Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

On Hiring Secretaries with Stochastic Departures

Thomas Kesselheim, Alexandros Psomas, Shai Vardi

informs.
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–6
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

On Hiring Secretaries with Stochastic Departures

Thomas Kesselheim,[a] Alexandros Psomas,[b] Shai Vardi[c,*]

[a] Institute of Computer Science, University of Bonn, 53115 Bonn, Germany; [b] Department of Computer Science, Purdue University, West Lafayette, Indiana 47907; [c] Krannert School of Management, Purdue University, West Lafayette, Indiana 47906

*Corresponding author

Contact: thomas.kesselheim@uni-bonn.de, https://orcid.org/0000-0002-9420-9424 (TK); apsomas@cs.purdue.edu (AP); svardi@purdue.edu, https://orcid.org/0000-0003-4720-6826 (SV)

Received: June 2, 2020
Revised: January 17, 2022; October 25, 2022
Accepted: March 16, 2023
Published Online in Articles in Advance: May 10, 2023

Area of Review: Optimization

https://doi.org/10.1287/opre.2023.2476

Abstract. We study a generalization of the secretary problem, where decisions do not have to be made immediately upon applicants' arrivals. After arriving, each applicant stays in the system for some (random) amount of time and then leaves, whereupon the algorithm has to decide irrevocably whether to select this applicant or not. The arrival and waiting times are drawn from known distributions, and the decision maker's goal is to maximize the probability of selecting the best applicant overall. Our first main result is a characterization of the optimal policy for this setting. We show that when deciding whether to select an applicant, it suffices to know only the time and the number of applicants that have arrived so far. Furthermore, the policy is monotone nondecreasing in the number of applicants seen so far, and, under certain natural conditions, monotone nonincreasing in time. Our second main result is that when the number of applicants is large, a single threshold policy is almost optimal.

Funding: A. Psomas is supported in part by the National Science Foundation [Grant CCF-2144208], a Google Research Scholar Award, and by the Algorand Centres of Excellence program managed by Algorand Foundation.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2023.2476.

Keywords: online algorithms • secretary problem • stopping problems

1. Introduction

The secretary problem is an archetypal stopping problem whose origin is still being debated, but is usually attributed to Gardner (1960), Lindley (1961), or Dynkin (1963). In its classical version, n applicants are interviewed in uniformly random order. Only one applicant can be hired, and the goal is to maximize the probability of hiring the best one. After each interview, the interviewer must make an immediate and irrevocable decision of whether to hire the applicant. Despite (or because of) its simplicity, the secretary problem has become the foundation of a host of theoretical decision-making models, including those for buying or selling an asset, finding an apartment (Zwick et al. 2003), selling advertisements (or other items) in online auctions (Hajiaghayi et al. 2004), and, naturally, hiring employees. Despite its wide-ranging appeal, the classical problem includes several assumptions that reduce the applicability of its solution to the problems it models: a decision has to be made immediately; the firm obtains no value from hiring any applicant that is not the best one; the number of applicants is known a priori; only one applicant can be hired.

In this paper, we address the assumption that a decision has to be made immediately. In the majority of scenarios to which the secretary problem is applied, the applicants *do not* instantaneously depart; nevertheless, most of the previous work on the secretary problem and its variations does not depart from its truly sequential nature: an irrevocable decision has to be made *immediately* upon seeing the applicant. In many situations of interest, however, it is reasonable to assume that the decision does not have to be immediate. When interviewing prospective employees, it is reasonable to expect that the candidates will not be hired by another employer immediately upon leaving the interview; when selling a commodity, potential buyers do not typically make instantaneous take-it-or-leave it offers. The assumption that decisions must be made instantaneously (or a similar assumption that there is no recall) is often explicitly mentioned as a weakness of the classical model (e.g., Freeman 1983, Goldstein et al. 2019). In other papers (e.g., Bearden et al. 2006), the space of real-world situations to which the results apply is severely restricted as a result of this assumption.[1] Relaxing this assumption is an important step toward expanding the space of problems to which many results related to the secretary problem apply. In the extant literature, only simple variants have been addressed, the main one being the "sliding window" variant, where a decision needs to be made only after some fixed number of other applicants has arrived (e.g., Goldys 1978, Ho and Krishnan 2015). This model is unrealistic for many applications. For example, a bidder on a house would not typically withdraw the bid after a fixed number of other bidders arrived; the amount of time before the bid is withdrawn is arguably better modeled using a distribution.

