The Last Success Problem with Samples

Toru Yoshinaga
The University of Tokyo, Japan

Yasushi Kawase
The University of Tokyo, Japan

Abstract

The last success problem is an optimal stopping problem that aims to maximize the probability of stopping on the last success in a sequence of independent n Bernoulli trials. In the classical setting where complete information about the distributions is available, Bruss [4] provided an optimal stopping policy that ensures a winning probability of $1/e$. However, assuming complete knowledge of the distributions is unrealistic in many practical applications. This paper investigates a variant of the last success problem where samples from each distribution are available instead of complete knowledge of them. When a single sample from each distribution is allowed, we provide a deterministic policy that guarantees a winning probability of $1/4$. This is best possible by the upper bound provided by Nuti and Vondrák [33]. Furthermore, for any positive constant ϵ, we show that a constant number of samples from each distribution is sufficient to guarantee a winning probability of $1/e - \epsilon$.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases The Last Success Problem, Secretary Problem, Sample Information Model, Optimal Stopping, Online Algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.105

Related Version *Full Version*: `https://arxiv.org/abs/2308.09356` [37]

Funding This work was supported by JST SPRING Grant Number JPMJSP2108, JST ERATO Grant Number JPMJER2301, JST PRESTO Grant Number JPMJPR2122, JSPS KAKENHI Grant Number JP20K19739, and Value Exchange Engineering, a joint research project between Mercari, Inc. and the RIISE.

Acknowledgements The authors would like to thank the anonymous reviewers for their valuable and insightful comments.

1 Introduction

Imagine you are driving down a street toward a movie theater, hoping to park in a space near the destination. You cannot know in advance which parking spaces are free. Each time you encounter an available space, you must decide whether to park there or continue driving. Your goal is to maximize the probability of parking in the available space nearest to the theater without going back. This scenario is studied as *the last success problem* in the literature on optimal stopping theory [15, 35].

The last success problem is an optimal stopping problem that aims to maximize the probability of stopping on the last *success* in a sequence of independent n Bernoulli trials, where the success probability of ith trial is known to be p_i for each $i \in \{1, 2, \dots, n\}$. The event of stopping on the last success is referred to as a *win*, whereas the other event is referred to as a *loss*.

This problem is a generalization of various optimal stopping problems, including the classical secretary problem, and has a wide range of applications such as parking [35], maintenance planning for production equipment [27, 31], and ethical choices for clinical trials [6]. The last success problem was originally studied by Hill and Krengel [25]. Subsequently,

32nd Annual European Symposium on Algorithms (ESA 2024).
Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 105; pp. 105:1–105:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Bruss [4] provided an optimal stopping policy for this problem. Bruss's policy is to stop at the first success from the index $\tau = \min\{t \in \{1, 2, \dots, n\} \mid \sum_{i=t+1}^{n} r_i < 1\}$, where r_i denotes the *odds* defined as $p_i/(1-p_i)$. Here, we treat r_i as ∞ if $p_i = 1$. In words, this policy stops at the first observed success within the range where the sum of the odds of future successes is less than one. Bruss's policy guarantees a winning probability of at least $1/e$ if the sum of the odds is at least one, i.e., $\sum_{i=1}^{n} r_i \geq 1$, and this lower bound is shown to be optimal [5]. Note that an assumption about the sum of the odds is needed to guarantee a winning probability. Without such an assumption, there can be a situation where no successes occur (i.e., $p_1 = \cdots = p_n = 0$), leading to an inevitable loss of any policy.

However, the assumption that the success probabilities are completely known is too restrictive in certain cases. For example, in the parking problem, it seems unnatural to know the probability of each parking place being available. Nevertheless, if you have previously driven to the movie theater, you would be aware of whether each parking space was available at that time. To model such a situation, it is appropriate to assume unknown distributions along with finite data obtained from them. To address real-world applications, optimal stopping problems under limited information have recently been extensively studied [1, 2, 3, 10, 13, 12, 18, 19, 21, 23, 28, 29, 32].

In this paper, we investigate the last success problem in the sample model where the decision maker has only m samples for each (unknown) distribution. Very little is known about the problem of this variant, excluding the fact that no policy achieves a winning probability strictly better than $1/4$ in the single sample case where $m = 1$.

▶ **Theorem 1** (Nuti and Vondrák [33]). *For any positive real ϵ, there does not exist a (randomized) stopping policy for the single sample last success problem that guarantees a winning probability of $1/4 + \epsilon$ even when a success occurs with probability one.*

Nuti and Vondrák [33] provided this impossibility result while examining a setting named the *adversarial order single sample secretary problem*, which is an optimal stopping problem with the goal of maximizing the probability of stopping on the largest value in a sequence of n independent real-valued random variables. Observing a tentatively maximum value in this problem can be interpreted as observing a success in our problem. They also proved that the following simple policy achieves a winning probability of $1/4$: set the largest value from the samples as a threshold and stop at the first index whose value beats this threshold. However, this fact does not directly imply the existence of a stopping policy that attains a winning probability of $1/4$ for the last success problem. Indeed, for the last success problem, the winning probability for any algorithm is 0 if the success probabilities are 0. Note that, in the adversarial order single sample secretary problem, such a case does not exist because the first variable must be a tentatively maximum value.

Our Contributions

We first examine the no-sample model, where $m = 0$, and show that no policy can guarantee a positive winning probability even when a success occurs with probability one (Theorem 5).

For the single sample model, where $m = 1$, we propose two natural deterministic policies, which we call the *from the last success (FLS)* policy and the *after the second last success (ASLS)* policy, and evaluate their winning probabilities. Each of the two policies first determines a threshold index from the sample sequence and then stops at the first success from the index. The FLS policy and the ASLS policy select the index of the last success and the *next* index of the second-last success in the sample sequence as the threshold index, respectively. We show that the FLS policy cannot guarantee a winning probability of better than $(1-e^{-4})/4 \approx 0.2454$ even in instances where a success occurs with probability one (Theorem 6). On the other hand, we prove that the ASLS policy guarantees a winning probability of $1/4$ for every instance where the sum of the odds is at least $(\sqrt{3}-1)/2 \approx 0.3660$ (Theorem 7). We also demonstrate that the ASLS policy is nearly optimal for any case with a restriction on the sum of the odds.

Moreover, we analyze a *randomized* policy derived from a policy proposed by Nuti and Vondrák [33] for the adversarial order single sample secretary problem. This policy randomly selects the threshold index to be either the index of the last success or the next index of the last success in the sample sequence, each with a probability of $1/2$. We conducted this analysis for comparison with our policies. Our analysis reveals that this policy guarantees a winning probability of $1/4$ if the sum of the odds is at least $1/2$ (Theorem 8) for our problem setting. However, it fails to guarantee a winning probability of $1/4$ if the sum of the odds is slightly less than $1/2$. Thus, we conclude that the ASLS policy is superior in both performance and deterministic nature.

