The Secretary Problem with Competing Employers on Random Edge Arrivals

Xiaohui Bei[1], **Shengyu Zhang**[2]
[1] Nanyang Technological University
[2] Tencent
xhbei@ntu.edu.sg, shengyzhang@tencent.com

Abstract

The classic secretary problem concerns the problem of an employer facing a random sequence of candidates and making online hiring decisions to try to hire the best candidate. In this paper, we study a game-theoretic generalization of the secretary problem where a set of employers compete with each other to hire the best candidate. Different from previous secretary market models, our model assumes that the sequence of candidates arriving at each employer is uniformly random but independent from other sequences. We consider two versions of this secretary game where employers can have adaptive or non-adaptive strategies, and provide characterizations of the best response and Nash equilibrium of each game.

1 Introduction

The *secretary problem* represents a canonical setup for studying online algorithms. Consider a player trying to hire the best secretary among a totally-ordered set of n candidates. The candidates arrive in a random order. At time step t, a candidate a_t arrives. The employer learns this candidate's relative order among all candidates arrived so far, but does not know how to compare it with the future candidates that he has not seen. The employer needs to make an irrevocable decision on whether to hire the current candidate a_t or not. If he hires, then the game is over, and the employer wins if the hired a_t is the best among all the n candidates. If the employer does not hire, then the next candidate a_{t+1} arrives and the game goes on. The secretary problem asks how to design a strategy for the employer to maximize the chance of hiring the best candidate. The optimal strategy for this problem turns out to be passing the first $(1/e)$-fraction of the candidates regardless of their relative rankings, and then hiring the first candidate afterwards that is better than the best he has seen so far. This optimal strategy first appeared almost six decades ago (Dynkin 1963) and has been discussed in reviews (Ferguson 1989). The secretary problem has also been generalized along various directions. We refer to (Freeman 1983; Ajtai, Megiddo, and Waarts 2001) for some good surveys.

Among these generalizations, one natural extension is when there are multiple employers competing for hiring the best candidate from a pool of candidates. This setting effectively converts the problem of online decision making into an online secretary *game* played among multiple employers. In this secretary game, a random stream of candidates arrives one by one. Upon the arrival of each candidate, a set of employers all learn the relative rankings of this candidate and then each makes a hiring decision. The goal of each employer is again to maximize the probability of hiring the best candidate. A sequence of works (Dynkin 1967; Szajowski 1993, 1995; Fushimi 1981) study the secretary problem with two competing employers and analyze the equilibrium behavior. These works differ in their assumptions on the payoff structure and how the conflicts are resolved when multiple employers make offers to the same candidate. Sakaguchi (Sakaguchi 1995) reviews the game from an optimal stopping strategy perspective. Immorlica et al (Immorlica, Kleinberg, and Mahdian 2006) consider a general setting with multiple indistinguishable employers and argue that the timing of the earliest offer decreases as the number of employers increases. Karlin and Lei (Karlin and Lei 2015) consider the secretary game with ranked employers and show how to compute subgame-perfect Nash equilibrium strategies of the game. Ezra et al. (Ezra, Feldman, and Kupfer 2020) consider a variant of the setting that allows deferred selections from the employers. Note that all these works discussed above use a *vertex arrival* model. That is, the employers are interviewing the same random sequence of the candidates, and when a candidate arrives, this candidate applies to all employers in a certain (e.g. fixed or random) order. It is also assumed that all employers have the same preferences over the candidates. This fits the scenarios where the candidates possess some intrinsic values that all employers can observe from an interview, or that the only information of available to the employers to make decisions on a candidate is a score from some standard test.

In this work, we consider an *edge arrival* model of this problem. More precisely, suppose that we still have n candidates $1, 2, \ldots, n$, associated with intrinsic values $v(1), v(2), \ldots, v(n)$, respectively. There are $k \geq 2$ employers competing for hiring the best of these candidates. Consider a complete bipartite graph (L, R, E) where L is the set of k employers, R is the set of n candidates, and E is initialized to $L \times R$. In an edge arrival model, at each time step, a random pair $(i, j) \in E$ arrives, corresponding to candidate

j applying for job provided by employer i. The employer i sees the value $v(j)$ and needs to decide irrevocably whether to hire candidate j on the spot. If he does not hire, then we update $E \leftarrow E - \{(i,j)\}$ and the game continues to the next time step. If the employer hires and the current candidate is indeed the best among the n candidates, then the employer wins the game (and the game ends). If the employer hires but the hired one is not the best, then the employer is out of the game and the rest employers continue; formally, we update $L \leftarrow L - \{i\}$ and $E \leftarrow E - \{(i,j') : j' \in R\}$.

