximum of all remembered values within $\{S_j, \dots, S_n, X_1, \dots, X_{j-1}\}$. Note that there may be no such value; then we use the convention that $\max \emptyset = -\infty$. We accept X_j if and only if it exceeds the threshold (so we definitely accept X_j if the threshold is $-\infty$).

We compare the behavior of this algorithm with that of the algorithm in [10]. To do so, let $j^* := \lfloor (1 - \epsilon/2) \cdot n \rfloor$. Note that, if

- (i) one of ℓ largest order statistics of $S_1, \dots, S_n, X_1, \dots, X_n$ is among S_{j^*}, \dots, S_n ,
- (ii) and $j \leq j^*$,

our computed threshold is equal to $\max\{S_j, \dots, S_n, X_1, \dots, X_{j-1}\}$, that is, the threshold that the algorithm in [10] sets. The reason is that this value is either the largest order statistic of $S_1, \dots, S_n, X_1, \dots, X_{j-1}$ that is in S_{j^*}, \dots, S_n , which we remembered by (i), or it is an even larger order statistic of $S_1, \dots, S_n, X_1, \dots, X_{j-1}$, which we also remembered.

Using the union bound, we bound the probability that the algorithm in [10] accepts $\max\{X_1, \dots, X_n\}$ but our streaming algorithm does not by ϵ (for an appropriate choice of ℓ). Clearly, $\max\{X_1, \dots, X_n\} = X_j$ for $j \geq j^*$ happens with probability $2\epsilon/3$ for n large enough since

$$\lim_{n \rightarrow \infty} \frac{(1 - \epsilon/2) \cdot n - j^*}{n} = 0;$$

for smaller n we can afford to behave identically as the algorithm in [10] by remembering the entire stream. So it is enough to choose ℓ such that the event in (i) happens with probability $1 - \epsilon/3$. We note that

$$\begin{aligned} \Pr[\text{event in (i) does not occur}] &\leq (1 - \epsilon)^\ell \\ &\leq \exp(-\epsilon\ell), \end{aligned}$$

so choosing $\ell = 1/\epsilon \cdot \ln(3/\epsilon)$ completes the proof. □

References

- [1] 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.
- [2] 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.
- [3] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1):137–147, 1999.
- [4] 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, 2018.
- [5] 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, 2018.
- [6] B. A. Berezovskiy and A. V. Gnedin. *The Best Choice Problem (In Russian)*. Nauka, Moscow, 1984.
- [7] 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.
- [8] D. Conlon, J. Fox, and B. Sudakov. Hypergraph Ramsey numbers. *Journal of the American Mathematical Society*, 23(1):247–266, 2010.

- [9] 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, 2019.
- [10] 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.
- [11] J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 169–186, 2017.
- [12] E. Dynkin. The optimum choice of the instant for stopping a Markov process. *Soviet Mathematics, Doklady*, 4:627–629, 1963.
- [13] T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- [14] M. Gaber, A. Zaslavsky, and S. Krishnaswamy. Mining data streams: A review. *SIGMOD Record*, 34(2): 18–26, 2005.
- [15] J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.
- [16] 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.
- [17] R. Gupta and T. Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.
- [18] T. Hill and R. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics*, 125:191–207, 1992.
- [19] 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.
- [20] 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.
- [21] R. P. Kertz. Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis*, 19:88–112, 1986.
- [22] H. W. Kuhn. Extensive games and the problem of information. In *Annals of Mathematics Studies*, volume 28, pages 193–216. Princeton University Press, 1953.
- [23] K.-H. Li. Reservoir-sampling algorithms of time complexity $o(n(1+\log(n/n)))$. *ACM Transactions on Mathematical Software*, 20:481–493, 1994.
- [24] J.-F. Mertens, S. Sorin, and S. Zamir. *Repeated Games*. Cambridge University Press, 2015.
- [25] 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.
- [26] F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 30:264–286, 1930.
- [27] 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.
- [28] J. S. Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 11(1): 37–57, 1985.

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

A Tight Prophet Inequality for Exchangeable Random Variables

In this appendix we consider sequences X_1, \dots, X_n of *exchangeable* random variables from a known joint distribution D . A sequence of random variables is exchangeable if its joint distribution is invariant under permutations of the sequence.

Examples of exchangeable random variables with a known joint distribution can be found in the context of urn models: Drawing n times *with replacement* from an urn with any number of m known values in it leads to a sequence X_1, \dots, X_n of exchangeable random variables. In this case the random variables will also be independently and identically distributed. If instead we draw n times *without replacement* from an urn with $m \geq n$ known values in it, then the resulting sequence of random variables X_1, \dots, X_n will still be exchangeable, but it is no longer independently and identically distributed.