For the multiple sample model, we provide a policy that guarantees a winning probability of $1/e - \epsilon$ with a constant m for any positive constant ϵ if the sum of the odds is at least 1 (Theorem 9). Notably, this result does not depend on the number of trials n. A natural policy would be to estimate the probability of success for each distribution and apply Bruss's policy. However, such a method may lead to errors that depend on the number of trials n. Instead, our policy determines a threshold index from the samples by estimating an index i such that $\prod_{k=i}^{n}(1-p_k) \leq 1/e + \delta$ and $\prod_{k=i+1}^{n}(1-p_k) \geq 1/e$ for a small $\delta > 0$. To evaluate the winning probability of this policy, we first demonstrate that its winning probability is at least $1/e - \delta$ if it correctly identifies an index i that satisfies these conditions. Subsequently, we show that an index i meeting the conditions can be selected with high probability by utilizing a martingale property. Moreover, we prove that no policy can guarantee a winning probability of exactly $1/e$ if m is a finite number (Proposition 4). Due to space limitations, most of the proofs are omitted and can be found in a full version [37].

1.1 Related work

One of the most fundamental problems in optimal stopping is the *secretary problem* (see the survey by Ferguson [22] for a detailed history). In the classical setting, applicants are interviewed one by one in a random order. After each interview, the interviewer must make an immediate and irrevocable decision to either hire or reject the candidate. The interviewer can only rank the candidates among those interviewed up to that point. The goal is to maximize the probability of hiring only the best candidate. The secretary problem can be viewed as a special case of the last success problem of known distributions where $p_i = 1/i$ for $i = 1, 2, \dots, n$. This is because the occurrence of a best candidate so far in the secretary problem is analogous to the occurrence of a success in the last success problem.

Another fundamental problem is the *prophet inequality problem*. In this problem, a decision-maker observes a sequence of individual real-valued random variables, one by one. For each observation, an irrevocable decision must be made to either select the current variable or wait for the next one. The objective is to select the variable with the best value. It is known that a 1/2-competitive algorithm exists for this problem, and it is best possible [30].

Models that allow limited sample access to unknown distributions have been proposed for online optimization in recent years. Azar et al. [3] introduced a sample model in which inputs are drawn from unknown distributions, but the decision maker can access some samples from each distribution. They applied this model to the prophet inequalities under constraints, guaranteeing a constant competitive ratio for each setting. Rubinstein et al. [34] proposed a stopping policy with a competitive ratio of $1/2$ for the single sample prophet inequality. Remarkably, this competitive ratio is optimal even when the distributions are fully known in advance [30]. Building on these works, online optimization problems with a limited number of samples have been studied extensively for the past decade [8, 12, 9, 11, 14, 16, 20, 24, 28, 33].

Several studies have also been conducted on the last success problem in the unknown distribution setting. Bruss and Louchard [7] considered a variant of the last success problem where the distributions are unknown but identical (i.e., $p_1 = \cdots = p_n = p$). They investigated a stopping policy by estimating the success probability p from observed random variables. The special case of our problem with $m = 1$ can be reduced to the *single sample secretary problem in an adversarial order* by considering the no success scenario as a win (refer to the subsection 3.3 in details). Nuti and Vondrák [33] proposed a policy for the problem that ensures a winning probability of $1/4$ and showed that this is best possible by using Theorem 1.[1]

2 Preliminaries

2.1 Model

We formally define the last success problem with m samples. For a positive integer n, we denote the set $\{1, \dots, n\}$ by $[n]$. Suppose that there are n random variables $X_1, \dots, X_n$, following independent and non-identical Bernoulli distributions. A trial i is called *success* if $X_i = 1$. For each trial $i \in [n]$, the *success probability* $p_i = \Pr[X_i = 1]$ is unknown. Instead, m samples independently drawn from the ith Bernoulli distribution are available for each $i \in [n]$. We sequentially observe realizations of the variables $X_1, \dots, X_n$. Upon observing a success at trial i, an immediate and irrevocable decision must be made to either halt the observation or continue to the subsequent trial. This decision must be based only on the observed values $X_1, \dots, X_i$ and the samples. The result is a *win* if stopping the observation on the last success (i.e., $X_i = 1$ and $X_{i+1} = \cdots = X_n = 0$). Note that, if no successes are observed by the end of the sequence, the result must be a loss. Our goal is to design a stopping policy that maximizes the probability of a win, that is, the probability of stopping at the last success.

We will evaluate policy performance with a worst-case analysis. However, for instances where successes never occur (i.e., $p_1 = \cdots = p_n = 0$), the probability of a win is zero no matter what policies are used. To conduct a meaningful analysis, we use the *sum of the odds* $R = \sum_{i=1}^n r_i$ as a parameter in our analysis, where $r_i = p_i/(1 - p_i)$ is the odds of ith trial. This parameter R plays a crucial role in the context of the last success problem [4]. Let $\mathcal{I}_{m,R}$ be the set of instances of the last success problem with m samples such that the sum of the odds is at least R. For a stopping policy $\mathcal{P}$ and an instance $I \in \mathcal{I}_{m,R}$, let $\mathcal{P}(I)$ be the winning probability when we apply $\mathcal{P}$ to I. Then, we call $\inf_{I \in \mathcal{I}_{m,R}} \mathcal{P}(I)$ the winning probability of $\mathcal{P}$ for *the* (m, R)*-last success problem*. The (m, R)-last success problem is easier than the (m, R')-last success problem when $R > R'$ since $\mathcal{I}_R \subsetneq \mathcal{I}_{R'}$. Thus, the easiest case is the (m, ∞)-last success problem.

2.2 Basic observation

The sum of the odds R is associated with the probability of no success as follows.

[1] Nuti and Vondrák [33] provided this lower bound even for a slightly more general problem called the *adversarial order two-sided game of googol*.

▶ **Lemma 2.** *For* $p_1, \ldots, p_n \in [0, 1]$*, let* $R = \sum_{i=1}^{n} p_i/(1-p_i)$ *and* $Q = \prod_{i=1}^{n}(1-p_i)$*. Then, we have* $1/e^R \leq Q \leq 1/(1+R)$*.*

We can derive an upper bound of the winning probability for the last success problem with samples by referencing the upper bound in the complete information setting.

▶ **Proposition 3** (Bruss [4])**.** *The winning probability of any (randomized) stopping policy for the* (m, R)*-last success problem is at most* R/e^R *if* $R \leq 1$ *and at most* $1/e$ *if* $R \geq 1$*.*

For the prophet inequality problem [30], it has been demonstrated that the optimal performance of 1/2, which can be achieved with complete information, can be attained using just a single sample [34]. In contrast, for the last success problem, the optimal winning probability $1/e$, which can be achieved with complete information, cannot be attained using a finite number of samples.

