Secretaries with Advice

PAUL DÜTTING, Google Research

SILVIO LATTANZI, Google Research

RENATO PAES LEME, Google Research

SERGEI VASSILVITSKII, Google Research

The secretary problem is probably the purest model of decision making under uncertainty. In this paper we ask which advice can we give the algorithm to improve its success probability?

We propose a general model that unifies a broad range of problems: from the classic secretary problem with no advice, to the variant where the quality of a secretary is drawn from a known distribution and the algorithm learns each candidate's quality on arrival, to more modern versions of advice in the form of samples, to an ML-inspired model where a classifier gives us noisy signal about whether or not the current secretary is the best on the market.

Our main technique is a factor revealing LP that captures all of the problems above. We use this LP formulation to gain structural insight into the optimal policy. Using tools from linear programming, we present a tight analysis of optimal algorithms for secretaries with samples, optimal algorithms when secretaries' qualities are drawn from a known distribution, and a new noisy binary advice model.

CCS Concepts: • **Theory of computation** → **Online algorithms**.

Additional Key Words and Phrases: secretary problem, machine learning advice, prophet inequality

ACM Reference Format:

Paul Dütting, Silvio Lattanzi, Renato Paes Leme, and Sergei Vassilvitskii. 2021. Secretaries with Advice. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3465456.3467623>

1 INTRODUCTION

The secretary problem captures one of the purest forms of decision making under uncertainty: n candidates arrive one at a time, and must either be accepted or rejected on the spot. The goal is to design a stopping policy that maximizes the probability of selecting the best candidate.

Some additional information is necessary to find policies that have provable performance guarantees. A standard assumption is that candidates arrive in random order, in which case one can achieve a $1/e \approx 0.36$ competitive ratio by hiring the first candidate who is better than the first n/e applicants.

However, pure random arrival is not the only plausible paradigm, and previous work considered other information augmentation schemes. For instance, Gilbert and Mosteller [16] posit that each candidate has a quality in $[0, 1]$, which is drawn independently from some fixed and known distribution. Gilbert and Mosteller claimed the optimality of a certain sequence of decreasing thresholds, and showed numerically that this leads to a winning probability of ≈ 0.58 [16]. Follow-up work formally showed that this policy is optimal and that the winning probability of the optimal

This work is licensed under a Creative Commons Attribution International 4.0 License.

EC '21, July 18–23, 2021, Budapest, Hungary.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8554-1/21/07.

<https://doi.org/10.1145/3465456.3467623>

policy converges to $e^{-c} + (e^c - c - 1) \int_1^\infty x^{-1} e^{-cx} dx \approx 0.580164$ where $c \approx 0.804352$ is the unique solution to $\sum_{j=1}^\infty c^j / (j!j) = 1$ [6, 17, 33].

Another natural form of advice are samples. Indeed there has been a flurry of recent work on stopping problems, especially the prophet inequality variety, with limited information about an underlying distribution in the form of samples (e.g., [4, 9–11, 21, 32]). In a model popularized by Kaplan et al. [21], for example, an adversary chooses $k + n$ numbers. A random subset of k of these numbers are revealed to the decision maker (as samples) at the outset. Afterwards, the decision maker gets to inspect the remaining n numbers in an online fashion, and in random order.

While equally well motivated fairly little is known about these sampling models for the secretary objective. The only exemption is an elegant paper by Correa et al. [9], which studies a variant of the Kaplan et al. model with an additional independence assumption that serves to increase the mathematical tractability of the problem. In this variant of the model, an adversary chooses n numbers, and each number is marked independently with probability p as a sample. In this setting, they show that the optimal policy for any n and p is a threshold policy, obtain closed form solutions for thresholds when $n \rightarrow \infty$, and use this to recover the optimal approximation guarantees of the classic no advice model of $1/e$ (when $p \rightarrow 0$) and of the known distributions setting of ≈ 0.58 (when $p \rightarrow 1$).

A different kind of information augmentation occurs frequently in real job markets in the form of recommendation letters. Suppose each candidate comes with a letter that makes claims about the candidate being the best in the pool. Obviously if the recommender is known to be 100% accurate, the problem becomes easy. However, if the recommender can be wrong with some probability, the choice of the optimal policy is less obvious.

All of these scenarios—assumptions on random arrival, availability of quality scores, and recommendation letters—can be seen as a kind of “advice” given to the algorithm. In this work we take a general view of this problem, and explore secretary problems with advice.

1.1 Our Contribution

We unify disparate information augmentation settings for the secretary problem into a single framework that is rich enough to capture the classic random arrival model [13], the Gilbert-Mosteller i.i.d. setting [16], its Markovian generalizations [2, 12, 18, 34], and the recently introduced sample-based variants [9, 21].

In all of these settings, previous work painstakingly showed existence and optimality of threshold policies, which essentially specify the minimum hiring bar for every time step. We identify structural properties of the advice that explain the optimality of such policies as well as give an algorithm for quickly computing the optimum thresholds.

To demonstrate the utility of the framework we investigate a new machine-learning inspired advice for this problem, prove that the optimal policy is a threshold policy, and give tight bounds on its performance.

Setup. Consider a setting where each secretary i has a rank r_i . In the generic case without ties, the ranks form a permutation of $[n] = \{1, \dots, n\}$; we assume that higher is better so that $r_i = n$ is the secretary we seek to hire. We pair each secretary with an abstract signal s_i from some signal space $\mathcal{S}$, and assume that ranks and signals are drawn from a known joint distribution F .

The crux of the problem is that although the algorithm knows this distribution, it cannot observe the rank of the arriving secretary directly. Instead, it can observe the relative order of the secretaries that have arrived so far, along with their signals, and use this information to deduce the likelihood of the current secretary being the best.

$$\begin{array}{l|l}
 \max \sum_i \sum_s z(i, s) \cdot a(i, s) \text{ s.t.} & \min \sum_i \sum_s u(i, s) \text{ s.t.} \\
 z(i, s) \leq 1 - \sum_{j < i} \sum_{s'} z(j, s') \cdot c(i, s, j, s') & u(i, s) + \sum_{j > i} \sum_{s'} u(j, s') \cdot c(j, s', i, s) \geq a(i, s) \\
 z(i, s) \geq 0 & u(i, s) \geq 0
 \end{array}$$

Fig. 1. Factor revealing LP for secretaries with advice and its dual.

Techniques. We identify two properties of signalling schemes, *Non-Filtering (NF)* and *History-Irrelevance (HI)*, that allow us to write down a factor revealing linear program that quantifies the performance of the best policy. The two properties restrict the joint distribution, F , on the signals and the ranks. Roughly speaking, Non-Filtering implies that future signals do not change the information state available to the algorithm when processing a particular candidate. History-Irrelevance states that, conditioned on the current candidate being best so far, the probability of the candidate being globally optimum is independent of previously seen signals. We describe these precisely in Section 2.2.

We show the factor revealing LP and its dual in Figure 1. The linear program is a generalization of the one proposed by Buchbinder et al. [7] for the classic secretary program, with additional considerations to capture the effect of the signals. In the primal, the variable $z(i, s)$ denotes the probability of accepting a candidate with signal s at time i provided that it is the best candidate seen thus far. The coefficients $a(i, s)$ and $c(i, s, j, s')$ depend only on the joint distribution, F . In particular, $a(i, s)$ is the probability of seeing the top rank element in position i with signal s , and $c(i, s, j, s')$ is the probability that the candidate in position j with signal s' is the best candidate thus far, given that candidate in position $i > j$ with signal s is also locally optimum.

Next, we characterize the policies captured by our linear program formulation. Here we show that under the NF and HI restrictions any stopping policy is captured by the primal, and any solution to the primal can be converted to a stopping rule.

PROPOSITION A (RESTATEMENT OF PROPOSITION 9). *For any policy for the secretary problem with a signaling scheme satisfying NF and HI there is a set of values $z(i, s) \geq 0$ such that the objective of the factor revealing linear program corresponds to the probability that the highest ranked secretary is selected.*

The converse is also true:

PROPOSITION B (RESTATEMENT OF PROPOSITION 10). *For any signaling scheme satisfying NF and HI, given feasible values $z(i, s) \geq 0$ there is a policy for the secretary problem which picks the highest ranked secretary with probability equal to the objective of the factor revealing linear program.*

We delve deeper to understand the kinds of optimal policies generated by the LP. Let *memoryless* policies be those that make the hiring decision at time i based only on the signal of the secretary i , and *not* on the signals of previous candidates. Memoryless policies are a natural class, since their implementation does not require tracking previously seen signals, and can be computed in constant space.

We prove that any policy captured by the LP in Figure 1 is a *memoryless* policy.