Before describing our model and results, we summarize the optimal solution for the classical problem, as it will serve as a basis for comparison. Call any applicant that is the best out of all the applicants seen thus far the (current) *candidate*. Interview some fixed number of applicants, without accepting any, and thereafter accept the first applicant that arrives that is a candidate. The precise number of applicants that are only observed is determined by backward induction (Lindley 1961, Gilbert and Mosteller 1966) and is asymptotically n/e; the probability of hiring the best applicant is asymptotically $1/e$. The optimal strategy is attractive because of its simplicity: the decision of whether to hire a candidate does not depend on the complete history; it depends only on the number of applicants interviewed so far.

To relax the constraint that the decision of whether to hire each applicant must be immediate, we consider the scenario in which the applicants arrive according to some distribution, and remain in the system for an amount of time drawn from some other distribution. Observe that, in the classical problem, the precise arrival distribution is irrelevant (as long as it does not have a point mass), as the only parameter of the decision (other than whether the applicant is the best-so-far) is the number of applicants that have arrived. We show that, in contrast, in the stochastic departure case, the arrival distribution affects the optimal strategy. Our main results are as follows.

1. The optimal policy has the following characteristics:
   (a) In order to decide whether to hire a candidate, the optimal strategy needs to take into account only two parameters: the number of applicants that have arrived and the time that has passed. We also show that this is unavoidable; that is, the optimal policy *must* consider both parameters.
   (b) The optimal policy is monotone nondecreasing in the number of applicants, in the following sense: for any fixed time, there is some threshold such that if the number of applicants that have arrived is above the threshold, then the policy accepts; otherwise, it rejects.
   (c) In the case that the arrival distribution is uniform, the optimal policy is monotone *nonincreasing* in the time: for any fixed number of applicants that have arrived, the policy is *less* likely to accept if more time has passed. We also show this is not necessarily the case for general arrival distributions.
2. When the arrival distribution is continuous and the number of applicants is large, the decisions do not have to depend on the number of applicants: a single-threshold policy achieves an asymptotically optimal success probability.

1.1. Overview of Proofs and Techniques

When facing a decision, the optimal policy accepts a candidate if and only if this gives a higher success probability than rejecting. The success probability after a rejection again depends on the policy. As a result, the success probabilities are frequently bounded in different conditional probability spaces. In order to compare these success probabilities, we use simulation arguments and coupling of the probability spaces. One such argument is used to show result 1(a): If the result is not true, then there must be two histories, h and h', with the same t (time) and k (number of applicants) on which the optimal policy makes different decisions. We then define a new policy that follows the decisions of the optimal policy for h whenever it encounters h'. We couple the future events to show that our new policy does better than the optimal policy on h', a contradiction.

Some arguments require a more detailed look at the probability spaces generating the events. For example, to show result 1(b), we would like to argue that after having seen k applicants at time t, if the success probability for accepting the current candidate is higher than for rejecting, then the same holds having seen $k+1$ applicants at time t. In this case, the simulating policy acts as if $k+1$ applicants have already arrived, and only $n-k-1$ applicants remain. It does so by ignoring a random future arrival. The difficulty comes from the fact that future observations conditioned on having seen k applicants are not identically distributed to the future observations conditioned on having seen $k+1$ applicants. We overcome this by coupling with a suitably chosen conditional probability space to get an upper bound on the probability of a future observation.