▶ **Proposition 4.** *For any (randomized) stopping policy and any positive integer* m*, the winning probability for the* (m, ∞)*-last success problem is strictly less than* $1/e$*.*

Of course, this proposition also holds for the (m, R)-last success problem with any positive real R.

For the no-sample setting where $m = 0$, we can demonstrate that no (randomized) stopping policy can guarantee a positive constant probability of winning. Indeed, consider a scenario where the adversary selects an index j uniformly at random from $[n]$ and sets the success probabilities as $p_1 = \cdots = p_j = 1$ and $p_{j+1} = \cdots = p_n = 0$. In this scenario, we need to correctly guess the index j to win, which can only be done with a probability of $1/n$. Hence, by Yao's principle [36], any stopping policy can win with probability at most $1/n$ for an instance with n trials. Note that the sum of the odds of the instance is ∞ by $p_1 = 1$. By considering the limit as n approaches infinity, we obtain the following theorem.

▶ **Theorem 5.** *The winning probability of any (randomized) stopping policy is zero for the* $(0, \infty)$*-last success problem.*

3 Single sample model

In this section, we concentrate on the special case where $m = 1$ for the last success problem with m samples, which we refer to as the *single sample last success problem*. Throughout this section, we write $Y_i \in \{0, 1\}$ to denote the sample of ith trial for each $i \in [n]$.

3.1 Estimating the sum of the odds

When all success probabilities are known in advance, the optimal stopping policy for the last success problem is Bruss's policy [4], wherein one stops on the first success for which the sum of the odds for the future trials is at most 1. Given this, a natural approach to our single sample setting is to compute an estimated success probability $\hat{p}_i$ for each trial $i \in [n]$ from the sample Y_i and then apply Bruss's policy based on these estimated success probabilities. Let α_0 and α_1 be real numbers in the range $[0, 1]$, and consider estimating $\hat{p}_i$ as α_0 if $Y_i = 0$ and $\hat{p}_i$ as α_1 if $Y_i = 1$. We will then show that regardless of the values of α_0 and α_1 used for estimation, we cannot achieve a winning probability of 1/4 even when a success occurs with probability 1.

We examine three cases: (i) $\alpha_0 > 0$, (ii) $\alpha_1 < 1/2$, and (iii) $\alpha_0 = 0$ and $\alpha_1 \geq 1/2$.

For the first case, where $\alpha_0 > 0$, we consider an instance with $n = \lceil 1/a_1 \rceil + 1$ trials such that the success probabilities are $p_1 = 1$ and $p_2 = \cdots = p_n = 0$. Note that the last success is the first trial, which means that we must stop at the first trial to win.

Here, the estimated success probabilities must be $\hat{p}_1 = \alpha_1$ and $\hat{p}_2 = \cdots = \hat{p}_n = \alpha_0$. As $\sum_{i=2}^{n} \hat{p}_i/(1-\hat{p}_i) = \lceil 1/\alpha_0 \rceil \cdot \alpha_0/(1-\alpha_0) > 1$, the threshold index must be at least 2. Therefore, the policy wins with a probability of 0.

For the second case, where $\alpha_1 < 1/2$, we examine an instance with $p_1 = p_2 = 1$ and $n = 2$. The estimated success probabilities must be $\hat{p}_1 = \hat{p}_2 = \alpha_1$. Further, the threshold index cannot be 2 because $\hat{p}_2/(1-\hat{p}_2) = \alpha_1/(1-\alpha_1) < 1$. Thus, the policy always stops at the first trial, resulting in a winning probability of 0.

Now, it remains to examine the last case where $\alpha_0 = 0$ and $\alpha_1 \geq 1/2$. In this case, the policy sets the index of the last success in the samples as a threshold because $\alpha_1/(1-\alpha_1) \geq 1$ and $\alpha_0/(1-\alpha_0) = 0$. We refer to this policy as the *From the Last Success (FLS)* policy, and we demonstrate that this policy cannot guarantee a winning probability exceeding $(1-e^{-4})/4 \approx 0.2454 < 1/4$ even when at least one success is guaranteed to occur, i.e., the sum of the odds R is positive infinity.

▶ **Theorem 6.** *The winning probability of the FLS policy for the* $(1,\infty)$*-last success problem is at most* $(1-e^{-4})/4$.

Proof sketch. For the instance where the success probabilities are $p_1 = 1$ and $p_2 = \cdots = p_n = 2/n$, the winning probability of the FLS policy can be computed as

$$\begin{aligned}
&\left(1-\frac{2}{n}\right)^{2n-2} + \sum_{k=2}^{n}(n-k+1)\cdot\frac{2}{n}\cdot\frac{2}{n}\cdot\left(1-\frac{2}{n}\right)^{2n-2k} \\
&= \left(1-\frac{2}{n}\right)^{\frac{n}{2}\cdot\frac{4(n-1)}{n}} + \frac{4}{n}\sum_{k=2}^{n}\left(1-\frac{k-1}{n}\right)\left(1-\frac{2}{n}\right)^{\frac{n}{2}\cdot 4(1-\frac{k}{n})} \\
&\to \frac{1}{e^4} + \int_0^1 4(1-x)\cdot\left(\frac{1}{e}\right)^{4(1-x)}\mathrm{d}x = \frac{1}{e^4} + \int_0^1 \frac{4x}{e^{4x}}\,\mathrm{d}x \\
&= \frac{1}{e^4} + \left[\frac{-(4x+1)}{4e^{4x}}\right]_0^1 = \frac{1}{e^4} - \frac{5}{4e^4} + \frac{1}{4} = \frac{1-e^{-4}}{4},
\end{aligned}$$

as n goes to infinity, where the limit is evaluated by $\lim_{x\to\infty}(1-1/x)^x = 1/e$ and interpreting the sum as a Riemann sum. ◀

3.2 Direct utilization of success positions

As demonstrated in the previous subsection, Bruss's policy with estimated success probabilities fails to achieve the winning probability of 1/4. Nevertheless, by appropriately setting the threshold, we propose a *deterministic* policy that guarantees a winning probability of 1/4 as long as the sum of the odds R is at least $(\sqrt{3}-1)/2$. Note that achieving a winning probability of 1/4 is best possible according to Theorem 1.

We consider stopping policies that determine the threshold index based only on the relative position from the success indices in the sample sequence. Then, the threshold needs to be set between the indices of the last success and the second-to-last success of the samples. Indeed, if the threshold is set after the last success index in the sample sequence, then the winning probability becomes zero for the instance with $n = 2$ and $(p_1, p_2) = (1, 0)$. Similarly, setting the threshold before or at the second last success index in the samples results in a winning probability of zero for the instance with $n = 2$ and $(p_1, p_2) = (1, 1)$. Setting the threshold at the last success (i.e., the FLS policy) is also not effective, as we showed in Theorem 6. Therefore, the threshold should be set between the indices of the last success and the second last success in the sample sequence.