THEOREM C (RESTATEMENT OF THEOREM 8). *If a signaling structure satisfies Non-Filtering and History-Irrelevance, then the optimal signaling policy is a memoryless policy.*

Fig. 2. Performance of the optimal policy for secretary with samples. Both plots are for $n = 20$. On the left, the x-axis shows k and the y-axis shows the probability of winning. On the right, the x-axis is $p = k/(k + n)$ and the y-axis is probability of winning.

An important subclass of memoryless policies are *threshold* policies. These policies associate an earliest acceptance time for every signal, and accept the first locally optimum candidate whose signal passes the test. We characterize algorithmically when threshold policies are optimal under HI and NF: They are optimal precisely when the dual to the linear program can be solved optimally by the greedy (backward induction) algorithm and satisfies a natural monotonicity assumption.

PROPOSITION D (RESTATEMENT OF PROPOSITION 11). *Assume NF and HI. If the greedy solution $u(i, s)$ for $i \in N$ and $s \in S$ constructed via backwards induction solution is an optimal solution to the dual LP and for all $s \in S$, $u(i, s)$ is non-decreasing in i , then the optimal policy is a threshold policy.*

and

PROPOSITION E (RESTATEMENT OF PROPOSITION 12). *Assume NF and HI. If the optimal policy is a threshold policy, then the greedy backwards induction solution is optimal for the dual LP.*

We also show that the monotonicity assumption in Proposition D is required: there are cases where the greedy algorithm leads to an optimal solution, but the solution violates monotonicity, and the optimal policy is not a threshold policy.

A key advantage of our approach is that it allows us to characterize the family of optimal policies fairly easily, in contrast with previous work where ad hoc lower bounds techniques were required. Furthermore, our factor-revealing linear program allows us to derive explicitly optimal policies and their success probability for several problems. We can obtain such results by a careful backward induction argument on the dual of our factor-revealing linear program. For example, in this way we can recover the optimal policy for the Gilbert and Mosteller setting, and for our binary setting with advice.

Applications. We highlight some of the results that can be derived using our framework.

First we focus on secretary algorithms for the sampling model of Kaplan et al. [21]. In this model an adversary chooses $n + k$ numbers, a random subset of k of these numbers are presented to the algorithm as *samples* at the outset, the remaining n numbers are presented to the algorithm one-by-one, in random order. This model smoothly interpolates between the classic secretary problem with no advice ($k = 0$) and the known distributions setting of Gilbert and Mosteller ($k \rightarrow \infty$).

In Section 5 we show how to capture this problem within our LP approach; as a consequence we obtain an LP formulation for the Gilbert-Mosteller model. We use the dual to show the optimality of thresholding policies for all number of samples k and all numbers of candidates n . The dual then has the interpretation that the $u(i, s)$ variables are the contribution to the optimal winning probability from secretary i when seeing signal s . The dual also gives rise to an efficient poly-time algorithm for computing the optimal policy and the corresponding winning probability. See Figure 2 for an exemplary set of these bounds.

Fig. 3. Left plot shows the performance of the optimal secretary with advice policy as a function of p . The right plot shows the thresholds t_Y and t_N as a function of p .

Furthermore, we obtain exact solutions to the dual for all n and k , and thus the thresholds used by the optimal policy. These can be used to derive exact analytic expressions of the optimal winning probability in the asymptotic regime when $n \rightarrow \infty$. We do this for the Gilbert-Mosteller setting, and thus recover the aforementioned analytic expression of the winning probability.

Our results for this setting strengthen and extend the results of both [9] and [16]. We use our framework to give an LP-duality based proof of the optimality of thresholding for all values of n and k (without the simplifying independence assumption of [9]), we obtain an efficient (poly-time) algorithm for computing the optimal policy and winning probability for all values of n and k , and derive closed formulas for the thresholds used by the optimal policy in both the asymptotic and non-asymptotic regime.

We then give an example of a new kind of advice that can be easily analyzed in this framework. Suppose each secretary arrives with a binary signal indicating whether the candidate is globally optimum, but the signal is incorrect with some probability $1 - p$. Anecdotally such a signal can model recommendation letter writers that claim that a candidate is “best in their class.”

Formally, we consider the random order arrival model, and restrict the signal space to $\mathcal{S} = \{Y, N\}$. For the best candidate the signal is Y with probability p ; for each other candidate the signal is N also with probability p , for $p \geq 1/2$. In Section 6 we consider the extended case where the probability of error is different for two classes, capturing potential false positive and false negative trade-offs in a real life classifier.

We show that the optimum solution in this setting is a threshold policy with two thresholds:

$$t_Y \approx n \left(\frac{1}{p} - 1 \right)^{1/p} e^{1-1/p} \quad \text{and} \quad t_N \approx ne^{1-1/p}.$$

In other words, the optimum policy waits the first t_Y steps; then accepts any candidate with a Y signal that is better than all previous candidates until time t_N , and then accepts any candidate that is better than all already rejected candidates. We use the structure of the dual LP and a backward-induction argument to solve for u_i^Y and u_i^N and substitute these into the dual objective. A detailed proof of the derivation is deferred to the full version of the paper. Note that when $p = 1/2$, and thus the signal provides no additional information, we recover Dynkin’s classic policy. However, as p grows, the thresholds t_Y and t_N diverge, and the competitive ratio grows to $\approx (1/p - 1)^{(1/p-1)} \cdot e^{(1-1/p)}$.

We plot both the growth of the competitive ratio as well as the two threshold t_Y and t_N in Figure 3.

1.2 Related Work

Our work is closely related to three broad active research directions. First, our problem formulation is inspired and aligned with the nascent area of algorithms with machine-learned advice. Second, our technique can be seen as a unifying framework that emphasizes the role of advice in the secretary and prophets literature. Finally, on a technical level, our work is related to the literature using factor revealing LPs in approximation and online algorithms. In the following we briefly review the most related papers in those three areas.

Algorithms with ML Advice. Traditionally, the design and analysis of algorithms has focused on provable guarantees for worst-case inputs. A growing body of work explores how “machine learned advice” can be leveraged. Thanks to its practical applications several problems have been studied through this lens. Examples range from building better data structures [22, 28], to improved competitive and approximation ratios for several online tasks [23, 24, 26, 29–31], to cases where advice has been used to speed-up algorithms [1, 5] or to reduce their space complexity [19]. Our work can be seen as a formalization of the classic secretary problem in this general framework.

Secretaries and Prophets Literature. The secretary problem is one of the most studied problems in online optimization. The classic formulation of the problem, introduced by Dynkin [13], succeeds with probability approaching $1/e$. The guarantee of $1/e$ is known to be best possible, even when the values of the secretaries are drawn from an unknown distribution (e.g., [15]).

The same objective has also been studied in the setting where the n candidates are i.i.d. draws from a known distribution (which can w.l.o.g. be the uniform distribution on $[0,1]$) [6, 16, 17, 33]. The optimal algorithm sets a decreasing sequence of thresholds, and it can be shown that it accepts the best secretary with probability at least ≈ 0.58 . A recent paper by Esfandiari et al. [14] has considered a non-i.i.d. variant of this problem, and shows how to obtain a $1/e$ approximation with a single threshold; and provides an example of non-i.i.d. distributions where this is best possible.

Another popular generalization of the Gilbert-Mosteller model (because of its application in finance) are random walk models. An early example is [18]. More recent results include [2, 12, 34]. The main take away from this line of work is known as the “bang-bang principle”: if the random walk is balanced then you might as well stop immediately, if it’s biased upward you should wait until the end, if it’s biased downwards you should accept immediately.

Another interesting line of work in this area studies the secretary and prophet problems in the presence of a limited number of samples [9–11, 21, 32]. Most relevant in this context are the Kaplan et al. paper [21], which is the model we adopt here, and the paper by Correa et al. [9] as it is the only prior work that looks at the secretary objective.

Our work offers a unifying lens that captures all these problems as secretary problems with advice; and extends the known LP formulation for the classic secretary problem to all of the other problems. In particular enables structural insights about the form the optimal policy takes, e.g. when and why backward induction yields optimal solutions.

Factor Revealing LPs. Factor revealing LPs have been used in a number of algorithmic analyses. They were introduced in the context of designing approximation algorithms for the facility location problem, in conjunction with the dual fitting technique [20]. The technique has been extended to strongly factor revealing LPs by [25], who used it to analyze the KVV ranking algorithm for bipartite matching in the random order model. Another variant called tradeoff revealing LPs was introduced in [27] to analyze a greedy algorithm for the the AdWords problem.

In the context of secretary problems, there are two main precursors: The first is [7], which describes a LP that recovers the optimal $1/e$ approximation guarantee for the classic secretary problem. The second one is [8], which extends this formulation to the (j, k) -secretary problem. In

this variant of the problem the algorithm is allowed to retain j elements and the goal is to maximize the expected number of elements that are among the k best secretaries.

Our LP formulation is inspired by [7], but is much more general, and in particular enables—for the first time—a unified treatment of the two classic secretary problems through the lens of LPs.

