

On a Competitive Secretary Problem with Deferred Selections

Tomer Ezra¹, Michal Feldman^{1,2}, Ron Kupfer³

¹Tel Aviv University

²Microsoft Research

³The Hebrew University of Jerusalem

tomer.ezra@gmail.com, michal.feldman@cs.tau.ac.il, ron.kupfer@mail.huji.ac.il

Abstract

We study the secretary problem in multi-agent environments. In the standard secretary problem, a sequence of arbitrary awards arrive online, in a random order, and a single decision maker makes an immediate and irrevocable decision whether to accept each award upon its arrival. The requirement to make immediate decisions arises in many cases due to an implicit assumption regarding competition. Namely, if the decision maker does not take the offered award immediately, it will be taken by someone else. We introduce a novel multi-agent secretary model, in which the competition is explicit. In our model, multiple agents compete over the arriving awards, but the decisions need not be immediate; instead, agents may select previous awards as long as they are available (i.e., not taken by another agent). If an award is selected by multiple agents, ties are broken either randomly or according to a global ranking. This induces a multi-agent game in which the time of selection is not enforced by the rules of the games, rather it is an important component of the agent's strategy. We study the structure and performance of equilibria in this game. For random tie breaking, we characterize the equilibria of the game, and show that the expected social welfare in equilibrium is nearly optimal, despite competition among the agents. For ranked tie breaking, we give a full characterization of equilibria in the 3-agent game, and show that as the number of agents grows, the winning probability of every agent under non-immediate selections approaches her winning probability under immediate selections.

1 Introduction

In the classic secretary problem [Ferguson, 1989] a decision maker observes a sequence of n non-negative real-valued awards v_1, \dots, v_n , which are unknown in advance, in a random order. At time step t , the decision maker observes award v_t , and needs to make an immediate and irrevocable decision whether or not to accept it. If she accepts v_t , the game terminates with value v_t ; otherwise, the award v_t is gone forever

and the game continues to the next round. The objective of the decision maker is to maximize the probability of choosing the maximal award. A tight competitive ratio of $1/e$ is well known for this problem (see, e.g., [Ferguson, 1989]).

This problem (and variants thereof) is an abstraction that captures many real-life scenarios, such as an employer who interviews potential workers overtime, renters looking for a potential house, a person looking for a potential partner for life, PhD students looking for an advisor, and so on. This problem also has interesting implications to mechanism design, auctions and pricing, for both welfare and revenue maximization, in various markets such as online advertising markets (see e.g. [Babaioff *et al.*, 2008; Babaioff *et al.*, 2007; Ezra *et al.*, 2018; Freeman, 1983; Kesselheim *et al.*, 2013; Kleinberg, 2005]).

In this work, we study the secretary problem in a multi-agent system by introducing a secretary model with multiple agents who compete with each other.

1.1 Competing Agents

Competition among agents is a fundamental component in many real-life scenarios. A recent line of work studies secretary settings with multiple agents, where a set of employers compete over a set of potential employees. The employees enter the labor market sequentially, and the employers need to decide whether to hire them or not. When a potential employee receives an offer, she accepts it; if she receives multiple offers, she chooses one among them.

Immorlica *et al.* [2006] and Karlin and Lei [2015] consider settings with multiple decision makers who compete over awards that arrive online. In these studies, as in the standard setting, decisions are *immediate* and *irrevocable*. One of the justifications for the requirement in the standard (single agent) secretary setting that decisions must be immediate arises from an implicit assumption about competition. Namely, if the decision maker does not take the current offered award, then it may be taken by someone else and gone forever. For example, an employer who does not make a job offer to a potential employee following an interview will lose her forever, as she will probably be hired by another firm.

In this paper, we introduce a multi-agent model in which the competition among agents over the arriving awards is *explicit*. Since the competition is explicitly captured by the multiplicity of agents, decisions need no longer be immediate

mediate. *Instead, every previously arriving award can be selected as long as it has not been taken by a different agent.* In our model, it is the explicit competition that may drive agents to make fast selections, rather than the rules of the game. That is, the time to select an award is part of an agent’s strategy. This model captures many real-life scenarios, where decisions need not be immediate; rather, the time in which agents make selections is part of their strategy, given the competition. For example, an apartment remains available in the market until it is rented by another potential renter. Similarly, a potential employee remains available until she accepts an offer from another company. In all of these scenarios, a decision need not be immediate, but delaying a decision might result in losing the object if it is taken by another agent in the meantime. This is the tension our model captures.