Additional examples include any sequence of i.i.d. random variables X_1, \dots, X_n that are drawn from a known distribution or, more generally, any sequence of random variables X_1, \dots, X_n that is generated by first choosing a distribution F^k from a set of known probability distributions $\{F^1, \dots, F^m\}$ according to a known probability distribution G on the index set $\{1, \dots, m\}$ and then drawing n times from F^k .

We consider stopping rules that sequentially observe the random variables X_1, \dots, X_n , and for each $i \in [n]$ decide whether to stop at X_i based on the values of X_1, \dots, X_i and the known joint distribution D . Such a stopping rule can be expressed as a family \mathbf{r} of functions r_1, \dots, r_n , where $r_i : \mathbb{R}_+^i \times \mathcal{D} \rightarrow [0, 1]$ for all $i = 1, \dots, n$. Here, for any $\mathbf{x} \in \mathbb{R}_+^n$ and $D \in \mathcal{D}$, $r_i(x_1, \dots, x_i, D)$ is the probability of stopping at X_i when we have observed the values $X_1 = x_1, \dots, X_i = x_i$, have not stopped at X_1, \dots, X_{i-1} , and when the joint distribution is D . The *stopping time* τ of such a stopping rule \mathbf{r} is thus the random variable with support $\{1, \dots, n\} \cup \{\infty\}$ such that for all $\mathbf{x} \in \mathbb{R}_+^n$ and $D \in \mathcal{D}$,

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

Theorem 5.

- (a) *There exists a stopping rule with stopping time τ such that for any sequence of exchangeable random variables X_1, \dots, X_n , $\mathbb{E}[X_\tau] \geq \frac{1}{e} \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}]$.*
- (b) *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, \dots, X_n such that for any stopping rule with stopping time τ , $\mathbb{E}[X_\tau] \leq (\frac{1}{e} + \delta) \cdot \mathbb{E}[\max\{X_1, \dots, X_n\}]$.*

Part (a) follows by applying the optimal solution to the classic secretary problem. To obtain Part (b) we cast the problem as a two-player zero-sum game, or equivalently as a min-max problem, where the first player chooses an exchangeable sequence of random variables 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 the distribution of the random variables, and player 1, the minimizer, gets to choose a difficult distribution with knowledge of the stopping rule. This max-min problem is, in fact, more difficult than the prophet inequality problem for i.i.d. random variables *from an unknown distribution* recently considered by Correa et al. [9]. The construction of Correa et al., however, relies on the infinite version of Ramsey’s theorem [26], 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., 24, 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. [9] 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. [8].

An interesting aspect of our argument is that the minimax theorem used in the above argument of course requires mixed strategies. This is clearly not a problem for player 2, but for player 1 this means 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 [22].

A.1 Hard Finite Instances for Unknown I.I.D. Random Variables

We first consider the setting of Correa et al. [9], where random variables are drawn independently from the same unknown distribution and we do not have access to any additional samples. For $p \in \mathbb{N}$, denote by B_p the set of probability measures on $[p]$. We prove the following result.

Proposition 1. *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, \dots, X_n are i.i.d. random variables drawn from b ,*

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

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 minmax theorem in Section A.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 called (ε, i) -value-oblivious on V if, there exists a $q_i \in [0, 1]$ such that, for all pairwise distinct $v_1, \dots, v_i \in V$ with $v_i > \max\{v_1, \dots, v_{i-1}\}$, it holds that $r_i(v_1, \dots, v_i) \in [q_i - \varepsilon, q_i + \varepsilon]$.*
- *A stopping rule \mathbf{r} is called 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 $v_1, \dots, v_j \in \mathbb{R}_+$ and all permutations π of $[j - 1]$, $r_i(v_1, \dots, v_j) = r_i(v_{\pi(1)}, \dots, v_{\pi(j-1)}, v_j)$.*

The cornerstone of our proof is the following lemma.

Lemma 5. *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 (see for example Conlon et al. [8]).