We give a few remarks about vertex arrival and edge arrival models. The vertex arrival model is suitable for the situation where a candidate applies to *all* secretary positions upon coming to the job market. But it is rarely the case in reality, as in reality one usually just applies very few or even one position first, and only applies to more if failing to get an offer in the previous application. This is closer to the edge arrival model studied in this paper. Another scenario that is better characterized by the random edge arrival model is the dating market, where men and women encounter in a rather random pattern, and surely it is not the case that a person announces his or her availability to the whole other side. Such edge arrival model has also been considered in (Cownden and Steinsaltz 2014). Their model differs from ours in a crucial point that in their model, a candidate rejected by any player will immediately leave the market and won't be seen by any other players. In our motivating examples, this is clearly not the case.

Sometimes the secretary problem is also studied in situations where one or several players pick from a sequence of items. In that language, the players correspond to the employers in the hiring game, and the items to the candidates. We may use the two sets of words interchangeably.

1.1 Results

In this work, we analyze the Nash equilibria of the online secretary game with competing employers in the edge arriving model. We focus on a continuous-time setting that captures the essence of the game but avoids some messy integrality issue in computations.

Similar to many models from previous works, in these games, each player's strategy at any given time can be described by a *threshold*, which specifies the time until when the player is not making any offers to any candidates, regardless of their ranks, and after the threshold time, the player will hire the first candidate that is better than all candidates arrived so far. We consider two models in which the employers can be either adaptive or non-adaptive. When employers are adaptive, they are allowed to update their threshold during the hiring process based on other employers' actions. On the other hand, non-adaptive employers can only stick to one threshold time throughout the process.

Adaptive players. We first consider adaptive players in Section 3. We provide a complete characterization of the Nash equilibria of the general k-player game with adaptive strategies. Specifically, we show that the game again has a unique Nash equilibrium, and actually it is symmetric and can be described by k thresholds $(r_k^*, r_{k-1}^*, \ldots, r_1^*)$, with $r_k^* < r_{k-1}^* < \cdots < r_1^* = 1/e$. Here r_i^* is the threshold time for any player when there are i players remaining in the game. In other words, in this unique Nash equilibrium, all k players start by setting the same threshold time r_k^*. As time goes by, whenever some player makes an incorrect hire (and therefore leaves the game), the remaining players will all change their strategy by switching to the next threshold time r_{k-1}^* as the game continues. In the end, when there is only one player left in the game, that player will use her threshold time to the classic one-player threshold r_1^*.

One may naturally like to know what these threshold values exactly are. We manually computed $r_1^* = 1/e, r_2^* \approx 0.29574$, and $r_3^* \approx 0.24006$. While it seems unlikely to have a closed-form formula for r_k^*, we give an algorithm with running time polynomial in k to compute these threshold values.

Non-adaptive players. In the second part, we consider players with non-adaptive strategies. We start with the two-player game and prove that the game has a unique Nash equilibrium in which both players set the threshold time to $r^* \approx 0.29533$, a value slightly smaller than $r_2^* \approx 0.29574$ in the *adaptive* two-player game.

Next, for the general k-player game, we focus on symmetric Nash equilibrium and prove that the equilibrium threshold r^* is the unique root of an equation. We further analyze the properties of this symmetric Nash equilibrium, and show that when the number of players k increases, this equilibrium threshold converges to 0 in the order of $1/k$.

The rest of the paper is organized as follows. In the next section, we formally define the one-player and multiple-player games, and review the threshold strategy for the single player case. We then study adaptive and non-adaptive strategies in the multiple-player games in Section 3 and 4, respectively. Finally we conclude in Section 5 with some open questions. Due to space constraints, the full proofs of many results below are omitted and can be found in the full version of this paper.