One issue that arises in this model is how to resolve ties among agents. That is, who gets the award if several agents select it. We consider two natural tie-breaking rules; namely, random tie breaking (where ties are broken uniformly at random) and ranked tie-breaking (where agents are a-priori ranked by some global order, and ties are broken in favor of higher ranked agents). Random tie breaking fits scenarios with symmetric agents, whereas ranked tie breaking fits scenarios where some agents are preferred over others, according to some global preference order. For example, it is reasonable to assume that a higher-position/salary job is preferred over lower-position/salary job, or that firms in some industry are globally ordered from most to least desired. Random and ranked tie-breaking rules were considered by Immorlica *et al.* [2006] and Karlin and Lei [2015], respectively, in secretary settings with immediate and irrevocable decisions.

Two natural objectives have been considered in settings with competition. The first is maximizing the probability of receiving the maximal award (see, e.g., [Immorlica *et al.*, 2006; Karlin and Lei, 2015]). The second is outperforming the competitors (see, e.g., the dueling framework studied by Immorlica *et al.* [2011]). We consider an extension of the latter objective, where an agent wishes to maximize the probability to win the 1st place, then to win the 2nd place, and so on.

Our goal is to study the structure and quality of equilibria in these settings.

1.2 Our Results and Techniques

Random Tie-Breaking

For the random tie-breaking rule, we characterize the equilibria of the induced game, and show that the expected social welfare in equilibrium is nearly optimal, despite competition between the agents. This is cast in Theorems 3.1 and 3.2; a simplified statement follows:

Theorem. (Theorems 3.1 and 3.2) *In every k -agent game with random tie-breaking, there exists a simple time-threshold strategy that guarantees each agent a winning probability of $\frac{1}{k}$, regardless of the strategies of the other agents. The strategy profile where all agents play this strategy is a subgame perfect equilibrium (SPE). Moreover, the expected social welfare of this SPE is at least a $\frac{k-1}{k}$ fraction of the sum of the top k awards.*

In particular, we show that each of the k agents can guarantee herself a winning probability of $\frac{1}{k}$ following a simple time-threshold strategy that depends only on the current time, the number of active agents (i.e., agents who have not yet received an award), and whether the maximal award so far is available. By symmetry, this is the maximal possible guarantee. This guarantee is then used to fully characterize the set of subgame perfect equilibria of the game.

We then establish that in equilibrium, the expected social welfare is at least $\frac{k-1}{k}$ fraction of the sum of the top k awards which is the optimal welfare (i.e., we bound the *price of competition*). We do so by using the following two observations: First, we show that in equilibrium the expected number of selected awards among the top k awards is high. Second, we observe that the probability of an award to be selected in equilibrium is monotone in its rank among the awards.

We complement this result with a matching upper bound (up to a constant factor), which is derived by observing that in equilibrium there is a constant probability that the first selected award is not one of the top k awards.

Ranked Tie-Breaking

For the ranked tie-breaking rule, we show that for a sufficiently large number of agents, the winning probabilities under immediate- and non-immediate selections are roughly the same.

Theorem. (Informal Theorem 4.2) *Under the ranked tie-breaking rule, for every rank i , as the number of agents grows, the winning probabilities of the i^{th} ranked agent under non-immediate selections approaches her winning probability under immediate selections.*

To prove this result, we use observations from [Matsui and Ano, 2016; Karlin and Lei, 2015; Ezra *et al.*, 2018] to show that in the immediate decision model, the probability that the maximal award is allocated goes to 1 as the number of agents grows. Since an agent in the non-immediate decision model can always mimic the strategy of an agent in the immediate decision model, her guarantee for her winning probability (which equals to her probability of receiving the maximal award) in the non-immediate model is at least her guarantee of receiving the maximal award in the immediate decisions game.

We therefore deduce that the winning probabilities in the non-immediate model converge to those in the immediate model. This claim essentially formalizes the intuition that as competition grows, the urgency to select awards faster grows.

In addition, we fully characterize the equilibria of the three-agent game.

Theorem. (Theorem 4.1) *In every equilibrium of the 3-agent game, agent 1 wins with probability $\frac{\ln 4}{2+\ln 4} \approx 0.41$ while each of agents 2, 3 wins with probability $\frac{1}{2+\ln 4} \approx 0.295$.*

Notice that agent 1 (the highest-ranked agent) can always guarantee herself a winning probability of at least $e^{-1} \approx 0.37$ by acting according to the optimal strategy in the classical secretary problem. The last theorem quantifies the benefit that agent 1 derives due to her ability to postpone decisions in a 3-agent setting. As implied by Theorem 4.2, this benefit shrinks as the number of agents grows.