We consider the policy of setting the threshold at the *next* index of the second last success in the sample sequence. We refer to this policy as the *After the Second Last Success (ASLS)* policy. If there are fewer than two successes in the samples, the threshold is assumed to be index 1. We demonstrate that the ASLS policy wins with a probability of at least 1/4 if the sum of the odds R is at least $(\sqrt{3}-1)/2$ (≈ 0.3660).

▶ **Theorem 7.** *The winning probability of the ASLS policy for the* $(1, R)$*-last success problem is at least* $1/4$ *if* $R \geq (\sqrt{3}-1)/2$ *and at least* $\frac{R(4+3R)}{4(1+R)^2}$ *if* $0 \leq R \leq (\sqrt{3}-1)/2$.

Proof. We give the proof only for the cases where $0 \leq p_i < 1$ for all $i \in [n]$. The other cases where $p_i = 1$ for some $i \in [n]$ can be proved by considering the continuity of the winning probability with respect to $p_1, \ldots, p_n$.

To simplify the analysis, we add two virtual trials that ensure success before the actual trials. Specifically, we place them at the 0th and -1st positions with $p_0 = p_{-1} = 1$ and $X_0 = X_{-1} = Y_0 = Y_{-1} = 1$. With these virtual trials, both sample and actual sequences must contain at least two successes. Let i_1 and i_2 be the indices of the last and the second last success observed in the sample sequence $(Y_i)_{i=-1}^n$, respectively. Similarly, let j_1 and j_2 be the indices of the last and the second last success observed in $(X_i)_{i=-1}^n$, respectively. We then analyze the winning probability of the ASLS policy based on the relative positions of i_1, i_2, j_1, j_2.

We classify possible outcomes into 16 cases of (a)–(p) as illustrated in Figure 1. Let p denote the probability that case (a) occurs. By symmetry, the probabilities that cases (b), (c), and (d) occur are also p each. However, the probability that case (g) (and also (h)) occurs is different from p. This difference arises because the actual sequence in case (g) may contain successes between i_1 and j_2. Let us denote q as the probability of case (e) occurring. By symmetry, the probability of case (f) occurring is q as well. Additionally, let r represent the probability of case (g) occurring, then the probability that case (h) occurs is also r due to symmetry.

Consider a realization of the sample and actual sequences. If the realization is classified as being in case (g), then swapping the realizations of X_{j_2} and Y_{j_2} results in case (a), (c), or (e). Conversely, if the realization is classified as being in case (a), (c), or (e), then swapping the realizations of X_{i_1} and Y_{i_1} leads to case (g). As such swaps yield sequences with the same realization probability, we can conclude that

$$r = 2p + q. \tag{1}$$

Similarly, let us denote the probability of case (i) (or case (j)) occurring as s and the probability of case (k) (or case (l)) occurring as t, respectively. The relative position of case (m) matches that of case (k), but we categorize them separately since the ASLS policy wins for case (m). We also consider case (n) separately from case (k) for the convenience of the analysis. Let the probability of cases (m) and (n) occurring be denoted as u. Moreover, let v be the probability of cases (o) occurring. However, in the case of no success, we separate it and denote it as case (p) with probability w since the ASLS policy cannot win. Note that there are two or more non-virtual successes for cases (a) to (i) and case (o). In all other cases, non-virtual successes are at most one. Since each realization is classified uniquely into one case in Figure 1, we obtain

$$4p + 2q + 2r + 2s + 2t + 2u + v + w = 8p + 4q + 2s + 2t + 2u + v + w = 1, \tag{2}$$

where the first equality holds by (1).

Note that $w = \prod_{i=1}^{n}(1-p_i)^2$ represents the probability that successes occur only in virtual trials. The probability of case (m) occurring can be expressed as

$$u = \left(\prod_{i_1=1}^{n}(1-p_{i_1})\right)\cdot\left(\sum_{j_2=1}^{n} p_{j_2}\prod_{i_2\in[n]\setminus\{j_2\}}(1-p_{i_2})\right)$$
$$= \prod_{i=1}^{n}(1-p_i)^2\cdot\sum_{j=1}^{n}\frac{p_j}{1-p_j} = wR. \tag{3}$$

The probability of having exactly one success in each of the actual and sample sequences, excluding the virtual trials, is

$$\left(\sum_{j=1}^{n} p_j\prod_{i\in[n]\setminus\{j\}}(1-p_i)\right)^2 = \left(\prod_{i=1}^{n}(1-p_i)\cdot\sum_{j=1}^{n}\frac{p_j}{1-p_j}\right)^2 = wR^2.$$

This situation is classified as case (e), (f), or (o). Thus, we have

$$wR^2 \le 2q + v. \tag{4}$$

The probability of the last successes in both the actual and sample sequences having the same index ($i_1 = j_1$) is

$$2s + v = \sum_{i=1}^{n} p_i^2 \prod_{j=i+1}^{n}(1-p_j)^2, \tag{5}$$

where the right-hand side is obtained by directly computing the probability, and the left-hand side is derived from the probabilities that cases (i), (j), and (o) occur. The probability that the indices of the second last successes of the actual and the sample sequences are the same (i.e., $i_2 = j_2$) is

$$2q + v - wR^2 = \sum_{i=1}^{n} p_i^2 \prod_{j=i+1}^{n}(1-p_j)^2\left(\sum_{j=i+1}^{n} r_j\right)^2 \tag{6}$$

where the right-hand side is obtained by directly computing the probability, and the left-hand side is derived from the probabilities that cases (e), (f), and (o) occur. Moreover, the probability t can be expressed as

$$t = \sum_{i=1}^{n} p_i^2 \prod_{j=i+1}^{n}(1-p_j)^2\left(\sum_{j=i+1}^{n} r_j\right). \tag{7}$$

By combining (5), (6), and (7), we have

$$t \le \frac{(2s+v) + (2q+v-w\cdot R^2)}{2} = q + s + v - \frac{w\cdot R^2}{2}, \tag{8}$$

since $(1 + (\sum_{j=i+1}^{n} r_j)^2)/2 \ge \sum_{j=i+1}^{n} r_j$ holds by the AM-GM inequality for every $i \in [n]$.

Since w represents the probability of no success in both the sample and the actual sequence, we have

$$w \le \left(\frac{1}{1+R}\right)^2 \tag{9}$$

by Lemma 2.