2 Preliminaries

We consider a setting where k employers compete to hire among n candidates, Sometimes we say that players pick items for employers hire candidates. We use $[n]$ to denote the set $\{1, 2, \ldots, n\}$. Let us introduce the model and preliminary methods starting from the single player case.

2.1 Single Player

Discrete-time model. There are n candidates arriving in an order chosen from all $n!$ permutations uniformly at random. The employer has a strict preference between any two candidates, and we use notation $a \succ b$ to denote that candidate a is preferred to candidate b. At each discrete time step $t = 1, 2, \ldots, n$, one candidate arrives. The employer can compare this candidate with any of the previous $t-1$ candidates, but cannot compare these t arrived candidates with the future $n-t$ ones that he has not seen. The employer needs to decide whether to hire the current candidate based on the candidates seen so far (including the current one). The game ends once the employer hires someone; if the employer does not hire any of the n candidates, the game also ends. The employer wins if and only if she hires the best candidate among the n candidates. The employer aims to maximize her winning probability.

Continuous-time Model. It is sometimes convenient to study a *continuous-time model* that captures the limit case when the number of candidates n approaches infinity. In this model, we assume the candidates form the set of natural numbers $\mathbb{N}$ with the usual ordering. Candidate i is better than candidate j if $i < j$, thus candidate 1 is the best. Each candidate $i \in \mathbb{N}$ arrives at time $t(i)$, which is sampled from $[0, 1]$ independently and uniformly at random. This also implies that at any time $t \in [0, 1]$, with probability 1, the fraction of all candidates that the employer has seen is also t.

The continuous-time model keeps the essence of the problem without some messy details in the discrete case due to the integrality, and it is the model that we adopt in this paper.

For convenience of notation, let us make the following definitions.

Definition 2.1. *We make the following definitions about candidates.*

- $\mathsf{candidate}(t)$ = *the candidate arriving at time t*
- $\mathsf{best}(I)$ = *the best candidate among all that arrive during time interval I*
- $\mathsf{best}(t) = \mathsf{best}([0, t])$.

Note that in the above definition, we ignore the event that two or more candidates arriving at the same time, because it happens with probability 0. Also note that $\mathsf{best}(1)$ is the best candidate among all.

Strategy. A *strategy* is a decision function mapping time $t \in [0, 1]$ and all historical information by time t to a binary decision in $\{0, 1\}$, indicating whether or not to hire the current candidate arriving at time t. The decision can depend on the rank information about all candidates that have arrived so far. Note that the employer aims to hire the best candidate $\mathsf{best}(1)$, thus she will consider hiring at time t only if $\mathsf{candidate}(t) = \mathsf{best}(t)$.

A particular type of strategies is that of the *threshold strategies*. The t^*-threshold strategy is the following one: The employer hires $\mathsf{candidate}(t)$ if and only if $t > t^*$ and $\mathsf{candidate}(t) = \mathsf{best}(t)$. It is easily seen that if the employer follows the t^*-threshold strategy, then for any $t > t^*$, $\mathsf{candidate}(t)$ is better than $\mathsf{best}([0, t))$ if and only if $\mathsf{candidate}(t)$ is better than $\mathsf{best}([0, t^*))$. It turns out that the optimal strategy for this hiring problem is always a threshold strategy. One can derive the optimal threshold in the discrete model by standard dynamic programming (Freeman 1983). In the same spirit, one can derive the optimal threshold in the continuous-time model, as shown in the following lemma.

Lemma 2.2. *The winning probability of a player using the t^*-threshold strategy is $t^* \ln \frac{1}{t^*}$. This is maximized at $t^* = 1/e$, for which the winning probability is also $1/e$.*

2.2 Multiple Players

Next we consider the case with multiple employers competing to hire the best candidate. Suppose that there are k employers, and the candidates form the set of natural numbers $\mathbb{N}$ with the usual ordering as before. Note that in this way, we implicitly assume that all players have the same ranking preference over the candidates. This assumption was also adopted in previous studies of multiple employer case in vertex-arrival models (Fushimi 1981; Immorlica, Kleinberg, and Mahdian 2006; Karlin and Lei 2015). Each candidate i arrives to player j at a random time $t_{ij} \in [0, 1]$. All of these random variables are independent, and each t_{ij} follows the uniform distribution over $[0, 1]$. When a player hires the best candidate, that player wins and the game ends for all players. When a player hires a sub-optimal candidate, that player is out of the hiring process (without winning), and this information is announced to all the other players who remain in the game.