1.3 Related Work

The classical secretary problem and variants thereof have attracted broad interest and have resulted in a vast amount of literature over the years. For a comprehensive survey, see, e.g., [Freeman, 1983].

Competing Agents. The closest papers to our work are the studies by Karlin and Lei [2015] and Immorlica *et al.* [2006], who study secretary settings with competing agents, with the ranked- and random tie breaking rules, respectively. The main difference between their models and ours is that they consider multi-agent settings where agents must make decisions immediately, while in our model the competition is endogenous; namely, past awards can be selected as long as they are available. Karlin and Lei [2015] show that under the ranked tie-breaking rule, the optimal strategy for each agent is a time-threshold strategy, which is given in the form of a recursive formula (albeit not in a closed form). Immorlica *et al.* [2006] characterize the Nash equilibria under the random tie-breaking rule. Another related work is the dueling framework by Immorlica *et al.* [2011], who considered, among other settings, a dueling scenario between two secretary algorithms, whose objective is to outperform the opponent algorithm.

Matroid Secretaries. In this paper we derive insights from studies on secretary variants in which a decision maker can choose multiple awards, based on some feasibility constraints. Babaioff *et al.* [2007] introduced the matroid secretary problem, where a decision maker selects multiple awards under a matroid constraint. It has been shown that a constant competitive ratio can be achieved for some matroid structures, but the optimal competitive ratio for arbitrary matroids is still open. An interesting special case (which was also studied in earlier works such as [Kleinberg, 2005] and [Gilbert and Mosteller, 1966]) is one where the decision maker may choose up to k awards (also known as a k -uniform matroid constraint). Previous works [Gilbert and Mosteller, 1966; Sakaguchi, 1978; Matsui and Ano, 2016; Ezra *et al.*, 2018] studied secretary models in which a decision maker wishes to maximize the probability of getting the highest award, but may choose up to k awards. In Section 4 we draw interesting connections between these models and the ours.

Non-Immediate or Revocable Decisions. Other relaxations of the requirement to select immediately have been considered in the literature. Ho and Krishnan [2015] consider a *sliding-window* variant, where decisions may be delayed for a constant amount of time. A similar model is considered by Kesselheim *et al.* [2019], where decisions may be delayed for a randomized (not known in advance) amount of time. Ezra *et al.* [2018] study settings where the irrevocability assumption is relaxed. Specifically, they consider a setting where the decision maker can select up to ℓ elements immediately and irrevocably, but her performance is measured by the top k elements in the selected set. This work is complementary to ours in the sense that it relaxes the irrevocability assumption, while our work relaxes the immediacy assumption.

1.4 Paper’s Structure

Our model is presented in Section 2. In Sections 3 and 4 we present our results with respect to the random- and ranked

tie-breaking rules, respectively. We conclude this paper in Section 5, where we discuss future directions. All missing proofs are deferred to the full version [Ezra *et al.*, 2020].

2 Model

We consider a variant of the classical secretary setting, where a set of n arbitrary awards are revealed online in a uniformly random order. Let v_t denote the award revealed at time t . Unlike the classical secretary problem that involves a single decision maker, in our setting there are k agents who compete over the awards. Upon the revelation of award v_t , every agent who has not received an award yet may select one of the awards v_1, \dots, v_t that is still available. An award that is selected by a single agent is assigned to this agent. An award that is selected by more than one agent is assigned to one of these agents either randomly (hereafter, *random tie breaking*), or according to a predefined ranking (hereafter, *ranked tie breaking*). Agents who received awards are no longer active. Awards that were assigned are no longer available. Thus, at time t , the set of *available* awards is the subset of awards v_1, \dots, v_t that have not been assigned yet. The game continues as long as there are active agents. I.e., after time n , if active agents remain, the agents compete (without newly arriving awards) on the remaining available awards until all agents are allocated.

Given an instance of a game, the history at time t includes all the relevant information revealed up to time t ; i.e., v_1, \dots, v_t , and the assignments up to time $t - 1$ ¹. A strategy of agent i , denoted by S_i , is a function from the set of all possible histories to a selection decision (either selecting one of the available awards, or passing). A strategy profile is denoted by $S = (S_1, \dots, S_k)$. We also denote a strategy profile by $S = (S_i, S_{-i})$, where S_{-i} denotes the strategy profile of all agents except agent i . Every strategy profile S induces a distribution over assignments of awards to agents. For ranked tie breaking, the distribution is with respect to the random order of award arrival, and possibly the randomness in the agent strategies. For random tie breaking, the randomness is also with respect to the randomness in the tie breaking.