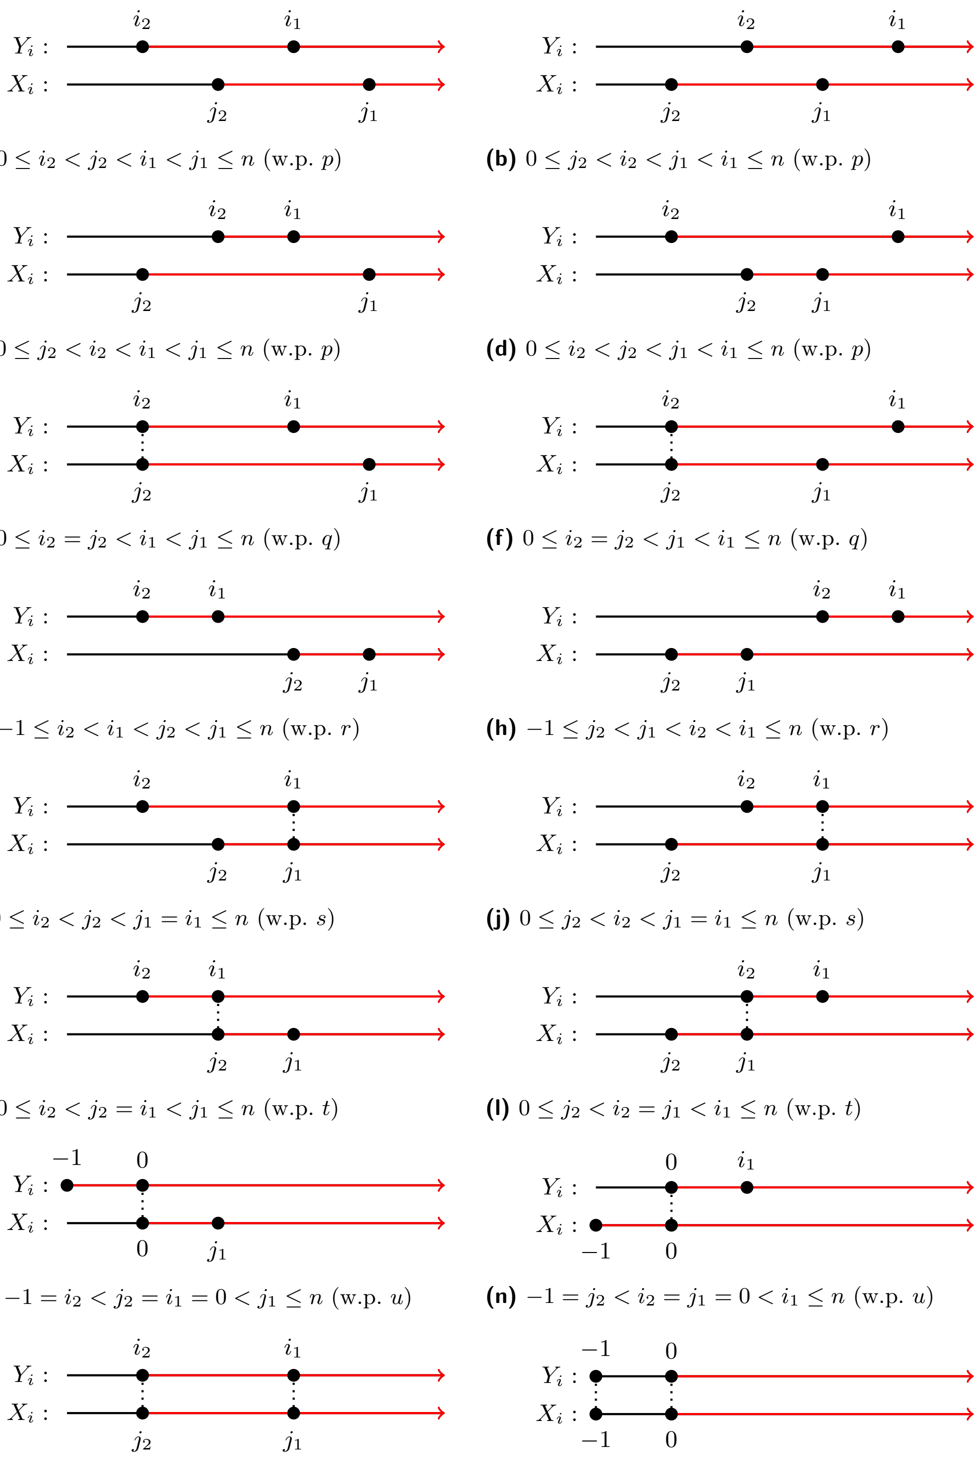

(a) $0 \le i_2 < j_2 < i_1 < j_1 \le n$ (w.p. p)
(b) $0 \le j_2 < i_2 < j_1 < i_1 \le n$ (w.p. p)
(c) $0 \le j_2 < i_2 < i_1 < j_1 \le n$ (w.p. p)
(d) $0 \le i_2 < j_2 < j_1 < i_1 \le n$ (w.p. p)
(e) $0 \le i_2 = j_2 < i_1 < j_1 \le n$ (w.p. q)
(f) $0 \le i_2 = j_2 < j_1 < i_1 \le n$ (w.p. q)
(g) $-1 \le i_2 < i_1 < j_2 < j_1 \le n$ (w.p. r)
(h) $-1 \le j_2 < j_1 < i_2 < i_1 \le n$ (w.p. r)
(i) $0 \le i_2 < j_2 < j_1 = i_1 \le n$ (w.p. s)
(j) $0 \le j_2 < i_2 < j_1 = i_1 \le n$ (w.p. s)
(k) $0 \le i_2 < j_2 = i_1 < j_1 \le n$ (w.p. t)
(l) $0 \le j_2 < i_2 = j_1 < i_1 \le n$ (w.p. t)
(m) $-1 = i_2 < j_2 = i_1 = 0 < j_1 \le n$ (w.p. u)
(n) $-1 = j_2 < i_2 = j_1 = 0 < i_1 \le n$ (w.p. u)
(o) $0 \le i_2 = j_2 < i_1 = j_1 \le n$ (w.p. v)
(p) $-1 = i_2 = j_2 < i_1 = j_2 = 0$ (w.p. w)

Figure 1 Classifications of the relative positions of i_1, i_2, j_1, j_2. Red lines indicate locations without success.

Now, we are ready to demonstrate that the ASLS policy guarantees a winning probability of 1/4. As depicted in Figure 1, the ASLS policy wins if the realization is classified as case (b), (c), (e), (f), (j), (m), or (o). Consequently, the winning probability is at least

$$\begin{aligned}
2p+2q+s+u+v &= \frac{8p+4q+2s+2t+2u+v+w}{4}+\frac{4q+2s-2t+2u+3v-w}{4} \\
&= \frac{1}{4}+\frac{4q+2s-2t+2u+3v-w}{4} && \text{(by (2))} \\
&\geq \frac{1}{4}+\frac{4q+2s-2(q+s+v-w\cdot R^2/2)+2u+3v-w}{4} && \text{(by (8))} \\
&= \frac{1}{4}+\frac{2q+2u+v+w(R^2-1)}{4} \\
&= \frac{1}{4}+\frac{2q+v+2wR+w(R^2-1)}{4} && \text{(by (3))} \\
&\geq \frac{1}{4}+\frac{wR^2+2wR+w(R^2-1)}{4} && \text{(by (4))} \\
&= \frac{1}{4}+\frac{w(2R^2+2R-1)}{4}.
\end{aligned}$$