Additional Related Work. In parallel to this work, Antonianidis et al. [3] have considered the value maximization variant of the secretary problem with advice (so a different objective than we consider here). The techniques and results of that paper are very different from those in this paper.

They consider, for example, the single choice problem, and as advice the anticipated quality ρ^* of the best secretary. They seek bounds that are at least $\alpha > 1/e$ when the advice is accurate (as measured by the Euclidean distance) and at least $\beta < 1/e$ when the advice is inaccurate. They obtain qualitatively similar results for more general combinatorial allocation problems, such as bipartite matching.

They do not characterize optimal policies and how their performance decays as the advice gets worse, and they also don't provide a general framework for studying different forms of advice.

2 MODEL OF SECRETARIES WITH ADVICE

In the original secretary problem, the algorithm's goal is hiring the best secretary from a set of n candidates. There is a total order on the candidates which is not known in advance. Candidates arrive in random order, and upon arrival, the algorithm is able to compare the candidate with each of the previously seen options. The algorithm must then irrevocably decide to either hire that secretary or pass, in which case that secretary is no longer available.

Mathematically, we can describe the problem as follows: Let $(r_1, \dots, r_n)$ be a permutation of $[n] = \{1, \dots, n\}$. We refer to r_i as the *rank* of the i -th arriving secretary. We say that $i \geq j$ (i.e. secretary i is at least as good as j) whenever $r_i \geq r_j$. Note that the best secretary has value n . The algorithm has no access to the ranks. Instead, at time i , it can only see the relative comparisons $i \geq j$ and/or $i \leq j$ for $j < i$. The goal is to maximize the probability with which we stop at the maximum.

While we described the model without the possibility of ties, in some of our applications it will be natural to allow for ties. In that case we will assume that $(r_1, \dots, r_n)$ is a vector with $r_i \in [n]$, and we will require that $r_i = n$ for at least one secretary i . Our goal will then be to stop at an i such that $r_i = n$ (of which there may be more than one).

We will use $\leq_{1..i}$ to represent the partial order induced on the first i elements.

Best-so-far Event. We define a probabilistic event that will play a key role in the analysis and definitions below. Let the *best-so-far* event T_i be:

$$T_i = \{i \geq j; \forall j < i\} \quad (1)$$

It is important to observe that since the goal of the algorithm is to pick the best secretary, one can assume without loss of generality that the algorithm only picks secretary i if it is the best-so-far.

Advice. We augment the secretary problem with an extra signal s_i for each arriving secretary. Let $\mathcal{S}$ be the space of signals. We assume that $(r_1, r_2, \dots, r_n, s_1, \dots, s_n)$ is drawn from a known joint distribution. In the no ties case, $r_1, \dots, r_n$ will just be a permutation of $[n]$ and $s_i \in \mathcal{S}$ for all i . With ties, we require that $r_i \in [n]$, $s_i \in \mathcal{S}$, and $r_i = n$ for some i .

In each period i the algorithm observes both the signal s_i and the relative comparisons $i \geq j$ or $i \leq j$ for $j < i$. The algorithm knows the joint distribution of ranks and signals, but it cannot observe ranks directly, it is limited to computing induced ranks.

The algorithm decides in each step i whether to stop or proceed. As before, the goal of the algorithm is to select the best candidate. Note that if there are two secretaries with the top rank, we can pick either one.

2.1 Examples of Secretary Problems with Advice

It is useful to keep some concrete examples in mind:

EXAMPLE 1 (SECRETARIES WITHOUT ADVICE [13]). If $\mathcal{S} = \{0\}$ and $(r_1, \dots, r_n)$ are distributed as a random permutation, we are back at the original secretary problem. The best optimal strategy for this problem (Dynkin's algorithm) finds the optimal secretary with probability $1/e \approx 0.37$.

EXAMPLE 2 (GILBERT-MOSTELLER [16]). Consider a fixed known distribution F over the real numbers. Let $\mathcal{S} = \mathbb{R}$, s_i is an independent sample from F and r_i represent the ranks induced by s_i , i.e., $r_i = k$ if s_i is the k -th largest value of among $(s_1, \dots, s_n)$. This stochastic version of the secretary problem is studied by Gilbert and Mosteller who show that with this extra information the algorithm can hire the best secretary with probability ≈ 0.58 .

A non-i.i.d. version of Gilbert-Mosteller model was studied by Esfandiari et al. [14] who show that when $s_i \sim F_i$ then it is possible to choose the optimal secretary with probability $1/e$; and this is best possible in the worst case. (Note that this does not follow from Dynkin's algorithm since the ranking induced by the random draws is no longer uniform random.)

EXAMPLE 3 (MARKOVIAN STOPPING). A generalization of the Gilbert-Mosteller setting is the following Markovian stopping problem: consider a Markov chain on space $\mathcal{S}$, i.e., a stochastic process $s_1, \dots, s_n$ where $\mathbb{P}[s_i \mid s_1, \dots, s_{i-1}] = \mathbb{P}[s_i \mid s_{i-1}]$ and let r_i be the ranks induced by signals assuming there is a total ordering defined on $\mathcal{S}$.

Markovian stopping problems are popular in finance, where they serve as proxies for investment problems. Hlynka and Sheahan [18], for example, study a simple "symmetric" random walk. The process starts on day zero with a reward of zero. Then on each of n days, with equal probability, either the reward is increased by one or it is decreased by one. The goal is to maximize the probability with which the process is stopped at the maximum reward of all days. They show that all strategies that skip a fixed number of $t \geq 0$ days, and then accept the first reward from day $t + 1$ onwards that is the highest so far actually achieve the exact same winning probability.

Subsequent work has identified this as the indifference case of what has become to be known as the "bang bang principle" [2, 12, 34]: In a random walk that is started at zero and in which the probability to move up by one is p and the probability to move down by one is $1 - p$ it is best to stop immediately when $p < 1/2$ and to wait until the end when $p > 1/2$.

EXAMPLE 4 (SECRETARIES WITH SAMPLES [9, 21]). An adversary writes down $n + k$ numbers. A random subset of size k is chosen and revealed to the algorithm as samples. Afterwards, the remaining n numbers are presented to the algorithm in an online fashion, in random order.

The algorithm can observe the relative order of all secretaries it has seen so far. So in addition to observing the best so far event T_i , the algorithm learns about the relative rank of the current secretary among the samples. The signal space is thus $\mathcal{S} = \{0, \dots, k\}$, where $s_i = j$ means that j of the samples are worse than the current secretary.

The rank r_i of a hireable secretary is its relative rank among the n hireable secretaries. So $r_i = n$ is the rank of the best secretary, and $r_i = 1$ is the rank of the worst secretary.

EXAMPLE 5 (BINARY CLASSIFIER). A natural ML-advice model is a classifier that given a secretary predicts whether it is the best secretary or not. The input to the machine learned system is a candidate with all of their features, and the output is a binary classification: $Y(es)$ or $N(o)$.

In ML it's common practice to evaluate the quality of a binary classifier using the following four metrics accuracy, precision, recall, and specificity. It will be more convenient for us to express our results in terms of recall (a.k.a. sensitivity) and specificity. Such metrics are depicted in Figure 4 as a function of True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN).

		Actual		
		Y	N	
Predicted	Y	TP	FP	Recall: $p = \frac{TP}{TP+FN}$
	N	FN	TN	Specificity: $p' = \frac{TN}{FP+TN}$

Fig. 4. Parameters of the binary classifier

This translates to a signal space $\mathcal{S} = \{Y, N\}$ and two parameters $p, p' \in [\frac{1}{2}, 1]$ corresponding to recall and specificity respectively. We will assume that $(r_1, \dots, r_n)$ is a random permutation and that:

$$\begin{aligned} \mathbb{P}[s_i = Y \mid r_i = n] &= p & \mathbb{P}[s_i = N \mid r_i = n] &= 1 - p \\ \mathbb{P}[s_i = Y \mid r_i = r] &= 1 - p' & \mathbb{P}[s_i = N \mid r_i = r] &= p' \end{aligned}$$

for every $r < n$. Essentially the advice suggest to hire the best secretary with probability p and suggest to hire any other secretary with probability $1 - p'$. When $p = p' = \frac{1}{2}$ we are back in the original secretary problem since the signal is clearly useless. When $p = p' = 1$ the ML model is perfect and we can hire the optimal secretary by simply following its suggestion, i.e., hiring whenever it says Y and not hiring otherwise.

2.2 Signal Structure Properties

The joint distribution in all of the examples we described above has two properties that will be key to our analysis: *Non-Filtering* and *History Irrelevance*. As we note below in Example ?? these are not universal and rule out certain types of advice.