We can make similar definitions about the candidate as in Definition 2.1, with a superscript (i) indicating that it is for player i. Thus $\mathsf{candidate}^{(i)}(t)$ is the candidate arriving at player i at time t. When the player is clear from context, we may also drop the superscript.

Adaptive strategies. At any time t, each player i in the game knows how many players are still in the hiring process. Her strategy can be based on this as well as all the historical information. For player i, an adaptive threshold strategy can be represented by a sequence of thresholds r_{ij}, and the strategy is that player i plays the r_{ij}-threshold strategy as long as the number of remaining players (including herself) is j. We will show that $0 \le r_{ik} \le \cdots \le r_{i1} \le 1$. The joint strategy for all k players is thus represented by the k^2 thresholds $\{r_{ij} : i, j \in [k]\}$.

An adaptive strategy $\{r_{ij} : i, j \in [k]\}$ is *symmetric* if $\forall j \in [k]$, $r_{1j} = r_{2j} = \cdots = r_{kj} = r_j$ for some r_j; that is, all players use the same adaptive strategy.

Nonadaptive strategies. We call player i's strategy $\{r_{ij} : j \in [k]\}$ *non-adaptive* if all these k thresholds are the same: $r_{i1} = r_{i2} = \cdots = r_{ik} = r_i$ for some r_i, namely player i does not adapt her action to the information of how many players are still in the game. If all players use non-adaptive strategies, then we say the joint strategy is non-adaptive.

A non-adaptive strategy $\{r_i : i \in [k]\}$ is *symmetric* if all r_i's are the same; that it, all players use the same r-threshold strategy.

3 Adaptive Strategies

In this section, we consider multiple players with adaptive strategies $\{r_{ij} : i, j \in [k]\}$, i.e. each player i uses the r_{ij}-threshold strategy when there are j players left in the game.

The following conditional winning probability is a key concept in later analysis.

Definition 3.1. *At time r, conditioned on that the best candidate has not been hired by any player yet, and there are j players in the game, let the winning probability of player i be $f_j^{(i)}(r)$. That is, we define the* conditional winning probability

$$f_j^{(i)}(r) = \Pr[\textit{Player } i \textit{ wins} | \textit{game has } j \textit{ players at time } r].$$

Basic properties. Here we list some basic properties of the secretary game with adaptive strategies. First, it is easy to see that similar to the single player case, the optimal strategy is a collection of threshold strategies.

Lemma 3.2. *No matter what strategies other players adopt, the optimal strategy for player i is a threshold strategy $\{r_{ij} : j \in [k]\}$, in which $r_{ik} \leq \cdots \leq r_{i1}$.*

Recall that when we introduce the notation r_{ij}, the game starts with k players, and r_{ij} is the threshold that player i takes when there are j players left in the game. Why we did not put k, the total number of players at the start of the game, inside the notation? The next lemma justifies our choice.

Lemma 3.3. *Consider a game with k players. For each $i, j \in [k]$, the best threshold r^*_{ij} for player i does not depend on the starting number k of the players. Moreover, the optimal threshold r_{i1} equals to $1/e$, regardless of the number of players at the beginning of the game.*

Two properties of Nash equilibria. Next we present two important properties regarding the Nash equilibria in a k-player game. The first one is that any Nash equilibrium is a set of fixed points of conditional winning probability functions.

Lemma 3.4. *In any pure Nash equilibrium $\{r^*_{ij} : i, j \in [k]\}$, we have*

$$f_j^{(i)}(r^*_{ij}) = r^*_{ij},\ \forall i, j \in [k].$$

The second property of Nash equilibria is that they are symmetric strategies.

Theorem 3.5. *Any pure Nash equilibrium $\{r^*_{ij} : i, j \in [k]\}$ is symmetric in the sense that $\forall i, j \in [k]$, we have $r^*_{ij} = r^*_j$ for some r^*_j.*

Theorem 3.5 marks a notable difference between the edge arrival model and the previously studied vertex arrival model. More specifically, pure Nash equilibria in the vertex arrival setting, regardless of what the tie-breaking rule is, are always asymmetric (Sakaguchi 1995; Fushimi 1981; Immorlica, Kleinberg, and Mahdian 2006; Karlin and Lei 2015), whereas in the adaptive edge arrival model we only have symmetric Nash equilibria.