We say that agent i wins the j^{th} place in the game if she receives the j^{th} highest award among all allocated awards. Let $p = (p_1, \dots, p_k)$ be a k -dimensional vector, where p_j is the probability to win the j^{th} place. As natural in competition settings, every agent cares only about her relative performance with respect to her competitors. Consequently, we find the following objective function the most natural in this setting: every agent wishes to maximize her probability to win the first place; upon equality she wishes to maximize the probability to win the second place, and so on. Thus, given two vectors $p, \bar{p} \in R^k$, p is preferred over \bar{p} , denoted $\bar{p} \prec p$, if p is lexicographically greater than \bar{p} . Similarly, p is weakly preferred over \bar{p} , denoted $\bar{p} \preceq p$, if p is lexicographically greater or equal to \bar{p} .

A strategy profile S induces a k -dimensional probability vector $p^i(S)$ for each agent i , where $p_j^i(S)$ is the probability

¹In our setting, additional information, such as the history of selections (in contrast to assignments) is irrelevant for future decisions.

that agent i wins the j^{th} place under S . Agent i derives higher (respectively, weakly higher) utility from strategy profile S than strategy profile \bar{S} , denoted $\bar{S} \prec_i S$ (resp., $\bar{S} \preceq_i S$), if $p^i(\bar{S}) \prec p^i(S)$ (resp., $p^i(\bar{S}) \preceq p^i(S)$). We use $\bar{S} \prec_i S$ and $S \succ_i \bar{S}$ interchangeably, and similarly for p and \bar{p} .

Note that $p_1^i(S)$ is the probability that agent i wins the first place under strategy profile S . We sometimes refer to it as agent i 's *winning probability* under S .

We consider the following equilibrium notions:

- A strategy profile $S = (S_1, \dots, S_k)$ is a **Nash equilibrium** (NE) if for every agent i and every strategy S'_i , it holds that $(S'_i, S_{-i}) \preceq_i (S_i, S_{-i})$.
- A strategy profile $S = (S_1, \dots, S_k)$ is a **Subgame Perfect Equilibrium** (SPE) if it is a NE for every subgame of the game. I.e. for every initial history h , S is a NE in the game induced by history h .

Every SPE is a NE, but not vice versa.

We now illustrate our models in the following example:

Example 1. Consider a setting with 5 houses for rent and 3 potential renters. The houses “enter the market” (i.e., become available) in a random order and remain available until rented by some renter.

Random tie-breaking captures scenarios where the house owners are indifferent between different renters. Upon the entrance of the first house, the renters do not select it since the probability that this house is the best one among all 5 houses is only $\frac{1}{5}$, which is too low (the threshold here is $1/3$, see Theorem 3.1). Upon the entrance of the second house, the better among the two houses is the best among all five houses with probability $\frac{2}{5} > \frac{1}{3}$, and so all three agents select it. Assume that renter #1 was assigned the selected house (by the random assignment process), and further assume that the third house is the highest among the first three houses. In this case, the two remaining renters will select this house and it will be assigned to one of them at random. The one remaining renter will select the maximal available house at the end of the process. By symmetry, each one of the three agents wins with probability $\frac{1}{3}$.

Ranked tie breaking captures scenarios where the house owners have identical preferences over renters (e.g., ranking according to the credit history of the renters).

Upon the entrance of the first house, the renter ranked #3 selects it and guarantees herself a winning probability of $\frac{1}{5}$ (one can show that the other renters can guarantee themselves a higher winning probability, and that if renter #3 does not select it, then her winning probability in SPE is upper bounded by $\frac{1}{6}$). Suppose next that the second house is worse than the first one. Then, it is not selected by any of the remaining renters. Suppose further that the third house is the best among the first three houses. Then, both remaining renters select it, and it is being assigned to renter #1 (by ranked tie-breaking). A full analysis of the SPE shows that the renters' winning probabilities are $\frac{31}{60}$, $\frac{17}{60}$, and $\frac{12}{60}$, respectively.

3 Random Tie-Breaking

In this section, we study the setting of the random tie-breaking rule. We characterize the SPEs and give simple time

threshold strategies with optimal utility guarantees. We show that in the SPE where all agents play according to this optimal guarantee strategy, at least $\frac{k-1}{k}$ of the optimal social welfare is achieved in expectation.