DEFINITION 6 (NON-FILTERING). We say that a signaling structure satisfies *Non-Filtering (NF)* if given any $j < i$ we have:

$$\mathbb{P}[\leq_{1..j}, s_1, \dots, s_j \mid T_j, s_j] = \mathbb{P}[\leq_{1..j}, s_1, \dots, s_j \mid T_j, s_j, T_i, s_i]$$

where T_i is the best-so-far event defined in equation (1).

The property essentially means that if at both i and j we see secretaries that are the best-so-far the information we see later in i does not affect the conditional distribution of information the algorithm gets at time j .

Note that it holds in Examples 2, 3 simply by the Markovian property: conditioned on T_j, s_j , we have that $(\leq_{1..j}, s_1, \dots, s_j)$ and $(s_j, s_{j+1}, \dots, s_n)$ are independent. Since s_i, T_i is a function of $(s_j, s_{j+1}, \dots, s_n)$, property NF automatically follows. A similar, but more delicate argument which we defer to Section 5, can be used to argue that Example 4 also satisfies NF. In Example 5 if $\mathbb{P}'(N) = p'$ and $\mathbb{P}'(Y) = 1 - p'$ we have:

$$\mathbb{P}[\leq_{1..j}, s_1, \dots, s_j \mid T_j, s_j, T_i, s_i] = \frac{1}{(j-1)!} \prod_{t=1}^{j-1} \mathbb{P}'(s_t)$$

since after we know that j is the best so far, each element before j cannot be the top ranked element and hence must have its signals drawn from $\mathbb{P}'$.

Non-Filtering captures the effect of future arrivals on the information state available to the algorithm. In contrast, history irrelevance imposes structure on the relationship between signals previously observed.

DEFINITION 7 (HISTORY-IRRELEVANCE). *We say that a signaling structure satisfies History-Irrelevance (HI) if conditioned on (T_i, s_i) the variable $(\prec_{1..i}, s_1, \dots, s_i)$ representing the information available at round i and the event $r_i = n$ are independent, i.e. $P(r_i = n \wedge (\prec_{1..i}, s_1, \dots, s_i) | T_i, s_i) = P(r_i = n | T_i, s_i)P(\prec_{1..i}, s_1, \dots, s_i | T_i, s_i)$.*

Fig. 5. Graphical model representing property HI

In graphical model language, we can say that the variables and event above satisfy the graphical model in Figure 5. It is useful to observe some implications of this fact. The first implication is that in order to determine the probability that each i is the highest ranked secretary, it is enough to look at its signal and to whether it is the highest ranked so far:

$$\mathbb{P}[r_i = n | \prec_{1..i}, s_1, \dots, s_i] = \mathbb{P}[r_i = n | T_i, s_i] \quad (2)$$

A second implication is that if X is an event that depends only on the the information available at period i (i.e., depends only on $(\prec_{1..i}, s_1, \dots, s_i)$) then:

$$\mathbb{P}[X | s_i, r_i = n] = \mathbb{P}[X | T_i, s_i] \quad (3)$$

since $r_i = n$ implies T_i .

Again it is simple to see that all the example discussed satisfy it. For Examples 2 and 3 it again follows from the same argument: $(\preceq_{1..j}, s_1, \dots, s_i)$ and $(s_i, s_{i+1}, \dots, s_n)$ are conditionally independent given T_i, s_i and $r_i = n$ depends only on $(s_i, s_{i+1}, \dots, s_n)$ whenever T_i holds. For Example 5 the probability that secretary i is the highest overall only depends on the first i elements and $s_1, \dots, s_i$ through T_i and s_i . Then, conditional on T_i , we distinguish two cases, namely $s_i = Y$ or $s_i = N$. The probability that $r_i = n$ conditional on T_i and $s_i = Y$ is $Pr[r_i = n | T_i]p / (Pr[r_i = n | T_i]p + (1 - Pr[r_i = n | T_i])(1 - p'))$. Similarly, for the case where $s_i = N$, the probability that $r_i = n$ conditional on T_i and $s_i = N$ is $Pr[r_i = n | T_i](1 - p) / (Pr[r_i = n | T_i](1 - p) + (1 - Pr[r_i = n | T_i])p')$. Now, $Pr[r_i = n | T_i] = i/n$ because $Pr[r_i = n] = 1/n$, $Pr[T_i] = 1/i$, and $Pr[r_i = n] = Pr[r_i = n | T_i]Pr[T_i]$. So the probability that $r_i = n$ conditional on T_i and $s_i = Y$ is $p / (p + ((n - i)/i)(1 - p'))$ and the probability that $r_i = n$ conditional on T_i and $s_i = N$ is $(1 - p) / (1 - p + ((n - i)/i)p')$. For Example 4 we again defer the discussion to Section 5.

3 MEMORYLESS POLICIES AND LP FORMULATION

Non-Filtering (NF) and History-Irrelevance (HI) are *natural* properties of a signalling scheme. In this section we investigate which policies are optimal under Non-Filtering (NF) and History-Irrelevance (HI).

Memoryless Policies. A generic policy is a map from the available information at period i , which consists of $(\preceq_{1..i}, s_1, \dots, s_i)$ to a stopping probability. In general the optimal generic policy can be quite complicated, specially when there are lots of correlations among signals. An interesting class of policies are the one where the decision is based on the advice given for the current secretary only and where the decision does not depend on the relative order among the first $i - 1$ elements,

but instead depends only on whether i is the best-so-far or not, i.e., whether we are in the event T_i . Putting it all together, we say that a *memoryless policy* is a map that outputs the probability of stopping in each round i conditioned on T_i and s_i .

Memoryless policies are well-motivated and a natural assumption in ML applications. The resulting policies are simpler, and require less space. They offer increased privacy and may even be a legal necessity (e.g., when data protection laws regulate which data may be stored and for how long).

Under the following natural conditions (satisfied by Examples 1, 2, 3, 4, and 5) the optimal signaling policy is memoryless:

THEOREM 8. *If a signaling structure satisfies Non-Filtering and History-Irrelevance, then the optimal signaling policy is a memoryless policy.*

Even under these nice conditions, it is not trivial to show the theorem above. As one tries to modify a generic policy to make it be oblivious to the history, one needs also to account for probability of not having stopped at any given point which is complicated to track.

LP Formulation and Proof of Theorem 8. The main tool we will use is a linear programming formulation for the secretary problem with advice subject to the NF and HI properties. Our formulation generalize the linear program of Buchbinder, Jain, and Singh [7] to our more challenging setting.

Assume the signaling structure satisfies NF and HI and consider a generic policy mapping $(\leq_{1..i}, s_1, \dots, s_i)$ to a stopping probability for each i . For each $i \in [n]$ and $s \in \mathcal{S}$, we define:

$$z(i, s) = \mathbb{P}[i \text{ is picked} \mid T_i, s_i = s]$$

We furthermore define for each $i \in [n]$ and $s \in \mathcal{S}$, $a(i, s) = \mathbb{P}[r_i = n, s_i = s]$ and for each $i, j \in [n]$ with $i > j$ and $s, s' \in \mathcal{S}$, $c(i, s, j, s') = \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s]$.

We will show that if $z(i, s)$ is “feasible” (to be defined soon), then it contains all the relevant data to reconstruct the policy.

Which constraints must $z(i, s)$ satisfy in order for us to be able to recover a policy from it? Besides the trivial requirement that $z(i, s) \geq 0$ the other requirement we will ask is that there is enough probability left to choose i when i arrives. Even if i is the top among the first i elements, an earlier element may have been chosen preventing us from choosing i later on. In other words, we have:

$$\mathbb{P}[i \text{ is picked} \mid T_i, s_i = s] + \sum_{j < i} \mathbb{P}[j \text{ is picked} \mid T_i, s_i = s] \leq 1.$$

Since each $j < i$ is only picked in the event T_j (j is the best so far) then we can re-write the expression above using the law of total probability as:

$$\mathbb{P}[i \text{ is picked} \mid T_i, s_i = s] + \sum_{j < i} \sum_{s' \in \mathcal{S}} \mathbb{P}[j \text{ is picked} \mid T_j, s_j = s', T_i, s_i = s] \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s] \leq 1.$$

Now we can apply property NF to argue that:

$$\mathbb{P}[j \text{ is picked} \mid T_j, s_j = s', T_i, s_i = s] = \mathbb{P}[j \text{ is picked} \mid T_j, s_j = s'].$$

Since the fact that j is picked depends only on $(\leq_{1..j}, s_1, \dots, s_j)$ and the conditional distribution of $(\leq_{1..j}, s_1, \dots, s_j)$ is the same given T_j, s_j or T_j, s_j, T_i, s_i we obtain the last display equation. Substituting it above and replacing the definition of $z(i, s)$ and $c(i, s, j, s')$ we obtain:

$$z(i, s) \leq 1 - \sum_{j < i} \sum_{s' \in \mathcal{S}} z(j, s') \cdot c(i, s, j, s'). \quad (4)$$

It is important to note that the term $c(i, s, j, s') = \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s]$ depends only on the joint distribution of $(r_1, \dots, r_n, s_1, \dots, s_n)$ and not on the policy itself.