Here, $(2R^2+2R-1)$ is non-negative if and only if $R \geq (\sqrt{3}-1)/2$. Thus, the winning probability is at least 1/4 if R is at least $(\sqrt{3}-1)/2$ by $w \geq 0$. Also, if $R < (\sqrt{3}-1)/2$, the winning probability is at least $\frac{1}{4}+\frac{2R^2+2R-1}{4(1+R)^2}=\frac{R(4+3R)}{4(1+R)^2}$ by (9). ◀

For the $(1,R)$-last success problem, the winning probability of any policy is at most 1/4 (Theorem 1) and at most R/e^R if $R \leq 1$ (Proposition 3). Introducing $\alpha \approx 0.357$ such that $\alpha/e^\alpha = 1/4$, the upper bound is summarized as R/e^R for $R \leq \alpha$ and 1/4 for $R \geq \alpha$. Consequently, the ASLS policy attains nearly the best possible winning probability for any R, as illustrated in Figure 2.

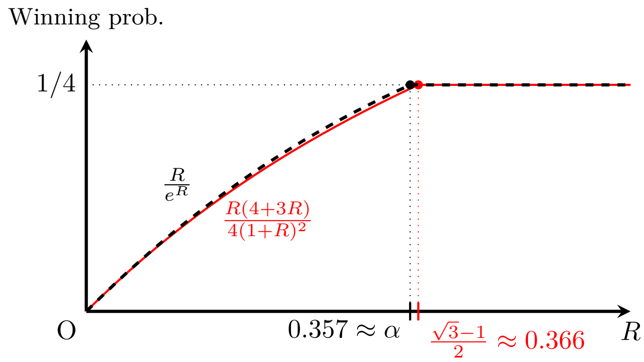

Figure 2 The winning probability of the ASLS policy (red line) and an upper bound of any policy for the $(1,R)$-last success problem (black dashed line).

3.3 Applying a policy by Nuti and Vondrák

Nuti and Vondrák [33] considered an optimal stopping problem called the adversarial order single sample secretary problem and gave a simple optimal stopping policy. Their policy can be converted into a stopping policy for the $(1,R)$-last success problem using randomization. This policy is a natural candidate that is expected to outperform the ASLS policy. However, this section observes that this is not the case. More precisely, we demonstrate that the policy achieves a winning probability of 1/4 only if R is at least 1/2. This means that the ASLS policy is superior in both performance and deterministic nature.

The adversarial order single sample secretary problem is described as follows: A sequence of independent real-valued random variables $V_1, V_2, \ldots, V_n$ is revealed one by one. Upon observing the value of V_i, an immediate and irrevocable decision must be made to either halt the observation or continue with other subsequent observations. It is assumed that the distribution that each random variable V_i follows is unknown, but a prior sample W_i that follows the same distribution as V_i is given for each $i \in [n]$. The goal is to maximize the probability of stopping at the highest value among V_i's. The single sample last success problem can be reduced to the adversarial order single sample secretary problem by considering a distribution where $V_i = i$ with probability p_i and 0 with probability $1 - p_i$, if we treat the no-success scenario as a win (or at least one success appears with probability 1).

For the adversarial order single sample secretary problem, Nuti and Vondrák [33] proposed the following simple policy: set the maximum value in the sample sequence as the threshold, and then stop when a value exceeding the threshold is observed. Here, if two values are the same, we make a random tie-break.[2] This policy wins with a probability of at least 1/4 as follows [33]: suppose the largest two numbers that appear in $V_1, \ldots, V_n, W_1, \ldots, W_n$ are a_1, a_2. Then, the policy wins if a_2 comes from $W_1, \ldots, W_n$ and a_1 comes from $V_1, \ldots, V_n$. Thus, it wins with probability 1/4 if a_1, a_2 come from different trials, and with probability 1/2 if a_1, a_2 come from the same trial.

This policy leads to the following randomized policy for the single sample last success problem: set the threshold index τ either to the index of the last success or to the next index of the last success in the sample sequence, each with a probability of 1/2. If no success appears in the sample sequence, it sets τ to be 1. We refer to this policy as the *From the Last Success, Randomized (FLSR)* policy. We demonstrate that the FLSR policy guarantees a winning probability of 1/4 only if $R \geq 1/2$, which is a stronger condition than the one of the ASLS policy. We remark that the above elegant analysis conducted by Nuti and Vondrák is not applicable to the FLSR policy if $R < \infty$. This is because a no-success scenario ($X_1 = \cdots = X_n = 0$) results in a loss in the last success problem, whereas its corresponding scenario ($V_1 = \cdots = V_n = 0$) leads to a win in the reduced adversarial order secretary problem.

▶ **Theorem 8.** *The winning probability of the FLSR policy for the* $(1, R)$*-last success problem is* $1/4$ *when* $R \geq 1/2$ *and at least* $\frac{1}{4} - \frac{1-2R}{4(1+R)^2}$ *when* $R \leq 1/2$*. Moreover, the winning probability is strictly less than* $1/4$ *if* $R < 1/2$*.*

4 Multisample model

In this section, we examine the last success problem with m samples for general m. Let $\mathrm{OPT}(R)$ be the optimal winning probability of the last success problem with complete information about the distributions when the sum of the odds is at least R. As mentioned in Proposition 3, $\mathrm{OPT}(R) = R/e^R$ if $R \leq 1$ and $\mathrm{OPT}(R) = 1/e$ if $R \geq 1$. For any positive constant ϵ, we propose a stopping policy that guarantees a winning probability of $\mathrm{OPT}(R) - \epsilon$ with a constant number of samples. Formally, we prove the following theorem.

▶ **Theorem 9.** *There exists a stopping policy that guarantees a winning probability of* $\mathrm{OPT}(R) - O(1/\sqrt[4]{m})$ *for the* (m, R)*-last success problem.*

[2] See, e.g., the paper of Rubinstein et al. [34] for more details on this tie-break. Note that without such a random tie-break, the winning probability of this policy is at most $(1 - e^{-4})/4$ by Theorem 6.

We emphasize that the number of samples required does not depend on the number of trials n. In addition, by Proposition 4, no policy can guarantee a winning probability of $1/e$ with a finite number of samples.