3.1 Recursive Relation for Symmetric Strategies

Given Theorem 3.5, we can focus on symmetric strategies and study Nash equilibria of the game. As all players use the same strategy, their conditional winning probabilities $f_j^{(i)}(r)$ are also the same function (in time r). For this reason, we will drop the superscript (i) in $f_j^{(i)}(r)$, and write $f_j(r)$ only.

First we give a recursion of conditional winning probability $f_k(r)$.

Lemma 3.6. *The conditional winning probability in a k-player hiring process $f_k(r)$ has the following recursive relation*

$$f_1(r) = -r \ln r,$$
$$f_k(r) = c_k r^k - r^k(k-1) \int \frac{1-r}{r^{k+1}} f_{k-1}(r) dr + \frac{r}{k-1}, \quad \forall k \geq 2, \forall r \geq r_k \quad (1)$$

where c_k is a constant of integration.

The idea of the proof is we can consider player i, time r, and an infinitesimal interval $[r, r + dr]$. Under the condition that player i has not hired by time r, we can analyze several different scenarios and sum up their contributions to the winning probability. The detailed proof can be found in the the Supplementary Material.

Next we take a closer look at this condition winning probability function $f_k(r)$. It takes different forms in different ranges $[r_k, r_{k-1}], \ldots, [r_2, r_1], [r_1, 1]$. Denote by $f_{k,j}(r)$ the section in $[r_j, r_{j-1}]$, namely $f_{k,j}(r) = f_k(r)$ for $r \in [r_j, r_{j-1}]$. The general form is hard to compute analytically for general j, but the two ends have a simple form.

Lemma 3.7. *For any symmetric strategy with thresholds $r_1, \ldots, r_k$, the conditional winning probability $f_k(r)$ of each player, in the range $r \geq r_1$ and $r \leq r_{k-1}$, is given as follows.*

1. *When $r \geq r_1$: $f_k(r) = f_{k,1}(r) = -(\frac{1}{k} r^k \ln^k(r) + r^{k-1} \ln^{k-1}(r) + \cdots + r \ln(r))$.*
2. *When $r \in [r_k, r_{k-1}]$: $f_k(r) = f_{k,k}(r) = c_k \cdot r^k + \left(\frac{1}{k-1} - f_{k-1}(r_{k-1})\right) \cdot r + \frac{k-1}{k} f_{k-1}(r_{k-1})$.*
3. *When $r \in [0, r_k]$: $f_k(r) = f_{k,k}(r_k)$.*

Now we are ready to analyze the optimal thresholds. There is no simple and closed-form formula for each r^*_k. However, it turns out it can be defined and computed by a recursive equation.

Theorem 3.8. *The optimal k-th threshold r^*_k is the unique solution to the following equation*

$$c_k \cdot r^k - \left(\frac{k-2}{k-1} + r^*_{k-1}\right) \cdot r + \frac{k-1}{k} r^*_{k-1} = 0.$$

Proof. First, we can derive the following conditional winning probability function $f_{k,k}(r)$ in range $[r^*_k, r^*_{k-1}]$ from Lemma 3.7. Now by Lemma 3.4, let us further set $f_{k,k}(r) = r$, and we have

$$f_{k,k}(r) - r = c_k \cdot r^k - \left(\frac{k-2}{k-1} + r^*_{k-1}\right) \cdot r + \frac{k-1}{k} r^*_{k-1} = 0,$$

the solution of which is r^*_k.

It remains to show that this equation has a unique solution. We let $g(r) = f_{k,k}(r) - r$, and take its derivative,

$$g'(r) = c_k \cdot k \cdot r^{k-1} - \left(\frac{k-2}{k-1} + r^*_{k-1}\right).$$

Let r_k be the smallest solution to $g(r) = 0$. We can then simplify $g'(r_k)$ as

$$\left(\frac{k-2}{k-1} + r^*_{k-1}\right) \cdot k - \frac{k-1}{r_k} r^*_{k-1} - \left(\frac{k-2}{k-1} + r^*_{k-1}\right)$$
$$= (k-1) \left(\frac{k-2}{k-1} + r^*_{k-1} - \frac{r^*_{k-1}}{r_k}\right)$$

Next we consider two cases on the sign of c_k.