We can also write the performance of the policy in terms of the $z(i, s)$ variables:

$$\begin{aligned}
 \text{Obj} &= \sum_{i=1}^n \mathbb{E}_{\leq_{1..i}, s_1, \dots, s_i} [\mathbb{P}[i \text{ is picked} \wedge r_i = n \mid \leq_{1..i}, s_1, \dots, s_i]] \\
 &= \sum_{i=1}^n \mathbb{E}_{\leq_{1..i}, s_1, \dots, s_i} [\mathbb{P}[i \text{ is picked} \mid \leq_{1..i}, s_1, \dots, s_i] \cdot \mathbb{P}[r_i = n \mid \leq_{1..i}, s_1, \dots, s_i]] \\
 &= \sum_{i=1}^n \mathbb{E}_{\leq_{1..i}, s_1, \dots, s_i} [\mathbb{P}[i \text{ is picked} \mid \leq_{1..i}, s_1, \dots, s_i] \cdot \mathbb{P}[r_i = n \mid T_i, s_i]] \\
 &= \sum_{i=1}^n \mathbb{E}_{s_i} [\mathbb{P}[i \text{ is picked} \mid T_i, s_i] \cdot \mathbb{P}[r_i = n \mid s_i]],
 \end{aligned}$$

where the first equality follows from independence, the second from HI (equation (2)) and the third follows from the law of conditional probability. Substituting $z(i, s)$ and $a(i, s)$ we get:

$$\text{Obj} = \sum_{i=1}^n \sum_{s \in S} z(i, s) \cdot a(i, s). \quad (5)$$

We note again that the term $a(i, s) = \mathbb{P}[r_i = n, s_i = s]$ depends only on the signaling structure and not on the policy itself. We showed the following statement:

PROPOSITION 9. *For any policy for the secretary problem with a signaling scheme satisfying NF and HI there is a set of values $z(i, s) \geq 0$ satisfying (4) such that the objective in equation (5) corresponds to the probability that the highest ranked secretary is selected.*

The converse is also true:

PROPOSITION 10. *Assume again that the signaling structure satisfies NF and HI. Given values $z(i, s) \geq 0$ satisfying (4) then there is a policy for the secretary problem which picks the highest ranked secretary with probability equal to the objective defined in (5).*

PROOF. Consider the policy that upon seeing secretary i with signal s chooses that secretary with probability:

$$q(i, s) = \frac{z(i, s)}{1 - \sum_{j < i} \sum_{s' \in S} z(j, s') \cdot \mathbb{P}[T_j \text{ and } s_j = s' \mid T_i \text{ and } s_i = s]} \quad (6)$$

if i is the best-so-far (in other words, if T_i happens) and zero otherwise. Now we need to argue that the probability that this policy chooses the highest ranked secretary is equal to the objective in equation (5).

Before we do that, we show that under the reconstructed policy, the probability that we pick the i -th secretary conditioned on $T_i, s_i = s$ is indeed $z(i, s)$. We will show that recursively. Assume for now it is true for all $j < i$. Under the constructed policy the probability that we pick i is the probability that we reach that step without picking any of the previous secretaries times the probability we choose i at that step:

$$\begin{aligned}
 \mathbb{P}[i \text{ is picked} \mid T_i, s_i = s] &= \mathbb{P}[\text{reach step } i \mid T_i, s_i = s] \cdot q(i, s) \\
 &= \left(1 - \sum_{j < i} \mathbb{P}[j \text{ is picked} \mid T_i, s_i = s]\right) \cdot q(i, s).
 \end{aligned}$$

We can now use the induction hypothesis to evaluate the probability that j is picked:

$$\begin{aligned} \mathbb{P}[j \text{ is picked} \mid T_i, s_i = s] &= \sum_{s' \in \mathcal{S}} \mathbb{P}[j \text{ is picked} \mid T_j, s_j = s', T_i, s_i = s] \cdot \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s] \\ &= \sum_{s' \in \mathcal{S}} \mathbb{P}[j \text{ is picked} \mid T_j, s_j = s'] \cdot \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s] \\ &= \sum_{s' \in \mathcal{S}} z(j, s') \cdot \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s], \end{aligned}$$

where the last equality follows from the induction hypothesis. Now taking the two previous display equations together and substituting the formula for $q(i, s)$ we obtain that:

$$\mathbb{P}[i \text{ is picked} \mid T_i, s_i = s] = z(i, s).$$

Equipped with that we can now bound the performance of the policy:

$$\text{Alg} = \sum_{i=1}^n \mathbb{P}[r_i = n] \cdot \left(\sum_{s \in \mathcal{S}} \mathbb{P}[s_i = s \mid r_i = n] \cdot \mathbb{P}[i \text{ is picked} \mid s_i = s, r_i = n] \right).$$

Since the probability that the algorithm picks i depends only on $(\leq_{1..i}, s_1, \dots, s_i)$ we can use property HI (equation (3)) to get that:

$$\mathbb{P}[i \text{ is picked} \mid s_i = s, r_i = n] = \mathbb{P}[i \text{ is picked} \mid T_i, s_i = s].$$

Putting it all together we get that:

$$\text{Alg} = \sum_{i=1}^n \mathbb{P}[r_i = n] \cdot \left(\sum_{s \in \mathcal{S}} \mathbb{P}[s_i = s \mid r_i = n] \cdot z(i, s) \right) = \text{Obj},$$

as claimed. $\square$

We now can show the proof of Theorem 8 as a corollary:

PROOF OF THEOREM 8. Given any policy use Proposition 9 to obtain variables $z(i, s)$. Now use Proposition 10 to convert this back to a policy. Observe that the policy we obtain from this transformation uses only s_i, T_i when deciding whether to stop. Hence, it is a memoryless policy. $\square$

Since all of the examples in Section 2 satisfy the NF and HI conditions, Theorem 8 implies that they all have an LP formulation, and moreover, they all have optimal memoryless policies.

We conclude this section observing that there exist examples where the signalling scheme does not satisfy NF and HI, and the optimal policy is not memoryless (for an example refer to the journal version).

4 OPTIMALITY OF THRESHOLD POLICIES

We showed that under NF and HI the optimal policy is memoryless. Here we investigate when the optimal strategy has the even simpler form of a threshold strategy. It will also be a good opportunity to study the structure of the dual LP.

Threshold Policies. The optimal policy for the secretary problem without advice (Example 1) is to wait until we have seen n/e secretaries and then pick the first secretary that is the best so far. This is a special case of what we call a *threshold policy*, which conceptually is just a further restriction of memoryless policies. Recall that we defined the *best-so-far* event T_i as:

$$T_i = \{i \geq j; \forall j < i\}.$$

Now we say that a stopping time τ is a threshold policy if there is a threshold function $t : \mathcal{S} \rightarrow [n]$ such that:

$$\tau = \min\{i \text{ s.t. } T_i \text{ and } i \geq t(s_i)\}.$$

Dynkin's algorithm for the secretary problem without advice has $t(\emptyset) = n/e$. It is useful to see what a threshold policy would look like for Yes/No advice (Example 5). A threshold policy in that case should specify threshold $t(N)$ and $t(Y)$. Assume for now that $t(Y) \leq t(N)$ (we will prove it should be the case later). Then the policy would say:

- No item is picked for $i < t(Y)$.
- For items arriving $t(Y) \leq i < t(N)$ they are chosen if they are the best so far and the advice is Yes.
- For items arriving $t \geq t(N)$ they are chosen if they are the best so far regardless of the advice.

For $p = 1$ the optimal policy is clearly a threshold policy with $t(Y) = 1$ and $t(N) = n$. For $p = \frac{1}{2}$ the optimal policy is the same as in Dynkin's algorithm: $t(Y) = t(N) = n/e$.

Characterization via Dual LP. We will see that whether or not the optimal policy is a threshold policy is closely related to our ability to solve the *dual of the LP* in the previous section with a greedy algorithm. More precisely, we will show that under NF and HI the optimal policy is a threshold policy if we can solve the dual with the greedy algorithm and the resulting solution satisfies a natural monotonicity assumption. This will be the case in both of our case studies. The reverse implication is always true: If the optimal policy is a threshold policy, then we can find it with the greedy algorithm.

To state our result more formally, let's first recall the LP from the previous section, and let's also derive its dual. Under NF and HI we have the following primal-dual pair:

$$\begin{array}{l|l} \max \sum_i \sum_s z(i, s) \cdot a(i, s) \text{ s.t.} & \min \sum_i \sum_s u(i, s) \text{ s.t.} \\ z(i, s) \leq 1 - \sum_{j < i} \sum_{s'} z(j, s') \cdot c(i, s, j, s') & u(i, s) + \sum_{j > i} \sum_{s'} u(j, s') \cdot c(j, s', i, s) \geq a(i, s) \\ z(i, s) \geq 0 & u(i, s) \geq 0 \end{array}$$

where $a(i, s)$ and $c(i, s, j, s')$ are coefficients in $[0, 1]$ that depend only on the joint distribution of signals and ranks. In particular, $a(i, s) = \mathbb{P}[s_i = s, r_i = n]$ and $c(i, s, j, s') = \mathbb{P}[T_j, s_j = s' \mid T_i, s_i = s]$.