To prove result 2, we explicitly construct a threshold policy. At any time t, the policy compares two probabilities: (i) the probability that the overall best applicant arrives by t, and (ii) the success probability of an optimal policy that only accepts applicants arriving after t. Whenever (i) is greater than (ii), accept the candidate. Note that, by definition, (i) is nondecreasing and (ii) is nonincreasing in t, so this policy is a threshold policy. To prove asymptotic optimality, we observe that the optimal policy actually compares the same probabilities but in the probability space conditioned on the number of applicants that have arrived and the candidate at time t leaving at t. Using concentration bounds, we show that the conditional probabilities are in most cases close to the unconditional ones, and the newly constructed policy differs from the optimal one only when accepting and rejecting yield similar success probabilities. To bound the loss in success probability by these errors, we compare conditional and unconditional probabilities using coupling arguments. The key difference from the proofs of the first three results is that we care not only about which policy has a higher success probability but also by how much they differ.

1.2. Related Work

The secretary problem and its variants have received much attention since the later part of the 20th century. We refer the reader to Freeman (1983) and Ferguson (1989) for surveys on the classical literature on secretary problems and variations thereof.

In this work, we address a major criticism of the classical secretary problem: that decisions have to be made immediately. Many other criticisms of the model have been addressed in the literature. Bearden et al. (2006) consider variants in which the firm does not obtain value only from hiring the best applicant, but instead receives a payoff that increases with the quality of the applicant. In a follow-up work, Palley and Kremer (2014) compare the situations in which the decision maker only has access to pairwise comparisons between the applicants, and when they observe the actual value. Goldstein et al. (2019) consider a variant in which the value of the applicant is sampled from some distribution, and this distribution can be learned as the game is repeated. Cownden and Steinsaltz (2014) study the scenario where multiple players seek to employ applicants from a common labor pool. Alpern and Baston (2017) consider the situation in which applicants are interviewed by a search committee and may have different values to different members of the committee.

Special cases of not making an immediate decision have been addressed in the literature: Goldys (1978) showed that the expected rank of the accepted applicant tends to ≈2.57 as n tends to infinity, when one is allowed to choose either the current applicant or the previous one. This is in contrast to the expected rank of ≈3.87, when one is only allowed to choose the current applicant, shown by Chow et al. (1964). In this setting, the value of an applicant is $n-i$ when the ith best applicant is accepted, as opposed to 1 if the best applicant is accepted and 0 otherwise in the classical setting. The scenario considered by Goldys is sometimes called a *sliding window*. In the online setting, a sliding window of size x means that after seeing an applicant, the algorithm does not need to make a decision until it has seen x more applicants. Goldys's setting is a sliding window of size 1. Smith and Deely (1975), and much later and independently Ho and Krishnan (2015), consider sliding windows of length $1 < m < n$, and show that the optimal policy is a thresholding policy. Ho and Krishnan (2015) also give a recursive (nonexplicit) formula for the probability of hiring the best applicant using sliding windows of size n/i for constant i; and they give an asymptotic bound when the window size is at least $n/2$.

Vardi (2015) considered the scenario where applicants arrive k times each, and the arrival order is uniform over the $(kn)!$ possible arrival orders, and gave an optimal threshold-based strategy for accepting the best applicant and computed the success probability for $k=2$; Hoefer and Wilhelmi (2021) extended these results to some packing domains. Petruccelli (1981) considered the case when the interviewer is able to recall some applicant from the past with some probability $p > 0$.

We note that in all of the above cases—in contrast to our setting—the optimal policy is uni-variate, depending only on the number of observed applicants, and not on the time at which the decision is made.

2. Model and Preliminaries

A set $\mathcal{S}$ of n applicants arrive and depart over some time interval, which we normalize to [0, 1]. For each $i \in \mathcal{S}$, an *arrival time* a_i is drawn independently from a (known) arrival distribution $\mathcal{A}$ and a *waiting time* l_i (denoting how long the applicant stays in the system) is drawn independently and identically distributed (i.i.d.) from a (known) waiting time distribution $\mathcal{L}$. Applicant i leaves at time $d_i = \min\{a_i + l_i, 1\}$. There is a total order $\rho : [n] \to [n]$ on $\mathcal{S}$, and a decision maker wishes to select the best applicant; that is, $i \in [n]$ such that $\rho(i) = 1$. We call this problem the *stochastic departure secretary problem* and denote it by a triple $(n, \mathcal{A}, \mathcal{L})$.