Consider the following strategy σ_i for agent i , where ℓ_t denotes the number of active agents at time t (agent i included).

- If $t \geq n$, then select the maximal available award.
- If $\frac{n}{\ell_t} < t < n$, and the maximal award so far is available, then select it.
- If $t = \frac{n}{2}$, $\ell_t = 2$, and the maximal award so far is available, then select it.
- Else, pass (i.e., select no award).

We denote by \mathcal{S}_i the set of strategies in which agent i plays according to σ_i up to the following cases:

- If $\ell_t = 2$, $t = \frac{n}{2}$, and both the highest and the second highest awards so far are available, then the agent can either pass or select the highest award so far.
- If $\ell_t = 1$ and the highest award so far is available, then the agent can either pass or select this award.

We next show that strategies in \mathcal{S}_i are the only strategies that guarantee a utility of at least $(\frac{1}{k}, \dots, \frac{1}{k})$. By symmetry, there is no strategy σ such that $(\frac{1}{k}, \dots, \frac{1}{k}) \prec p^i(\sigma, S_{-i})$ for every S_{-i} . Thus, the strategy profiles where each agent i plays according to a strategy in \mathcal{S}_i , are the only SPEs.

Theorem 3.1. For every agent i , $\sigma \in \mathcal{S}_i$ and S_{-i} , it holds that $(\frac{1}{k}, \dots, \frac{1}{k}) \preceq p^i(\sigma, S_{-i})$. For every $\sigma' \notin \mathcal{S}_i$ there exists S_{-i} such that $(\frac{1}{k}, \dots, \frac{1}{k}) \succ p^i(\sigma', S_{-i})$. Moreover, the strategy profiles where each agent j plays according to a strategy in \mathcal{S}_j , are the only SPEs.

Before proving Theorem 3.1, we show the following:

Observation 3.1. For every time $t < n$, selecting an element that is not the maximal so far cannot guarantee a winning probability of $\frac{1}{k}$.

Thus, we can assume that agents do not select elements that are not the maximal so far up to time $t = n$. We now give lower bounds on the probabilities of winning first and second places in a strategy $\sigma \in \mathcal{S}_i$ given a time t , and whether the maximal and second maximal awards so far are available. For every $1 \leq \ell \leq k$, let $A_t^\ell \in [0, 1]^2$ be an ordered pair denoting a lower bound on the probabilities of agent i winning first and second places under strategy profile (σ, S_{-i}) , conditioned on the event that at time t (after observing the award v_t , but before making selections in time t) agent i is active, there are ℓ active agents (including agent i), and the maximal and second maximal awards up to time t are available. Similarly, let B_t^ℓ be a lower bound on the probabilities of agent i winning first and second places under strategy profile (σ, S_{-i}) , conditioned on the event that at time t , agent i is active, there are ℓ active agents (i included), and the maximal award up to time t is available, but the second maximal award is not available.

Let C_t^ℓ be a lower bound on the probabilities of agent i winning first and second places under strategy profile (σ, S_{-i}) , conditioned on the event that *after the allocations* of time t , agent i is active, there are ℓ active agents (including agent i),

and the maximal award up to time t is not available, but the second maximal award is available. Finally, let D_t^ℓ be a lower bound on the probabilities of agent i winning first and second places under strategy profile (σ, S_{-i}) , conditioned on the event that after the allocations of time t , agent i is active, there are ℓ active agents (i included), and none of the maximal and the second maximal awards up to time t are available.

In Lemma 3.1 we lower bound the above terms. The complete proof of the lemma is deferred to the full version.

Lemma 3.1. *For every t , and every $k \geq \ell > 1$, it holds that:*

$$\begin{aligned} \bullet A_t^\ell &\geq \left(\frac{1}{\ell}, \frac{1}{\ell}\right) & \bullet C_t^\ell &\geq \left(\frac{n-t}{n\ell}, \frac{n^2+t^2-tn-n}{n(n-1)\ell}\right) \\ \bullet B_t^\ell &\geq \left(\frac{1}{\ell}, \frac{(n-t)(n+t-1)}{n(n-1)\ell}\right) & \bullet D_t^\ell &\geq \left(\frac{n-t}{n\ell}, \frac{n-t}{n\ell}\right) \end{aligned}$$