Lemma 6. *There exists a function $R : \mathbb{N}^3 \rightarrow \mathbb{N}$ 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 Lemma 5. Fix $\varepsilon > 0$ and set $c = \lfloor (2\varepsilon)^{-1} \rfloor$. Define an integer sequence $(p_i)_{0 \leq i \leq n}$ 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 $j \in \{0, 1, \dots, n\}$ 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 $S_0 = \{n^{3s} : s = 0, \dots, p_0\}$ 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 Lemma 6, we construct a complete j -hypergraph H with vertex set S^{j-1} . Consider any set $v_1, \dots, v_j \subset S^{j-1}$ of cardinality j such that $v_j > \max(v_1, \dots, v_{j-1})$. Note that there exists a unique $u \in \{1, 2, \dots, c\}$ such that $r_j(v_1, \dots, v_j) \in [(2u-1)\varepsilon - \varepsilon, (2u-1)\varepsilon + \varepsilon)$, and color the hyperedge $\{v_1, \dots, v_j\}$ of H with color u . By Lemma 6, 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 $v_1, \dots, v_j \in S^j$ with $v_j > \max(v_1, \dots, v_{j-1})$. Since the edge $\{v_1, \dots, v_j\}$ in H has color u , $r_j(v_{\pi(1)}, \dots, v_{\pi(j-1)}, v_j) \in [q - \varepsilon, q + \varepsilon)$ for some permutation π of S^{j-1} . But since \mathbf{r} is order-oblivious, also $r_j(v_1, \dots, v_{j-1}, v_j) \in [q - \varepsilon, q + \varepsilon)$. So \mathbf{r} is (ε, j) -value oblivious on S^j . This completes the induction step. \square

With Lemma 5 at hand we are now ready to prove Proposition 1.

Proof of Proposition 1. Let $\delta > 0$ and $n \in \mathbb{N}$. Consider a stopping rule \mathbf{r} with performance guarantee $1/e + \delta$. Set $\varepsilon = 1/n^2$. By Lemma 5, 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 $V = \{v_1, \dots, v_{n^3}, u\}$. Denote by τ the stopping time of \mathbf{r} . By construction, we have $u \geq n^3 \max\{v_1, \dots, v_{n^3}\}$. 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

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

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, \dots, 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 [13]. 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. \square

A.2 From Unknown I.I.D. to Exchangeable Variables: A Minimax Argument

We now use a minimax argument to convert Proposition 1, which concerns the unknown i.i.d. case, into an upper bound (impossibility result) for exchangeable random variables *with a known joint distribution*. Fix $\delta > 0$ and let n_0 and p be as in Proposition 1. Fix $n \geq n_0$. We need the following definitions.

Definition 2.

- A *deterministic stopping rule* is a sequence $a = (a_1, \dots, a_n)$ such that $a_i : [p]^{i-1} \rightarrow \{0, 1\}$. 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 $\mathbf{r} = (r_1, \dots, r_n)$ such that $r_i : [p]^{i-1} \rightarrow [0, 1]$. 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 [22] 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 exchangeable problem with some fixed known joint distribution \mathcal{D} , where in the latter case we have omitted the dependence of a on \mathcal{D} .

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

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

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

$$g(x, y) = \int_A \int_B g(a, b) x(da) y(db).$$

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 γ in the exchangeable problem.

Let us first examine V^- . Let $a \in A$ and $y \in \mathcal{P}(B)$, and consider the n -tuple of exchangeable random variables Y_1, Y_2, \dots, Y_n obtained by first picking an element $b \in B$ according to y and then drawing Y_1, Y_2, \dots, Y_n independently from b . Let $\tilde{\tau}$ be the stopping time of the stopping rule a on Y_1, Y_2, \dots, Y_n . By linearity of expectation,

$$g(a, y) = \mathbb{E}[X_{\tilde{\tau}}] - \left(\frac{1}{e} + \delta\right) \cdot \mathbb{E}[\max\{Y_1, \dots, Y_n\}].$$

This implies the following result.

Proposition 2. *If V^- is non-positive, then $\gamma \leq \frac{1}{e} + \delta$.*

Proof. Assume that V^- is non-positive. Then there exists $y \in \mathcal{P}(B)$ such that no deterministic stopping rule provides a better guarantee than $(1/e + \delta)$ in the exchangeable problem with 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 [22]. It follows that $\gamma \leq 1/e + \delta$. \square

Moreover, the minimax theorem implies that V^+ and V^- are the same.

Proposition 3. $V^+ = V^-$.

Proof. The set A is finite, the set B compact metric. Moreover, for all $a \in A$, the mapping $b \rightarrow g(a, b)$ is continuous. By the minimax theorem [24, Proposition 1.17] it follows that the mixed extension of the normal-form game (A, B, g) has a value, i.e., that $V^+ = V^-$. \square

To complete the proof of Theorem 5 it is thus enough to show that $V_+ \leq 0$. First note that by Kuhn's theorem [22],

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

By definition of p , which we have chosen as in Proposition 1, for each stopping rule in C there exists $b \in B$ such that

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

It follows that $V_+ \leq 0$, as claimed.