An *instance* (which we will sometimes refer to as a *complete instance* for clarity) is represented by a triple $(\vec{a}, \vec{d}, \vec{\rho})$, where $|\vec{a}| = |\vec{d}| = |\vec{\rho}| = n$. $\vec{a} \in [0,1]^n$ is a vector of arrival times, $a_1 \le a_2 \le \cdots \le a_n$; $\vec{d} \in [0,1]^n$ is a vector of departure times, and $\vec{\rho}$ is the vector of the ranks of the applicants. Alternatively, we also denote an instance by a triple $(\vec{a}, \vec{d}, \vec{r})$, where $\vec{r}$ is a vector of *relative ranks* of the applicants. The relative rank of applicant i, $r_i \in \{1, \dots, i\}$, indicates the number of applicants among $1, \dots, i$ that are at least as good as i. Formally, $r_i = |\{j \le i | \rho(j) \le \rho(i)\}|$. This representation has the advantage that it matches the knowledge of the decision maker; after i arrivals, the decision maker knows $r_1, \dots, r_i$ but not how these applicants compare with future arrivals. At time t, we call applicant j where $j = \max\{i : a_i \le t, r_i = 1\}$ a *candidate*; at any time after the first arrival, there is a unique candidate. A history h is an event at which an applicant departs. It is denoted by a tuple $h = (\vec{a}, \vec{d}, \vec{r}, t)$, where $|\vec{a}| = |\vec{d}| = |\vec{r}| = n$; there is some $k \in \{1, \dots, n\}$ such that for $i \in \{k+1, \dots, n\}, a_i = \perp$, denoting that applicant i has not yet arrived by time t. Here, $\vec{a}$ is sorted: it holds that $0 \le a_1 \le \cdots \le a_k \le t$, for $i \le k$. For $i \in \{1, \dots, n\}$, $d_i \in [0, t] \cup \perp$, where $\perp$ indicates that the applicant has not departed by time t (clearly, $a_i = \perp \Rightarrow d_i = \perp$). Lastly, $r_i \in \{1, \dots, i\}$ for $i \le k$, $r_i = \perp$ for $i > k$.

Generally, at any point in time, an optimal policy makes a decision that maximizes the probability of success from this point onward. As ties can occur, the policy is not unique. Therefore, we only consider *lazy* policies. A policy is lazy if it rejects whenever acceptance and rejection have identical conditional success probabilities. There is always a lazy optimal policy and it is unique. Denote this policy by OPT. We first make two standard observations regarding the optimal lazy policy OPT (see, e.g., Gilbert and Mosteller 1966, Bruss 2000). Given a time t, let ACC_t be the policy that accepts only at time t and only in the event that the candidate at time t departs at exactly t. Let REJ_t be the policy that rejects all departing applicants up to and including time t and thereafter continues with the optimal policy (conditioned on having rejected all applicants up to this point). Given any policy POL, we denote by SUCCESS(POL) the event that POL selects the best applicant. The first observation follows immediately from the facts that any optimal policy makes decisions that maximize the probability of success and that OPT is lazy.

Observation 1. Given a time t, let h be a history until t in which the candidate departs at t. Then OPT accepts this candidate if and only if it has not accepted any applicant before and

$$\Pr[\text{SUCCESS}(\text{REJ}_t | h)] < \Pr[\text{SUCCESS}(\text{ACC}_t | h)].$$

Observation 2. Given a time t, let h be a history in which k applicants have arrived by t and the candidate departs at t. Then $\Pr[\text{SUCCESS}(\text{ACC}_t | h)] = \frac{k}{n}$.

To see why the second observation is true, note that, as the arrival order of applicants is chosen uniformly at random out of all permutations, the probability that the optimal is among the first k arrivals is exactly k/n.