Furthermore note that since the coefficients c are derived from a probability distribution then:

$$\sum_{s'} c(i, s, j, s') \leq 1. \quad (7)$$

Consider solving the dual LP with the following greedy algorithm: Set $u(n, s) = a(n, s)$ for all s . Then for $i < n$ assuming that we have set $u(j, s')$ for $j > i$ and all s' , set

$$u(i, s) = \max \left\{ 0, a(i, s) - \sum_{j > i} \sum_{s'} u(j, s') \cdot c(j, s', i, s) \right\}. \quad (8)$$

It is clear that this leads to a feasible dual solution.

More formally, using the tightness and by leveraging complementary slackness we prove the following two propositions.

PROPOSITION 11. *Assume NF and HI. If the greedy solution $u(i, s)$ for $i \in [n]$ and $s \in \mathcal{S}$ constructed via equation (8) is an optimal solution to the dual LP and for all $s \in \mathcal{S}$, $u(i, s)$ is non-decreasing in i , then the optimal policy is a threshold policy.*

PROOF. We proceed in two steps:

Step 1: For now, assume that $a(i, s)$ is *generic*. That is, in equation (8) we have $a(i, s) - \sum_{j>i} \sum_{s'} u(j, s') \cdot c(j, s', i, s) \neq 0$. In that case for each pair (i, s) we have that either (a) $u(i, s) = 0$ and the dual constraint is non-tight; or (b) $u(i, s) > 0$ and the dual constraint is tight. Define the threshold $t(s)$ to be the smallest i such that $u(i, s) > 0$. Then, because $u(i, s)$ is non-decreasing in i , for every $i < t(s)$ we are in case (a) and for every $i \geq t(s)$ we are in case (b).

By complementary slackness we have that in case (a) since the dual constraint is not tight, we must have $z(i, s) = 0$ and hence the probability of picking i given signal s and that i is the best so far should be $q(i, s) = 0$ (by equation (6)). In case (b) since $u(i, s) > 0$ the primal constraint should be tight and hence the probability of picking i should be 1 (again by equation (6)). This shows that the policy must be a threshold policy.

Step 2: The only missing part now is to argue what happens if $a(i, s)$ is not generic. (This will be a technical and not particularly algorithmic argument. The reader that doesn't care about corner cases may want to skip it). In that case we can appeal to a perturbation argument: consider $a_\epsilon(i, s) = a(i, s) + \epsilon \cdot n(i, s)$ where $n(i, s)$ is a random perturbation and ϵ is a small number (that we will send to zero). Then almost surely we will have that the LP will be generic (in the sense of Step 2). Hence the optimal solution z_ϵ will be a threshold policy. Take now a sequence of $\epsilon_t \rightarrow 0$ and consider the threshold policy solutions z_{ϵ_t} . Since the solutions z_{ϵ_t} live in a compact space, they must converge to some feasible solution z^* in the limit (passing to a subsequence if necessary). This solution must also be a threshold policy since the set of threshold policies is closed. To see that this is an optimal solution to the unperturbed LP, note that the perturbation only affects the objective function, hence z^* is feasible. Also note that for any feasible point z we have that $\langle a_\epsilon, z_\epsilon \rangle \geq \langle a_\epsilon, z \rangle$. Taking the limit as $\epsilon \rightarrow 0$ we get: $\langle a, z^* \rangle \geq \langle a, z \rangle$ for all feasible z , hence z^* is a solution to the unperturbed LP. $\square$

For the reverse direction we have:

PROPOSITION 12. *Assume NF and HI. If the optimal policy is a threshold policy, then the greedy backwards induction solution (equation (8)) is optimal for the dual LP.*

PROOF. Let $t^*(s)$ be the thresholds in the optimal policy and let u be any solution to the dual LP. Then the optimal primal solution is such that $z(i, s) = 0$ for $i < t^*(s)$ and the primal (i, s) -constraint is tight for $i \geq t^*(s)$.

This means in particular that for $i < t^*(s)$ the primal (i, s) constraint is slack, so by complementary slackness we must have $u(i, s) = 0$. For $i \geq t^*(s)$ we have $z(i, s) > 0$ which means that the (i, s) -constraint needs to be tight (again by complementary slackness). It means in particular that for $i \geq t(s)$ we must have $u(i, s) = a(i, s) - \sum_{j>i} \sum_{s'} u(j, s) \cdot c(j, s', i, s)$.

Now we argue that $u(i, s)$ must be equal to the solution obtained by backwards induction in equation (8). If not, let (i, s) be index with largest i such they differ. Note that both the dual solution u and the solution obtained by backwards induction are either zero or tight in each coordinate. If $a(i, s) - \sum_{j>i} \sum_{s'} u(j, s) \cdot c(j, s', i, s) > 0$ then neither can be zero (because of feasibility) so both must be tight. If $a(i, s) - \sum_{j>i} \sum_{s'} u(j, s) \cdot c(j, s', i, s) \leq 0$ then both must be zero. So they must also agree on (i, s) . $\square$

Finally we remark that there exist settings that satisfies HI and NF, but where the optimal policy is not a threshold policy (for an example please refer to the journal version).

5 APPLICATION: SECRETARY WITH SAMPLES

As our first case study we consider the secretary problem with samples. For concreteness, we focus on the model of Kaplan et al. [21] (Example 4). Recall that in this model an adversary picks $n + k$ numbers. Then k of these numbers chosen uniformly at random are shown to the algorithm as samples. Afterwards, the remaining n values are presented in random order.

LP Formulation. As signals we consider: s = “how many elements in the sample set are smaller than the current element”. So the signal space is in $\mathcal{S} = \{0, 1, 2, \dots, k\}$.

Now we need to calculate $a(i, s)$ and $c(i, s, j, s')$ and argue that both NF and HI are satisfied. The way to do that is based on the following way to sample:

- Start with a total order on k elements. Let's imagine that those are blue.
- Now we will add n black elements one by one in the following way:
For $t = 1, \dots, n$, there are $k + t$ positions where we can insert the t -th first black element. Insert in a random position.

This is the same sampling process as sampling the entire permutation and marking k of those blue and revealing the other ones in random order. The sampling procedure decides $\leq_{1..j}, s_1, \dots, s_j$ before deciding the position of later elements $i > j$ implying NF. Similarly, since the remaining elements are inserted after j , the relative order and signals of elements before j is irrelevant to whether j will be the top element conditioned on T_j, s_j , therefore verifying HI. Now, from this sampling process we can compute the coefficients:

$$a(i, s) = \mathbb{P}(r_i = n, s_i = s) = \mathbb{P}(s_i = s) \cdot \mathbb{P}(r_i = n \mid s_i = s) = \frac{1}{k+1} \cdot \prod_{t=1}^{n-1} \frac{s+t}{k+t+1}.$$

To see that imagine inserting first element i (it must be in one of the $k+1$ positions) and then inserting every other element and see what positions are available so that it is still the top so far.

For $c(i, s, j, s') = \mathbb{P}(T_j, s_j = s' \mid T_i, s_i = s)$ where T_i means that i is the top so far. Note that if $j \leq i$ this is only non-zero when $s' \leq s$. Subject to that we can again imagine inserting the elements one by one. We get:

$$c(i, s, j, s') = \frac{1}{s+1} \cdot \prod_{t=1}^{j-1} \frac{s'+t}{s+t+1}$$

The idea is again similar. Start by inserting j and then insert all the elements between 1 and $j-1$.

Plugging these formulas for the coefficients $a(i, s)$ and $c(i, s, j, s')$ into our generic LP, we obtain the following primal-dual pair:

$$\begin{array}{ll} \max \sum_i \sum_s z(i, s) \cdot \frac{1}{k+1} \cdot \prod_{t=1}^{n-1} \frac{s+t}{k+t+1} \text{ s.t.} & \min \sum_i \sum_s u(i, s) \text{ s.t.} \\ z(i, s) \leq 1 - \sum_{j < i} \sum_{s' \leq s} z(j, s') \cdot \frac{1}{s+1} \cdot \prod_{t=1}^{j-1} \frac{s'+t}{s+t+1} & u(i, s) + \sum_{j > i} \sum_{s' \geq s} u(j, s') \cdot \frac{1}{s'+1} \cdot \prod_{t=1}^{i-1} \frac{s+t}{s'+t+1} \\ & \geq \frac{1}{k+1} \cdot \prod_{t=1}^{n-1} \frac{s+t}{k+t+1} \\ z(i, s) \geq 0 & u(i, s) \geq 0 \end{array}$$

See Figure 2 for plots of the success probability of the optimal policy for $n = 20$ and varying k once as a function of k and once as a function of $p = k/(k+n)$.