Proof sketch. We observe that the winning probability of an agent depends on the time step t , the number of active agents ℓ , and whether the maximal award so far is available or not. If at time t an agent receives the maximal award up to time t , she wins with probability $\frac{t}{n}$ (which is the probability that this award is the global maximum). If another agent receives the maximal award up to time t , then by symmetry, each remaining active agent can guarantee a winning probability of $\frac{n-t}{n(\ell-1)}$. Thus, selecting the maximal award so far is better whenever $\frac{t}{n} > \frac{n-t}{n(\ell-1)}$, and passing is better whenever $\frac{t}{n} < \frac{n-t}{n(\ell-1)}$. In cases where $\frac{t}{n} = \frac{n-t}{n(\ell-1)}$, both passing and selecting the maximal award so far give a winning probability of $\frac{1}{\ell}$, and the agents break this tie based on the probability of winning the second place. For the four states of whether the maximal and second maximal awards so far are available, we establish lower bounds on the probabilities of winning first and second places, by induction on t and ℓ . \square

We are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. It follows from the proof of Lemma 3.1 that for every strategy that is not in \mathcal{S}_i , if each agent $j \neq i$ plays according to a strategy in \mathcal{S}_j , agent i 's utility is smaller than $(\frac{1}{k}, \frac{1}{k}, 0, \dots, 0, \frac{k-2}{k})$. It also shows that if agent i plays according to a strategy in \mathcal{S}_i , and there exists an agent j that plays according to a strategy not in \mathcal{S}_j , utility is greater than $(\frac{1}{k}, \frac{1}{k}, \frac{k-2}{k}, 0, \dots, 0)$.

Thus, the only SPEs are profiles in which each agent j plays according to a strategy in \mathcal{S}_j , and by symmetry, the utility of agent i is exactly $(\frac{1}{k}, \dots, \frac{1}{k})$, as desired. \square

In Theorem 3.1 we characterize the structure of all SPEs. Immorlica *et al.* [2006] prove that in the immediate decisions model, whenever $\frac{t}{n} \geq \frac{1}{\ell_t}$, and the current arriving award is the maximal so far, all active agents select it. However, they do not provide a full characterization of the cases where $\frac{t}{n} < \frac{1}{\ell_t}$. Indeed, in some cases agents do select the maximal award so far at such times t . In contrast, we show that in the non-immediate decisions model, the times where $\frac{t}{n} \geq \frac{1}{\ell_t}$ are the *only* times where the agents make selections.

We next study the social welfare (i.e., the sum of the awards received by all agents) obtained in an SPE. Let y_i denote the i^{th} maximal value among v_1, \dots, v_n . Then, the optimal social welfare is $\text{OPT} = \sum_{i=1}^k y_i$.

The following theorem asserts that despite the competition, the social welfare of the SPE where all agents play according to σ_i is at least $\frac{k-1}{k}$ of the optimal social welfare.

Theorem 3.2. *The expected sum of the allocated awards in the SPE profile $S = (\sigma_1, \dots, \sigma_k)$ is at least $\frac{k-1}{k} \cdot \text{OPT}$.*

We complement this result by showing an instance of awards y_1, \dots, y_n where the social welfare in every SPE is at most $\frac{k-\Omega(1)}{k} \cdot \text{OPT}$ for every $k > 1$.

Example 2. *Suppose $y_1 = \dots = y_k = 1$ and $y_j = 0$ for every j such that $k < j \leq n$. In every SPE, the first selection is made at time no later than $t = \lfloor \frac{n}{k} \rfloor + 1$. If none of the top k awards appeared up to time t , at least one of the agents gets an award of 0. The probability that none of y_1, \dots, y_k appeared by time t is approximately $(\frac{k-1}{k})^k = \Omega(1)$. Thus, the expected social welfare is at most $\frac{k-\Omega(1)}{k} \cdot \text{OPT}$.*

4 Ranked Tie-Breaking

In this section, we study competition under the ranked tie-breaking rule. We first claim that without loss of generality, for every agent the winning probability equals to the probability of receiving the highest award. To show this, we observe that whenever exactly one agent is active, she may as well wait until time $t = n$ and only then select the maximal award without harming her utility. Thus, it can be assumed that the maximal award is always allocated, and the winning agent receives it. Thus, we may assume that the first-order objective of every agent is to maximize the probability of receiving the maximal award, as in the standard secretary problem and previous multi-agent extensions.

In Section 4.1 we present general observations regarding equilibria in this setting. We then characterize the equilibrium in the 3-agent game in Section 4.2. In Section 4.3 we show that as the number of competing agents goes to infinity, the agents' probabilities of receiving the highest award (which equal to the agents' winning probabilities) converge to the corresponding probabilities in the immediate decisions model described by Karlin and Lei [2015].