We propose a threshold-based policy with such a winning probability. The threshold is determined by the samples as follows. For each $i \in [n]$, define Q_i to be $\prod_{k=i}^{n}(1-p_k)$, which is the probability that the event $X_i = X_{i+1} = \cdots = X_n = 0$ occurs. For each $i \in [n]$ and $j \in [m]$, let $Y_{i,j} \in \{0,1\}$ be jth sample from ith distribution. The sequence $(Y_{1,j}, Y_{2,j}, \ldots, Y_{n,j})$ is said to be jth sample sequence. For each $i \in [n]$, let T_i be the number of sample sequences where no success is observed from i to n, i.e., $T_i = |\{j \in [m] \mid Y_{i,j} = Y_{i+1,j} = \cdots = Y_{n,j} = 0\}|$. Since the event $Y_{i,j} = Y_{i+1,j} = \cdots = Y_{n,j} = 0$ happens with probability Q_i, the random variable T_i follows a binomial distribution characterized by number m and probability Q_i. We also define $\hat{Q}_i$ as T_i/m, which is an unbiased estimator of Q.

Let ϵ be a real such that $0 < \epsilon < 1/2$. Our policy first calculates the index $\hat{\imath} = \arg\min\{i \in [n] \mid \hat{Q}_{i+1} \geq 1/e + \epsilon\}$ from samples and then stops at the first success observed from $\hat{\imath}$ in the actual trials. Here, we assume that $\hat{Q}_{n+1} = 1$.

We prepare some lemmas to prove the theorem. Firstly, we show that the winning probability of a threshold-based policy is close to $1/e$ if the threshold is an index i such that Q_i is approximately $1/e$.

▶ **Lemma 10.** *Let $i \in [n]$ and let δ be a positive real. If $Q_i \leq 1/e + \delta$ and $Q_{i+1} \geq 1/e$, the winning probability of the stopping policy that stops at the first success from the index i is at least $1/e - \delta$.*

Secondly, we evaluate the probability that $\hat{\imath}$ satisfies the condition of Lemma 10 by showing that $\hat{Q}_i$ and Q_i are close with high probability. Let $D_i = (Q_i - \hat{Q}_i)/Q_i$, which represents the relative difference of $\hat{Q}_i$ and Q_i. To bound the maximum deviation, we use Doob's inequality (see, e.g., [17]), which states that, for any $\lambda > 0$ and any non-negative submartingale[3] $S_1, S_2, \ldots, S_n$, the following inequality holds:

$$\Pr\left[\max_{k=1}^{n} S_k \geq \lambda\right] \leq \mathbb{E}[S_n]/\lambda. \tag{10}$$

Let $i^* \in [n]$ be the minimum index i such that $Q_{i+1} \geq 1/e$, where we assume that $Q_{n+1} = 1$. Note that we have (i) $Q_1 \leq Q_{i^*} < 1/e$ or (ii) $Q_1 \geq 1/e$ and $i^* = 1$. Additionally, we have $p_{i^*+1}, p_{i^*+2}, \ldots, p_n < 1$ by $Q_{i^*+1} = \prod_{k=i^*+1}^{n}(1-p_k) \geq 1/e > 0$. We show that $D_n \ldots, D_{i^*+1}$ is a martingale.

▶ **Lemma 11.** *The sequence $D_n, \ldots, D_{i^*+1}$ is a martingale.*

From this lemma, we can conclude that the sequence of random variables $|D_n|, \ldots, |D_{i^*+1}|$ is a non-negative submartingale. By applying the Doob's inequality (10) to $|D_n|, \ldots, |D_{i^*+1}|$, we can obtain the following inequality.

▶ **Lemma 12.** $\Pr\left[\max_{k=i^*+1}^{n} D_k \geq \epsilon\right] \leq \sqrt{\frac{e}{m\epsilon^2}}$.

Furthermore, we have the following lemma by Hoeffding's inequality [26].

▶ **Lemma 13.** $\Pr\left[\hat{Q}_{i^*} \geq \frac{1}{e} + \epsilon\right] \leq \exp\left(-2\epsilon^2 m\right)$ *if $Q_1 < 1/e$.*

[3] A sequence of random variables $S_1, \ldots, S_n$ is called a *martingale* if $\mathbb{E}[S_{i+1} \mid S_1, \ldots, S_i] = S_i$ for all $i \in [n-1]$, and a *submartingale* if $\mathbb{E}[S_{i+1} \mid S_1, \ldots, S_i] \geq S_i$ for all $i \in [n-1]$.

Next, we demonstrate that $\hat{i}$ satisfies the condition of Lemma 10 if the events considered in Lemmas 12 and 13 occur.

▶ **Lemma 14.** *If* $Q_1 < 1/e$, $\max_{k=i^*+1}^n D_k < \epsilon$ *and* $\hat{Q}_{i^*} < 1/e + \epsilon$, *then* $Q_{\hat{i}+1} \geq 1/e$ *and* $Q_{\hat{i}} \leq 1/e + 3\epsilon$.

Finally, we use the union bound to show that the probability of $\max_{k=1}^n D_k < \delta$ or $\hat{Q}_{i^*} < \frac{1}{e} + \delta$ occurring is small. Then, by combining this fact with Lemma 10, we prove Theorem 9.

Proof of Theorem 9. We choose ϵ as the one that satisfies $m = e/\epsilon^4$, which leads to $\epsilon = O(1/\sqrt[4]{k})$. Since we perform an asymptotic analysis, we may assume that $0 < \epsilon < 1/2$.

We first consider the case where $Q_1 < 1/e$. In this case, $Q_{i^*} < 1/e$ by the definition of i^*. By Lemmas 12 and 13, $\max_{k=i^*+1}^n D_k \geq \epsilon$ or $\hat{Q}_{i^*} \geq 1/e + \epsilon$ happens with probability at most $\sqrt{\frac{e}{m\epsilon^2}} + \exp(-2\epsilon^2 m)$ by the union bound. Then, we have

$$\sqrt{\frac{e}{m\epsilon^2}} + \exp(-2\epsilon^2 m) = \epsilon + \exp\left(-\frac{2e}{\epsilon^2}\right) \leq \epsilon + \frac{\epsilon^2}{2e} \leq 2\epsilon,$$

where the first inequality holds since $e^{-x} \leq 1/(1+x) \leq 1/x$ by $e^x \geq 1 + x$, and the last inequality holds by the assumption that $\epsilon < 1/2$. Therefore, the probability that both $\max_{k=i^*+1}^n D_k < \epsilon$ and $\hat{Q}_{i^*} < 1/e + \epsilon$ happen is at least $1 - 2\epsilon$. Under these conditions, we have $Q_{\hat{i}+1} \geq 1/e$ and $Q_{\hat{i}} \leq 1/e + 3\epsilon$ by Lemma 14. Using such an index $\hat{i}$, the threshold policy wins with probability at least $1/e - 3\epsilon$ by Lemma 10. Thus, the overall winning probability is

$$(1 - 2\epsilon)\left(\frac{1}{e} - 3\epsilon\right) \geq \frac{1}{e} - 4\epsilon = \frac{1}{e} - O\left(\frac{1}{\sqrt[4]{m}}\right).$$