3. Optimal Stopping Rule

In this section, we characterize the optimal stopping rule. We use the following definition.

Definition 1. A policy for the stochastic departure secretary problem is *bivariate* if its decision whether to accept or reject a candidate departing at time t given a history h, depends only on t and k, the number of applicants that have arrived up to t. In other words, there exists a function $\Theta_n(t,k)$,

$$\Theta_n : [0,1] \times [n] \to \{\text{ACCEPT}, \text{REJECT}\},$$

such that whenever a candidate x departs at time t and k applicants have arrived prior to t, the policy accepts if and only if x is the candidate and $\Theta_n(t,k) = \text{ACCEPT}$.

Theorem 1. *There exists a bivariate optimal policy $\Theta_n(t,k)$ for the stochastic departure secretary problem; Θ_n is nondecreasing in k for fixed t, and if the arrival distribution $\mathcal{A}$ is uniform, then it is nonincreasing in t for fixed k.*

Theorem 1 implies that, for uniform arrivals, it is possible to succinctly[2] represent the optimal policy using n real numbers, $t_1, \ldots, t_n$: if at time t, k applicants have arrived and the best leaves, then accept if and only if $t \le t_k$. The following two examples show that both k and t are necessary; that is, Θ_n is indeed a function of both k and t, and not just of one of them.

Example 1 (The Optimal Policy Depends on k). It is trivial to confirm that for any time t, any $n > 2$, and any arrival and departure distributions, if $k = 1$ the optimal policy rejects, and if $k = n$, then the optimal policy accepts.

Example 2 (The Optimal Policy Depends on t). Let the arrival distribution be the uniform distribution, and assume that each applicant stays in the system for some very small fixed time $\epsilon \le 1/n^3$. If the number of applicants that have arrived by time $t = 4/9$ is $4n/9$, then the probability of success for accepting the candidate at this time is $4/9$ (by Observation 2). To bound the probability of success for rejecting, denote such a history by h, and denote the event that there will exist some time in $[t, 1]$ when there is more than one applicant in the system by $\mathbb{I}$. Then,

$$\begin{aligned}\Pr[\text{SUCCESS}(\text{REJ}_t | h)] &\le \Pr[\text{SUCCESS}(\text{REJ}_t | h, \mathbb{I})]\Pr[\mathbb{I} | h] \\ &\quad + \Pr[\text{SUCCESS}(\text{REJ}_t | h, \neg\mathbb{I})]\Pr[\neg\mathbb{I} | h] \\ &\le \Pr[\mathbb{I} | h] + \Pr[\text{SUCCESS}(\text{REJ}_t | h, \neg\mathbb{I})] \\ &\le 2/5.\end{aligned}$$

This is because $\Pr[\mathbb{I} | h] \le 1/n$ and $\Pr[\text{SUCCESS}(\text{REJ}_t | h, \neg\mathbb{I})]$ is exactly the probability of success in the classical secretary of rejection at time t conditioned on $4n/9$ applicants arriving by this time. The latter is at most the maximal probability of success for rejecting in the classical case, which is approximately $1/e$ (the exact value depends on the value of n). As the probability of success of accepting is greater than that of rejecting, we should accept. However, if at time $t = 1 - \epsilon$, $4n/9$ applicants have arrived, then we should reject: all the remaining applicants will arrive by time 1, and none of them will depart by this time. Therefore, the probability of success if we reject is equal to the probability that the best applicant is not one of the first $4n/9$ applicants (i.e., $5/9$), while the probability of accepting is $4/9$.

In order to prove Theorem 1, we prove three things: (a) the optimal policy depends only on t and k (not on the complete history at time t; (b) the optimal policy is monotone nondecreasing in k; and (c) the optimal policy is monotone nondecreasing in t for the uniform arrival distribution. The proofs of these results are given in the online appendix. The implicit requirement in (a)—that the applicants' arrival or waiting times are identically distributed—is necessary. If the applicants' arrival or waiting times are not identically distributed, then t and k might be insufficient to define the optimal policy might (i.e., (a) does not hold). This can be seen in the following examples.