4.1 General Observations

We first make observations about the structure of the subgame perfect equilibria (SPE) of the game.

Proposition 4.1. *A strategy profile $S = (S_1, \dots, S_n)$ is an SPE if for every agent i , S_i is described by a set of time thresholds T_j^ℓ for every j, ℓ such that $1 \leq j \leq \ell \leq k$. At time t , agent i selects the highest award so far if it is available and $t \geq T_j^\ell$, where the current number of active agents is ℓ and agent i is ranked j^{th} among them.² In addition, if ℓ agents are active and $t > n - \ell$, then the lowest-ranked active agent makes a selection, even if the highest award so far is not available.*

We proceed with several observations about the time thresholds in the SPE of the game. Since any agent can always mimic the strategy of an agent ranked lower than her,

²If $t = T_j^\ell$, then the agent is indifferent between selecting and passing.

in equilibrium a lower-ranked agent would be willing to receive any award that a higher-ranked agent would be willing to receive. In the threshold terminology, it means that:

Observation 4.1. *For any number of active agents ℓ , for every pair of ranks h, j such that $h < j \leq \ell$, without loss of generality it holds that $T_j^\ell \leq T_h^\ell$.*

The following observation gives bounds on the time threshold of the lowest-ranked active agent relative to the second-lowest ranked active agent.

Observation 4.2. *For any number of active agents ℓ , it holds that $T_{\ell-1}^\ell \geq T_\ell^\ell \geq T_{\ell-1}^\ell - 1$.*

Proof. By Observation 4.1 we have that $T_{\ell-1}^\ell \geq T_\ell^\ell$. On the other hand, the lowest-ranked active agent never makes a selection before time $\min_{j \neq \ell} T_j^\ell - 1$, because she can only benefit from waiting as long as no other active agent makes a selection. The claim now follows since, by Observation 4.1, $\min_{j \neq \ell} T_j^\ell = T_{\ell-1}^\ell$. \square

Recall that the winning probability of agent i under strategy profile S is denoted by $p_i^1(S)$. Throughout this section, we make two simplifications in notation. First, we omit S . Second, we omit the subscript 1, since we consider only the probability of winning the 1st place. Consequently, we denote the probability that agent i wins the 1st place in strategy profile S by p_i .

The following observation gives bounds on the winning probability of the lowest-ranked agent relative to the second-lowest agent.

Observation 4.3. *It holds that $p_{k-1} - \frac{1}{n} \leq p_k \leq p_{k-1}$.*

4.2 The 3-Agent Game

In a 2-agent game Observation 4.3 implies that $\frac{1}{n} \geq p_1 - p_2 \geq 0$. That is, both agents win with probability roughly a half. This symmetry breaks as more agents join the game and the setting becomes interesting already in the case of 3 agents. Notice that the highest-ranked agent can always guarantee herself a probability of at least $e^{-1} \approx 0.37$ to receive the highest award by adopting the optimal strategy in the classical secretary problem. An interesting question is whether the opportunity to make non-immediate decisions increases this probability for the highest-ranked agent.

We show the following:

Theorem 4.1. *In a setting with 3 agents, in any SPE, agent 1 wins with probability ≈ 0.41 , while each of agents 2, 3 wins with probability ≈ 0.295 .*

4.3 Immediate vs. Non-Immediate Selections

In this section, we compare the immediate and non-immediate models for games with a large number of agents. Let $p_{i,k}$ denote the probability that agent i wins in a k -agent game, with non-immediate selections, and let q_i denote the probability that agent i receives the highest award in a game with immediate selections. We note that under immediate selections, q_i is independent of the number of agents.

The main result here is that agents' winning probabilities in equilibrium under non-immediate selections approach their winning probabilities under immediate selections, as the number of agents grows.

Theorem 4.2. *For every i , $\lim_{k \rightarrow \infty} p_{i,k} = q_i$.*

The full proof of Theorem 4.2 is deferred to the full version. We give here the high-level idea of the proof. We use the following theorem by Karlin and Lei [Karlin and Lei, 2015] regarding the immediate decision model.