- If $c_k < 0$, $g'(r)$ is decreasing in r, thus it suffices to show that $g'(r_k) < 0$. Again from $g(r_k) = 0$ and $c_k < 0$, we have $\left(\frac{k-2}{k-1} + r^*_{k-1}\right) r_k - \frac{k-1}{k} r^*_{k-1} = c_k \cdot r_k^k < 0$, which implies $\left(\frac{k-2}{k-1} + r^*_{k-1}\right) < \frac{k-1}{k} \cdot \frac{r^*_{k-1}}{r_k}$. Plug this inequality into $g'(r_k)$, we therefore have

$$g'(r_k) \le (k-1)\left(\frac{k-1}{k} \cdot \frac{r^*_{k-1}}{r_k} - \frac{r^*_{k-1}}{r_k}\right) < 0$$

- If $c_k \ge 0$, $g'(r)$ is (weakly) increasing in r. In this case, if $g(r) = 0$ has another solution $r' \in [r^*_k, r^*_{k-1}]$, it must be that at point r^*_{k-1}, we have $g(r^*_{k-1}) = f_{k,k}(r^*_{k-1}) - r^*_{k-1} \ge 0$. However, we also know that $f_{k,k}(r^*_{k-1}) < f_{k-1,k}(r^*_{k-1}) = r^*_{k-1}$ by the definition of function f. This gives a contradiction. Therefore $g(r) = 0$ cannot have another solution.

□

Note that in a k-player game, we can use above theorem to compute all optimal thresholds $r^*_1, r^*_2, \ldots, r^*_i$, where r^*_i is the solution to the equation

$$c_i \cdot r^i - \left(\frac{i-2}{i-1} + r^*_{i-1}\right) \cdot r + \frac{i-1}{i} r^*_{i-1} = 0. \quad (2)$$

3.2 Computing Conditional Probability and Nash Equilibria

Theorem 3.8 allows us to compute the optimal threshold r^*_k for each k iteratively. When n is small, these optimal thresholds can be computed by hand. We list the first three optimal thresholds as follows.

Proposition 3.9. *The first three optimal thresholds are*

$$r^*_1 = e^{-1} \approx 0.36788$$
$$r^*_2 = \frac{-1 + \sqrt{1 - e + e^2}}{(e-1)e} \approx 0.29574$$
$$r^*_3 \approx 0.240061.$$

The details of how they are obtained can be found in the the Supplementary Material. However, as we can see, the recursion is quite tedious and messy to calculate even for small k. Therefore, it is desirable to compute them via an algorithm. In the following, we present an efficient algorithm to compute $f_k(r)$ and r^*_k for general k.

Theorem 3.10. *On any given input threshold $r_1, \ldots, r_k$, Algorithm 1 computes all functions $\{f_1(r), \ldots, f_k(r)\}$ in time $O(k^5)$. It also computes the optimal thresholds $r^*_1, \ldots, r^*_k$ in time $\sum_{i=1}^k T_i$, where T_i is the time to solve Equation* (2).

3.3 Best Response of General Strategies

In this section we consider the winning probabilities of and the best response to general strategies. It turns out that this is nontrivial even for two players. We will further fix r_{11} and r_{21} to be their best choice, namely $1/e$, and consider the freedom of the other two thresholds r_{12} and r_{22}, the two players' thresholds when both players are still in the game.