Example 3. Let $n = 100$. Partition the applicants into two sets: A has 49 applicants, and these arrive uniformly at random in $[0, 1]$; B has 51 applicants that arrive at time 0 with probability $x \in (0, 1)$ and at time 0.51 with probability $1 - x$. All applicants stay in the system for some small $\epsilon \leq 10^{-6}$. If at time 0.5 all of A but none of B have arrived and the candidate departs, then we should reject: the probability of success for accepting and rejecting is 0.49 and 0.51, respectively. If, on the other hand, 49 applicants from B but none from A have arrived by this time and the candidate departs, then we should accept. The probability of success for accepting is 0.49; to bound the probability of success for rejecting, denote this history by h, and denote the event that there will exist some time in [0.5, 1] when there is more than one applicant from A in the system by $\mathbb{I}_1$, and the event that the best candidate is one of the remaining applicants in B by $\mathbb{I}_2$. Then,

$$\begin{aligned}\Pr[\text{SUCCESS}(\text{REJ}_t|h)] \leq \Pr[\mathbb{I}_1|h] + \Pr[\mathbb{I}_2|h] \\ + \Pr[\text{SUCCESS}(\text{REJ}_t|h, \neg\mathbb{I}_1, \neg\mathbb{I}_2)] \\ < 0.49,\end{aligned}$$

as $\Pr[\mathbb{I}_1|h] \leq 0.01$, $\Pr[\mathbb{I}_2|h] \leq 0.02$ and $\Pr[\text{SUCCESS}(\text{REJ}_t| h, \neg\mathbb{I})]$ is equal to the probability of success in the classical secretary of rejection at time t conditioned on 49 applicants arriving by this time, which is approximately $1/e$, and at most (say) 0.45.

Example 4. A similar example shows that the waiting time needs to be identically distributed. Here, the arrival distribution is the uniform distribution for all applicants, and again we partition the applicants into two sets. Set A has 49 applicants that depart immediately upon arrival, and set B has 51 applicants that stay in the system for 0.5. If all of A and none of B have arrived by time 0.5, then we should reject. If 49 applicants from B have arrived, then we should accept.

Theorem 1 gives a sufficient condition for the optimal policy to be monotone nonincreasing in t. Giving a necessary condition is left as an open problem, but we note that the optimal policy is not monotone nonincreasing in t for all arrival distributions, as the following example shows.

Example 5. Consider the following arrival and departure distributions: each applicant arrives in $[0, 4\epsilon]$ or $(6\epsilon, 1]$ with probability ϵ, and in $(4\epsilon, 6\epsilon]$ with probability $1 - 2\epsilon$. Conditioned on arriving in each interval, the arrival distribution is uniform over the interval. Applicants remain in the system for 3ϵ. Assume that n is large. If $4n/9$ applicants have arrived by time 4ϵ, and the candidate leaves, then we should reject, as $\Pr[\text{SUCCESS}(\text{ACC}_t)] = 4/9$, $\Pr[\text{SUCCESS}(\text{REJ}_t)] \approx 5/9$. However, if $t = 4/9$, then we should accept, as the circumstances are almost identical to those of the classical continuous time secretary model, where agents arrive uniformly at random in [0, 1], and we should accept if $t > \frac{1}{e}$ (see, e.g., Freeman 1983).

4. Optimal Policy in the Limit

We now show that, for large n, optimal policies have an even simpler structure. To get asymptotically optimal performance, it is enough to define a time threshold t^* such that, irrespective of the number of arrivals, we always accept a departing candidate if they depart at time $t \geq t^*$. To formally define this policy, we first make some observations regarding the conditional success probabilities that a policy can consider when making accept/reject decisions.