Theorem 4.3 (Karlin and Lei [2015]). *For every n, k and every $i \in [k]$, there is a unique T_i (independent of k) such that agent i plays a T_i -threshold strategy in SPE; namely, wait until time T_i , then make a selection whenever a best-so-far award appears. It holds that $T_{i-1} \geq T_i$, and $q_i = \frac{T_i}{n}$, for all i . Moreover, threshold strategy T_i guarantees agent i a winning probability of q_i regardless of other agent strategies.*

We then show that in the non-immediate model, every agent i can mimic strategy T_i specified in Theorem 4.3 and guarantee herself the same guarantee of q_i , thus for all i, k , it holds that $p_{i,k} \geq q_i$. Moreover, using results by [Matsui and Ano, 2016] and [Gilbert and Mosteller, 1966], it follows that $\lim_{k \rightarrow \infty} \sum_{i \leq k} q_i = 1$. We show that this implies that $\lim_{k \rightarrow \infty} p_{i,k} \leq q_i$, since otherwise, there exists k such that $\sum_i p_{i,k} > 1$.

5 Discussion and Future Directions

We study secretary settings with competing decision makers. While in previous secretary settings, including ones where competition among multiple agents are considered, decisions must be made immediately, we introduce a model where the time of selection is part of the agent's strategy, and thus the competition is endogenous. In particular, decisions need not be immediate, and agents may select previous awards as long as they are available. These settings capture many real-world settings, where agents compete over "awards" that may remain available until taken by a competitor.

This work suggests open problems and directions for future research. For the ranked tie-breaking rule, we fully characterize the equilibria of a 3-agent game, and derive the corresponding utilities of the agents. Extending this characterization to any number of agents is an interesting open problem.

Below we list some future directions that we find particularly natural: (1) Study competition in additional problems related to optimal stopping theory, such as prophet and Pandora box settings. (2) Study competition in secretary settings under additional tie-breaking rules, such as random tie breaking with non-uniform distribution, and tie-breaking rules that allow to split awards among agents. (3) Study competition in secretary settings under additional feasibility constraints. For example, scenarios where agents can choose up to k awards, or other matroid constraints. (4) Extend the current study to additional objective functions.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132, 740282), and by the Israel Science Foundation (grant number 317/17).

The full version of this work can be found at [Ezra *et al.*, 2020].

References

- [Babaioff *et al.*, 2007] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 434–443, 2007.
- [Babaioff *et al.*, 2008] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *ACM SIGecom Exchanges*, 7(2):1–11, 2008.
- [Ezra *et al.*, 2018] Tomer Ezra, Michal Feldman, and Ilan Nehama. Prophets and secretaries with overbooking. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 319–320, 2018.
- [Ezra *et al.*, 2020] Tomer Ezra, Michal Feldman, and Ron Kupfer. On a competitive secretary problem with deferred selections. *arXiv preprint arXiv:2007.07216*, 2020.
- [Ferguson, 1989] Thomas S. Ferguson. Who solved the secretary problem? *STATISTICAL SCIENCE*, 4(3), 1989.
- [Freeman, 1983] P. R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review / Revue Internationale de Statistique*, 51(2):189–206, 1983.
- [Gilbert and Mosteller, 1966] John P. Gilbert and Frederick Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, 1966.
- [Ho and Krishnan, 2015] Shan-Yuan Ho and Abijith Krishnan. A secretary problem with a sliding window for recalling applicants. *arXiv preprint arXiv:1508.07931*, 2015.
- [Immorlica *et al.*, 2006] Nicole Immorlica, Robert Kleinberg, and Mohammad Mahdian. Secretary problems with competing employers. In *International Workshop on Internet and Network Economics*, pages 389–400. Springer, 2006.
- [Immorlica *et al.*, 2011] Nicole Immorlica, Adam Tauman Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz. Dueling algorithms. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 215–224. ACM, 2011.
- [Karlin and Lei, 2015] Anna Karlin and Eric Lei. On a competitive secretary problem. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [Kesselheim *et al.*, 2013] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *European symposium on algorithms*, pages 589–600. Springer, 2013.
- [Kesselheim *et al.*, 2019] Thomas Kesselheim, Alexandros Psomas, and Shai Vardi. How to hire secretaries with stochastic departures. In *Web and Internet Economics - 15th International Conference, WINE 2019*, volume 11920, page 343, 2019.
- [Kleinberg, 2005] Robert D Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, volume 5, pages 630–631, 2005.
- [Matsui and Ano, 2016] Tomomi Matsui and Katsunori Ano. Lower bounds for bruss’ odds problem with multiple stoppings. *Mathematics of Operations Research*, 41(2):700–714, 2016.
- [Sakaguchi, 1978] M Sakaguchi. Dowry problems and ola policies. *Rep. Stat. Appl. Res., JUSE*, 25:124–128, 1978.