Optimality of Thresholding. Before we show optimality of thresholding, we establish the following lemma that will allow us to work with expressions arising in the dual (which can be thought as a form of discrete integration):

LEMMA 13. *Given integers $a < b$ and $u < v$ the following equality holds:*

$$\sum_{s=u}^v \prod_{t=a+1}^b (s+t) = \frac{1}{b-a+1} \left[\prod_{t=a}^b (v+1+t) - \prod_{t=a}^b (u+t) \right]$$

PROOF. For some integer s we have:

$$\prod_{t=a}^b (s+1+t) - \prod_{t=a}^b (s+t) = [(s+b+1) - (s+a)] \prod_{t=a+1}^b (s+t) = (b-a+1) \prod_{t=a+1}^b (s+t)$$

The result is then obtained by telescoping the above equality for $s = u$ to v . $\square$

We can now prove:

THEOREM 14. *For every n and every k the optimal policy for secretaries with samples is a threshold policy.*

PROOF SKETCH (FULL PROOF IN THE JOURNAL VERSION). We prove the theorem through Proposition 11. We need to show that the greedy solution in equation (8) is optimal and non-decreasing in i . Monotonicity follows from observing that $a(i, s) = \frac{1}{k+1} \cdot \prod_{t=1}^{n-1} \frac{s+t}{k+t+1}$ does not depend on i and $c(j, s', i, s) = \frac{1}{s'+1} \cdot \prod_{t=1}^{i-1} \frac{s+t}{s'+t+1}$ is non-increasing in i .

For optimality assume that the optimal solution $u(i, s)$ does not satisfy the recursion that defines the greedy algorithm. Then there must be a largest index i^* and a signal s^* for which the equation is violated. Now we consider changing $u(i, s)$ to $u(i, s) + \Delta(i, s)$ where $\Delta(i^*, s^*) = -\delta$ for

$$\delta := u(i^*, s^*) - \max \left\{ 0, a(i^*, s^*) - \sum_{j>i^*} \sum_{s' \geq s} u(j, s') \cdot c(j, s', i^*, s^*) \right\}$$

and $\Delta(i, s) = 0$ for all other $i \geq i^*$. For $i < i^*$ we define $\Delta(i, s)$ recursively as follows

$$\Delta(i, s) = - \sum_{j=i+1}^{i^*} \sum_{s' \geq s} \Delta(j, s') \cdot c(j, s', i, s).$$

Furthermore, define

$$R_i := \sum_{j \geq i} \sum_{s \leq s^*} \Delta(j, s)$$

and note that R_i is the cumulative change to the objective function. We can now prove by backward induction that (see full proof in the journal version for the details):

(i) For $i = i^*$: $\Delta(i^*, s^*) = -\delta$ for $s = s^*$, and $\Delta(i^*, s) = 0$ for all $s \neq s^*$.

For $i < i^*$: $\Delta(i, s) = \delta \cdot \frac{1}{s^*+1} \prod_{t=1}^{i^*-2} \frac{s+t}{s^*+t+1}$ for all $s \leq s^*$, and $\Delta(i, s) = 0$ for all $s > s^*$.

(ii) For all $i \leq i^*$: $R_i = -\delta \cdot \frac{i-1}{i^*-1} \leq 0$.

In fact, by the definition of R_i and (i), we have

$$\begin{aligned} R_i &= \left(\sum_{j=i}^{i^*-1} \sum_{s=0}^{s^*} \Delta(j, s) \right) - \delta = \delta \cdot \left(\sum_{j=i}^{i^*-1} \sum_{s=0}^{s^*} \frac{1}{s^* + 1} \prod_{t=1}^{i^*-2} \frac{s+t}{s^* + t + 1} \right) - \delta \\ &= \delta \cdot \sum_{j=i}^{i^*-1} \left(\frac{\sum_{s=0}^{s^*} \prod_{t=1}^{i^*-2} (s+t)}{\prod_{t=1}^{i^*-1} (s^* + t)} \right) - \delta = \left(\frac{i^* - i}{i^* - 1} - 1 \right) \delta = -\delta \cdot \frac{i-1}{i^*-1} \leq 0, \end{aligned}$$

where the last equality follows from Lemma 13.

We now use (i) and (ii) to argue that the operation preserves feasibility and only improves the objective: We have chosen the recursion for the $\Delta(i, s)$ to satisfy the first constraint of the dual. From (i) we get that all the $\Delta(i, s)$ are non-negative, so we also satisfy the non-negativity constraints. From (ii) we get that $R_i \leq 0$ so we only decrease the objective. By repeatedly applying this operation we can conclude that the greedy solution is an optimal solution. $\square$

An immediate implication of Theorem 14 is an efficient (poly-time) algorithm for computing the optimal policy, and the winning probability of that policy.

Explicit solution to dual. Next we derive an explicit (non-recursive) formula for the optimal dual solution (Proof in the journal version).

THEOREM 15. *The following is an explicit solution to the dual recursion:*

For $i = n$ and all s :

$$u(n, s) = \frac{1}{k+1} \cdot \prod_{t=1}^{n-1} \frac{s+t}{k+t+1}.$$

For $i < n$ and all s :

$$u(i, s) = \max \left\{ 0, \frac{1}{k+1} \cdot \prod_{t=1}^{i-1} \left(\frac{s+t}{k+t+1} \right) \cdot \left(\prod_{t=i}^{n-1} \left(\frac{s+t}{k+t+1} \right) - \frac{\sum_{j=1}^{n-i} \left(\binom{n-i}{j} \cdot \frac{1}{j} \cdot \prod_{\ell=1}^j (k-s+\ell) \cdot \prod_{\ell=1}^{n-i-j} (s+i-1+\ell) \right)}{\prod_{t=i}^{n-1} (k+t+1)} \right) \right\},$$

where we use the convention that $\prod_{t=a}^b x_t = 1$ if $b < a$.

The dual solution has a natural interpretation: The positive term is the probability that secretary i with signal s is the best over all. This is the winning probability if we would accept. The negative terms are the probability with which we would win if we would pick the first secretary among the remaining $n-i$ secretaries that is better than the current one.

Gilbert and Mosteller as the limit when $k \rightarrow \infty$. In the model of Gilbert and Mosteller (Example 2) each secretary is associated with a sample from a known distribution F . The distribution can without loss of generality be thought as the uniform distribution over $[0, 1]$ since the algorithm can always process the quantiles $q_i = F^{-1}(s_i) \sim U[0, 1]$. Quantiles can be seen as a limit of the secretary with samples model with $k \rightarrow \infty$ by taking $q = s/(k+1)$ where $s \in \{0, \dots, k+1\}$ and take the limit $k \rightarrow \infty$. In the limit, $q \in [0, 1]$ corresponds to the quantile of the secretary. One can then take the limit of the LP coefficients:

$$a(i, q) = q^{n-1} dq \quad \text{and} \quad c(i, q, j, q') = \frac{1}{q} \left(\frac{q'}{q} \right)^{j-1} dq'$$

and obtain the following functional optimization problem in the limit:

$$\begin{array}{l|l} \max \sum_i \int_s z(i, q) \cdot q^{n-1} dq \text{ s.t.} & \min \sum_i \int_s u(i, q) dq \text{ s.t.} \\ z(i, q) \leq 1 - & u(i, q) + \\ \sum_{j < i} \int_{q' \leq q} z(j, q') \cdot \frac{1}{q} \cdot \left(\frac{q'}{q}\right)^{j-1} dq' & \sum_{j > i} \int_{q' \geq q} u(j, q') \cdot \frac{1}{q'} \cdot \left(\frac{q}{q'}\right)^{i-1} dq' \geq q^{n-1} \\ z(i, q) \geq 0 & u(i, q) \geq 0 \end{array}$$

Since this is a limit of secretaries with samples, we can again obtain a dual solution via backwards induction, which has a particularly nice form:

COROLLARY 16 (FROM THEOREM 14). *The optimal solution $u(i, s)$ to the dual LP satisfies:*

$$\begin{aligned} u(i, s) &= s^{n-1} && \text{for } i = n \text{ and all } s, \text{ and} \\ u(i, s) &= \max \left\{ 0, s^{i-1} \cdot \left(s^{n-i} - \sum_{k=1}^{n-i} \frac{1}{k} \binom{n-i}{k} s^{n-i-k} (1-s)^k \right) \right\} && \text{for } i < n \text{ and all } s. \end{aligned}$$

An immediate consequence of Corollary 16 is that the optimal policy can be determined by setting $s_n^* = 0$ and finding for each $i < n$ the s_i^* such that

$$(s_i^*)^{n-i} - \sum_{k=1}^{n-i} \frac{1}{k} \binom{n-i}{k} (s_i^*)^{n-i-k} (1 - (s_i^*))^k = 0, \quad (9)$$

and to then accept secretary i with signal s if it is the best so far and $s \geq s_i^*$. We note that this is precisely how Gilbert and Mosteller [16] define the optimal policy (Equation (3b-1) on p. 53).