Let $\mathcal{E}_t$ denote the event that the candidate leaves at time t. Let REJ_t^n be the policy for n applicants that rejects all applicants that depart up to and including time t, and thereafter continues with the optimal policy. Let REJA_t^n be the best policy for n applicants out of those that reject all applicants that *arrive* up to and including time t. Whereas REJA_t^n and REJ_t^n are different, they have the same probability of success conditioned on $\mathcal{E}_t$ (from Theorem 1); that is,

$$\Pr[\text{SUCCESS}(\text{REJA}_t^n|\mathcal{E}_t)] = \Pr[\text{SUCCESS}(\text{REJ}_t^n|\mathcal{E}_t)].$$

This is because if $\mathcal{E}_t$ occurs, then neither policy accepts any applicant that arrived before. Define $p_{n,t} = \Pr[\text{SUCCESS}(\text{REJA}_t^n)]$. In the online appendix, we show that the limit $p_t = \lim_{n\to\infty} p_{n,t}$ exists for any fixed t. Dini's theorem implies that this convergence is uniform, because $p_{n,t}$ is continuous for all n. Therefore, $\lim_{n\to\infty}\sup_{t\in\mathbb{R}}|p_t - p_{n,t}| = 0$. As a consequence, the function $\pi : t \to p_t$ is nonincreasing and continuous.

The cumulative distribution function of the arrival distribution ($\mathcal{A}$) and π are continuous, have the same domain, and their ranges are [0, 1] and $[0, \zeta]$ for some $\zeta \in (0, 1)$, respectively. Therefore, they must intersect. That is, there is some $t^* \in [0, 1]$ such that $p_t \geq \mathcal{A}(t)$ for $t < t^*$ and $p_t \leq \mathcal{A}(t)$ for $t > t^*$. This allows us to define our policy, denoted by POL^*.

Definition 2. Let $t^* \in [0, 1]$ be the intersection of $\mathcal{A}$ and $\pi : t \to p_t$. We have POL^* as follows. For any $t \in [0, 1]$, if the candidate departs at time t, then $\text{POL}^*(t) = \text{REJECT}$ if $t \leq t^*$ and $\text{POL}^*(t) = \text{ACCEPT}$ if $t > t^*$.

Our main result in this section is that POL^* is asymptotically optimal. The proof is given in the online appendix.

Theorem 2. *Given any continuous arrival distribution $\mathcal{A}$ and arbitrary waiting time distribution $\mathcal{L}$, there exists a policy POL^* defined by a threshold t^* that accepts a candidate that leaves at time t if and only if $t > t^*$. For every $\epsilon > 0$, there is an n_0 such that for all $n > n_0$, $\Pr[\text{SUCCESS}(\text{POL}^*)] \geq \Pr[\text{SUCCESS}(\text{POL}_n)] - \epsilon$, where POL_n is the optimal policy for n applicants.*

5. Future Directions

This paper offers a structural analysis of the secretary problem with stochastic departures. However, many natural questions are more quantitative in nature: How does the optimal policy's probability of success vary with parameters of the model, such as the waiting time distribution $\mathcal{L}$? Under what conditions on $\mathcal{L}$ is the probability of success guaranteed to be bounded away from $1/e$ as n tends to infinity? For example, does this hold whenever $\Pr(\mathcal{L} > 0)$ is positive? If we hold the arrival time distribution $\mathcal{A}$ fixed but vary $\mathcal{L}$ to make applicants more patient (i.e., substitute $\mathcal{L}$ with another distribution $\mathcal{L}'$ that stochastically dominates $\mathcal{L}$), does this always increase the probability of success?

Another interesting research direction is relaxing the assumption that the system knows when an applicant is about to depart. In this work, we assume that the optimal policy receives a signal when each applicant departs and is allowed to make a decision thereafter. What structure does the optimal policy have when we only receive a signal immediately *after* an applicant's departure?

Finally, the stochastic departure aspect of our model can be applied to virtually all variations of the secretary problem. Of particular interest is the effect of stochastic departures on the matroid secretary problem (Babaioff et al. 2007), which have a strong connection with online auctions (Babaioff et al. 2008).