Next, suppose that $Q_1 \geq 1/e$. In this case, we have $i^* = 1$. Then, by Lemma 12, we have $\max_{k=2}^n D_k < \epsilon$ with probability at least $1 - \sqrt{\frac{e}{m\epsilon^2}} = 1 - \epsilon$ by $m = e/\epsilon^4$. In what follows, we analyze the winning probability under the assumption that $\max_{k=2}^n D_k < \epsilon$. By the definition of $\hat{i}$, there are two cases where (i) $\hat{i} = 1$ or (ii) $\hat{i} \geq 2$ and $\hat{Q}_{\hat{i}} < 1/e + \epsilon$. If $\hat{i} = 1$, the winning probability is $Q_1 R$, which is at least $1/e$ if $R \geq 1$, and at least R/e^R if $R < 1$ by Lemma 2. If $\hat{i} \geq 2$ and $\hat{Q}_{\hat{i}} < 1/e + \epsilon$, we have $(Q_{\hat{i}} - \hat{Q}_{\hat{i}})/Q_{\hat{i}} = D_{\hat{i}} < \epsilon$, and hence

$$Q_{\hat{i}} < \frac{\hat{Q}_{\hat{i}}}{1-\epsilon} < \frac{1/e + \epsilon}{1 - \epsilon} < \left(\frac{1}{e} + \epsilon\right)(1 + 2\epsilon) = \frac{1}{e} + \left(\frac{2}{e} + 1\right)\epsilon + 2\epsilon^2 < \frac{1}{e} + 3\epsilon,$$

where the last two inequalities hold by $0 < \epsilon < 1/2$. Thus, the winning probability is at least $1/e - 3\epsilon$ by Lemma 10. In both cases (i) and (ii), the winning probability is at least $\mathrm{OPT}(R) - 3\epsilon$. Further, as the assumption of $\max_{k=2}^n D_k < \epsilon$ happens with probability $1 - \epsilon$, the overall winning probability is at least

$$(1 - \epsilon)(\mathrm{OPT}(R) - 3\epsilon) \geq \mathrm{OPT}(R) - 4\epsilon = \mathrm{OPT}(R) - O\left(\frac{1}{\sqrt[4]{m}}\right).$$ ◀

References

1. Saeed Alaei, Ali Makhdoumi, and Azarakhsh Malekian. Revenue maximization under unknown private values with non-obligatory inspection. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 27–28, 2021.

2 C. J. Argue, Alan Frieze, Anupam Gupta, and Christopher Seiler. Learning from a sample in online algorithms. In *Proceedings of the 36th Conference on Neural Information Processing Systems*, volume 35, pages 13852–13863, 2022.

3 Pablo D. Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th annual ACM-SIAM symposium on Discrete algorithms*, pages 1358–1377, 2014.

4 F. Thomas Bruss. Sum the odds to one and stop. *Annals of Probability*, pages 1384–1391, 2000.

5 F. Thomas Bruss. A note on bounds for the odds theorem of optimal stopping. *Annals of Probability*, 31(4):1859–1961, 2003.

6 F. Thomas Bruss. A mathematical approach to comply with ethical constraints in compassionate use treatments. *arXiv preprint arXiv:1804.04628*, 2018.

7 F. Thomas Bruss and Guy Louchard. The odds algorithm based on sequential updating and its performance. *Advances in Applied Probability*, 41(1):131–153, 2009.

8 Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. Single-sample prophet inequalities via greedy-ordered selection. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1298–1325, 2022.

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

10 José R. Correa, Andrés Cristi, Boris Epstein, and José A. Soto. Sample-driven optimal stopping: From the secretary problem to the i.i.d. prophet inequality. *Mathematics of Operations Research*, 49(1):441–475, 2024.

11 José R. Correa, Andrés Cristi, Laurent Feuilloley, Tim Oosterwijk, and Alexandros Tsigonias-Dimitriadis. The secretary problem with independent sampling. In *Proceedings of the 32nd ACM-SIAM Symposium on Discrete Algorithms*, pages 2047–2020, 2021.

12 José R. Correa, Paul Dütting, Felix Fischer, and Kevin 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.

13 José R. Correa, Paul Dütting, Felix Fischer, Kevin Schewior, and Bruno Ziliotto. Unknown iid prophets: Better bounds, streaming algorithms, and a new impossibility. *arXiv preprint arXiv:2007.06110*, 2020.

14 Andres Cristi and Bruno ZiliottoFeng. Prophet inequalities require only a constant number of samples. In *Proceedings of the 56th ACM Symposium on Theory of Computing*, pages 491–502, 2024.

15 Morris H. DeGroot. *Optimal statistical decisions.* McGraw-Hill, 1970.

16 Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91(C):318–333, 2015.

17 Rick Durrett. *Probability: Theory and Examples.* Cambridge University Press, 2019.

18 Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1452–1465, 2021.

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

20 Paul Dütting, Silvio Lattanzi, Renato Paes Leme, and Sergei Vassilvitskii. Secretaries with advice. *Mathematics of Operations Research*, 2023.

21 Yiding Feng, Jason D. Hartline, and Yingkai Li. Revelation gap for pricing from samples. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1438–1451, 2021.

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

23 Khashayar Gatmiry, Thomas Kesselheim, Sahil Singla, and Yifan Wang. Bandit algorithms for prophet inequality and pandora's box. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 462–500, 2024.

24 Nick Gravin, Hao Li, and Zhihao Gavin Tang. Optimal prophet inequality with less than one sample. In *Proceedings of the 18th International Conference on Web and Internet Economics*, pages 115–131, 2022.

25 Theodore P. Hill and Ulrich Krengel. A prophet inequality related to the secretary problem. *Contemporary Mathematics*, 125:209–215, 1992.

26 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, pages 409–426, 1994.

27 Benoît Iung, Eric Levrat, and Edouard Thomas. 'Odds algorithm'-based opportunistic maintenance task execution for preserving product conditions. *CIRP annals*, 56(1):13–16, 2007.

28 Haim Kaplan, David Naori, and Danny 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.

29 Haim Kaplan, David Naori, and Danny Raz. Online weighted matching with a sample. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1247–1272, 2022.

30 Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

31 Eric Levrat, Edouard Thomas, and Benoît Iung. Odds-based decision-making tool for opportunistic production-maintenance synchronization. *International Journal of Production Research*, 46(19):5263–5287, 2008.

32 Will Ma, David Simchi-Levi, and Chung-Piaw Teo. On policies for single-leg revenue management with limited demand information. *Operations Research*, 69(1):207–226, 2021.

33 Pranav Nuti and Jan Vondrák. Secretary problems: The power of a single sample. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2015–2029, 2023.

34 Aviad Rubinstein, Jack Wang, and S. Matthew 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.

35 Mitsushi Tamaki. An optimal parking problem. *Journal of Applied Probability*, 19(4):803–814, 1982.

36 Andrew C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 222–227, 1977.

37 Toru Yoshinaga and Yasushi Kawase. The last success problem with samples. *arXiv preprint arXiv:2308.09356*, 2023.