Winning Probability. More generally, the explicit dual solution (in Theorem 15 and Corollary 16) allows to deduce thresholds in signal space (as we just did for Gilbert and Mosteller), both in the non-asymptotic and asymptotic regimes. For Gilbert and Mosteller, for example, it is known that solutions to equation (9) in the asymptotic regime satisfy $b_i \rightarrow 1$ and $i(1 - b_i) \rightarrow c$, where b_i for $i = n - 1, \dots, 0$ is the threshold when there are i secretaries after the current one, and this can be used to obtain analytic expressions for the asymptotic winning probability [17, 33].

6 APPLICATION: ADVICE FROM A BINARY CLASSIFIER

We will now use the technology developed in the previous sections to derive the optimal policy for the secretary problem with advice from a binary classifier (from the intro, and defined formally in Section 2). Recall that in this example secretaries have uniform random ranks, and that for the top secretary (rank = n) we receive advice Y with probability p and advice N with probability $1 - p$. For all other secretaries we receive advice N with probability p' and advice Y with probability $1 - p'$. In ML speak, p corresponds to the precision of the advice and p' to its specificity.

Since Example 5 satisfies NF and HI we can instantiate our general linear programming formulation to the binary classifier case. The LP and its dual are given in the journal version. We can then use it to show:

THEOREM 17. *The optimal policy for the binary advice problem is a threshold policy.*

The Optimal Threshold Policy. Next we use the fact that the optimal policy is a threshold policy to obtain a closed-form understanding of the optimal policy and the performance it achieves.

Let t_N and t_Y denote the thresholds (in time) after which we start accepting the best-ranking secretary so far conditioned on the advice being No (N) or Yes (Y). Write $\text{Opt}(n, p, p')$ for the optimal solution to the primal LP, and let $\text{OPT}(p, p') = \inf_{n \geq 1} \text{Opt}(n, p, p')$.

THEOREM 18. *With advice from a binary classifier with recall p and specificity p' , the optimal policy is a threshold policy that has two thresholds $t_Y \leq t_N$ from which onwards it accepts the best secretary so far upon receiving advice Y or N , and the optimal objective is:*

$$\text{Opt}(p, p') \geq \lim_{n \rightarrow \infty} \text{Opt}(n, p, p') = \left(\frac{1 - p'}{p} \right)^{(1-p')/p} \cdot e^{p/p'-1/p'}.$$

The two thresholds are $t_Y = n \cdot \left(\frac{1-p'}{p} \right)^{1/p'} \cdot e^{p/p'-1/p'}$ and $t_N = n \cdot e^{p/p'-1/p'}$.

Note how with $p = p' = 1/2$ (no advice), both thresholds become $t_Y = t_N = n/e$ recovering Dynkin's algorithm. If $p = p' = 1$ (perfect advice) then the thresholds become $t_Y = 0$ and $t_N = n$ which means the policy should always pick whenever the signal is Y and never pick if the signal is N .

The proof of Theorem 18 proceeds by solving the dual LP, we defer it to the journal version.

7 DISCUSSION AND FUTURE WORK

In this work we present a unifying view of multiple versions of the classic secretary problem. The linear program that is at the heart of our analysis precisely isolates the effect of the signalling scheme from the combinatorial structure of the problem: the joint distribution on ranks and signals only affects the coefficients of the simple set of constraints. The formulations that are captured by our analysis are quite diverse, as we demonstrate through the examples in Section 2, and it is remarkable that all of these can be solved optimally by memoryless policies.

As we saw, however, some settings require simple, but not memoryless policies. One possible direction for future work is to explore *counting* policies, that do not rely on the permutation of signals observed, but only on their histograms. Another approach is to consider advice that does not satisfy NF or HI, but still has sufficient structure to reason about its efficacy.

More broadly, we hope that the advice lens can be used to abstract the modeling assumptions and find formal connections between related problems in other areas.

REFERENCES

- [1] D. Alabi, A. T. Kalai, K. Ligett, C. Musco, C. Tzamos, and E. Vitercik. Learning to prune: Speeding up repeated computations. In *Proceedings of the 2019 Conference on Learning Theory*, pages 30–33, 2019.
- [2] P. Allaart. A general ‘bang-bang’ principle for predicting the maximum of a random walk. *Journal of Applied Probability*, 47(4):1072–1083, 2010.
- [3] A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. In *Proceedings of the 2020 Conference on Neural Information Processing*, 2020.
- [4] P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1377, 2014.
- [5] M. Balcan, T. Dick, T. Sandholm, and E. Vitercik. Learning to branch. In *Proceedings of the 35th International Conference on Machine Learning*, pages 353–362, 2018.
- [6] B. A. Berezovskiy and A. V. Gnedin. The best choice problem (in russian). *Nauka, Moscow*, 1984.
- [7] N. Buchbinder, K. Jain, and M. Singh. Secretary problems via linear programming. *Mathematics of Operations Research*, 39(1):190–206, 2014.
- [8] T. H. Chan, F. Chen, and S. H.-C. Jiang. Revealing optimal thresholds for generalized secretary problem via continuous lp: Impacts on online k-item auction and bipartite k-matching with random arrival order. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 1169–1188, 2015.
- [9] J. Correa, A. Cristi, L. Feuilloley, T. Osterwijk, and A. Tsigonias-Dimitriadis. The secretary problem with independent sampling. In *Proceedings of the 31st SIAM Symposium on Discrete Algorithms*, 2020. Forthcoming.

- [10] J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto. The two-sided game of googol and sample-based prophet inequalities. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms*, pages 2066–2081, 2020.
- [11] J. R. Correa, P. Dütting, F. A. 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.
- [12] J. Du Toit and G. Peskir. Selling a stock at the ultimate maximum. *The Annals of Applied Probability*, 19:983–1014, 2009.
- [13] E. B. Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics*, 4:627–629, 1963.
- [14] H. Esfandiari, M. Hajiaghayi, B. Lucier, and M. Mitzenmacher. Prophets, secretaries, and maximizing the probability of choosing the best. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 3717–3727, 2020.
- [15] T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4:282–289, 1989.
- [16] J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.
- [17] A. V. Gnedin. On the full information best-choice problem. *Journal of Applied Probability*, 33:678–87, 1996.
- [18] M. Hlynka and J. Sheahan. The secretary problem for a random walk. *Stochastic Processes and their Applications*, 28(2):317 – 325, 1988.
- [19] P. Indyk, A. Vakilian, and Y. Yuan. Learning-based low-rank approximations. In *Proceedings of the 2019 Conference on Neural Information Processing Systems*, pages 7400–7410, 2019.
- [20] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *Journal of the ACM*, 50(6):795–824, 2003.
- [21] H. Kaplan, D. Naori, and D. Raz. Competitive analysis with a sample and the secretary problem. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms*, pages 2082–2095, 2020.
- [22] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, pages 489–504, 2018.
- [23] S. Lattanzi, T. Lavastida, B. Moseley, and S. Vassilvitskii. Online scheduling via learned weights. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 1859–1877, 2020.
- [24] T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *Proceedings of the 35th International Conference on Machine Learning*, pages 3302–3311, 2018.
- [25] M. Mahdian and Q. Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 597–606, 2011.
- [26] A. M. Medina and S. Vassilvitskii. Revenue optimization with approximate bid predictions. In *Proceedings of the 2017 Annual Conference on Neural Information Processing Systems*, pages 1858–1866, 2017.
- [27] A. Mehta, A. Saberi, U. U. Vazirani, and V. V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22–40, 2007.
- [28] M. Mitzenmacher. A model for learned bloom filters and optimizing by sandwiching. In *Proceedings of the 2018 Annual Conference on Neural Information Processing Systems*, pages 462–471, 2018.
- [29] M. Mitzenmacher. Queues with small advice. *CoRR*, abs/2006.15463, 2020.
- [30] M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In *Proceedings of the 2018 Annual Conference on Neural Information Processing Systems*, pages 9684–9693, 2018.
- [31] D. Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1834–1845, 2020.
- [32] 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*, volume 151, pages 60:1–60:10, 2020.
- [33] S. Samuels. Exact solutions for the full information best choice problem. *Purdue University Statistics Mimeo Series*, page 81-87, 1982.
- [34] S. C. P. Yam, S. P. Yung, and W. Zhou. Two rationales behind the ‘buy-and-hold or sell-at-once’ strategy. *Journal of Applied Probability*, 46(3):651–668, 2009.