Acknowledgments

Part of the work on this paper was done while the authors were visiting the Simons Institute for the Theory of Computing. The authors would like to thank the anonymous journal and conference reviewers for their careful reading and useful comments.

Endnotes

[1] Bearden et al. (2006) motivate their results using "tight housing markets."

[2] To the extent that real numbers can be represented succinctly.

References

Alpern S, Baston V (2017) The secretary problem with a selection committee: Do conformist committees hire better secretaries? *Management Sci.* 63(4):1184–1197.

Babaioff M, Immorlica N, Kleinberg R (2007) Matroids, secretary problems, and online mechanisms. *Proc. 18th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 434–443.

Babaioff M, Immorlica N, Kempe D, Kleinberg R (2008) Online auctions and generalized secretary problems. *ACM SIGecom Exchanges* 7(2):1–11.

Bearden JN, Rapoport A, Murphy RO (2006) Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Sci.* 52(9):1437–1449.

Bruss FT (2000) Sum the odds to one and stop. *Ann. Probab.* 28:1384–1391.

Chow Y, Moriguti S, Robbins H, Samuels S (1964) Optimal selection based on relative rank (the "secretary problem"). *Israel J. Math.* 2(2):81–90.

Cownden D, Steinsaltz D (2014) Effects of competition in a secretary problem. *Oper. Res.* 62(1):104–113.

Dynkin EB (1963) The optimal choice of the instant for stopping a Markov process. *Soviet Math. Doklady* 4:627–629.

Ferguson TS (1989) Who solved the secretary problem? *Statist. Sci.* 4:282–296.

Freeman P (1983) The secretary problem and its extensions: A review. *Internat. Statist. Rev.* 51(2):189–206.

Gardner M (1960) Mathematical games. *Sci. Amer.* 202(152):178–179.

Gilbert JP, Mosteller F (1966) Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.* 61(313):35–73.

Goldstein DG, McAfee RP, Suri S, Wright JR (2019) Learning when to stop searching. *Management Sci.* 66(3):1375–1394.

Goldys B (1978) The secretary problem: The case with memory for one step. *Demonstratio Math.* 11:789–799.

Hajiaghayi MT, Kleinberg RD, Parkes DC (2004) Adaptive limited-supply online auctions. *Proc. Fifth ACM Conf. Electronic Commerce (EC)* (ACM, New York), 71–80.

Ho S, Krishnan A (2015) A secretary problem with a sliding window for recalling applicants. Preprint, submitted August 31, https://doi.org/10.48550/arXiv.1508.07931.

Hoefer M, Wilhelmi L (2021) Packing returning secretaries. *Networks* 77:454–471.

Lindley D (1961) Dynamic programming and decision theory. *J. Royal Statist. Soc. Ser. C Appl. Statist.* 10:39–52.

Palley AB, Kremer M (2014) Sequential search and learning from rank feedback: Theory and experimental evidence. *Management Sci.* 60(10):2525–2542.

Petruccelli JD (1981) Best-choice problems involving uncertainty of selection and recall of observations. *J. Appl. Probab.* 18(2):415–425.

Smith MH, Deely JJ (1975) A secretary problem with finite memory. *J. Amer. Statist. Assoc.* 70(350):357–361.

Vardi S (2015) The returning secretary. *32nd International Symposium on Theoretical Aspects of Computer Science (STACS 2015)* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany), 716–729.

Zwick R, Rapoport A, Lo AKC, Muthukrishnan AV (2003) Consumer sequential search: Not enough or too much? *Marketing Sci.* 22(4):503–519.

Thomas Kesselheim is a professor in the Institute of Computer Science at the University of Bonn. His research considers problems of approximation algorithms, online algorithms, and algorithmic game theory.

Alexandros Psomas is an assistant professor of computer science at Purdue University. He is interested in theoretical computer science, working on a broad set of problems at the intersection of computer science and economics.

Shai Vardi is an assistant professor of management at Purdue University. His research spans mechanism design, online platforms, and dynamic allocation problems.