Algorithm 1: Algorithm for computing continuous winning probability in a k-player hiring process

Input: $k \in Z, 0 < r_k \le \cdots \le r_1 \le 1$
Output: $f_{i,j}(r)$ as a function of $r \in [r_{j-1}, r_j]$, for $i, j \in [k]$.

```
r_0 = 1
f_{1,1}(r) := -r ln r
for j = 2 to k do
    f_{1,j}(r) := f_{1,1}(r_1)
end for
for i = 2 to k do
    f_{i,1}(1) := 0
    for j = 1 to i do
        t_{i,j}(r) := (i-1) ∫ (1-r)/r^{i+1} f_{i-1,j}(r) dr
        c_{i,j} := f_{i,j-1}(r_{j-1}) · r_{j-1}^{-i} - 1/(i-1) · r_{j-1}^{1-i} + t_{i,j}(r_{j-1})
        f_{i,j}(r_j) := c_{i,j} r_j^i + r_j/(i-1) - t_{i,j}(r_j) · r_j^i
        f_{i,j}(r) := c_{i,j} r^i - t_{i,j}(r) · r^i + 1/(i-1) · r
    end for
    for j = i+1 to k do
        f_{i,j}(r) := f_{i,i}(r_i)
    end for
    Solve Equation (2) with c_i = c_{i,i} to get r*_i
end for
Output {f_{i,j}(r) : i, j ∈ [k]}
```

For notational convenience, let us denote $x = r_{12}$ and $y = r_{22}$. Without loss of generality, assume $x \le y \le 1/e$ (by Lemma 3.2 and 3.3).

We can compute the winning probabilities for any joint strategy (x, y).

Lemma 3.11. *When using the thresholds $(r_{11}, r_{12}, r_{21}, r_{22}) = (1/e, x, 1/e, y)$ with $x \le y$, the two players' winning probabilities are*

$$\Pr[\textit{P1 wins}] = \frac{1-e}{2}xy + \left(1 - \frac{1}{e}\right)x + \frac{x}{2ey} + x \ln\frac{y}{x} \quad (3)$$

$$\Pr[\textit{P2 wins}] = \frac{1-e}{2}xy + \left(1 - \frac{1}{e}\right)x - \frac{x}{2ey} + \frac{1}{e} + \frac{x}{e}\ln\frac{x}{y} \quad (4)$$

Using the above lemma, we can give the best responses x^* to y and y^* to x of the two players in the complete region of $(x, y) \in [0, 1/e] \times [0, 1/e]$.

Theorem 3.12. *Under the constraint of $x \le y$ and using the thresholds $(r_{11}, r_{12}, r_{21}, r_{22}) = (1/e, x, 1/e, y)$, the best response for player 1 is*

$$x^* = \begin{cases} y & \text{if } y \in (0, r^*], \\ y \cdot \exp(\frac{1-e}{2}y + \frac{1}{2ey} - \frac{1}{e}) & \text{if } y \in [r^*, 1/e], \end{cases}$$

where $r^* = \frac{-1+\sqrt{1-e+e^2}}{(e-1)e} \approx 0.29574$. The best response for player 2 is

$$y^* = \max\{r^*, x\} = \begin{cases} r^* & \text{if } x \in (0, r^*], \\ x & \text{if } x \in [r^*, 1/e]. \end{cases}$$

Proof. By Lemma 3.11, the best response of player 1 is the $x^* = \text{argmax}_{x \le y} f_2^{(1)}(x)$.

$$\frac{\partial}{\partial x}\Pr[\text{Player 1 wins}] = -\frac{1}{e} + \frac{1}{2ey} + \frac{1}{2}(1-e)y + \ln\frac{y}{x}$$

When $y \le r^* = \frac{-1+\sqrt{1-e+e^2}}{(e-1)e} \approx 0.29574$, this derivative is at least $-\frac{1}{e} + \frac{1}{2ey} + \frac{1}{2}(1-e)y$ which is always non-negative. Thus the optimal $x = y$ for y in this range. For $y \in [0.29574, 1/e]$, the optimal

$$\begin{aligned} x &= \exp\left(-\frac{1}{e} + \frac{1}{2ey} + \frac{1}{2}(1-e)y + \ln y\right) \\ &= y \cdot \exp\left(-\frac{1}{e} + \frac{1}{2ey} + \frac{1}{2}(1-e)y\right). \end{aligned}$$

The best response of player 2 in the range $[x, 1/e]$ is the $y^* = \text{argmax}_{x \le y \le 1/e} \Pr[\text{Player 2 wins}]$.

$$\begin{aligned} &\frac{\partial}{\partial y}\Pr[\text{Player 2 wins}] \\ &= \frac{1-e}{2}x + \frac{x}{2ey^2} - \frac{x}{ey} = x\left(\frac{1-e}{2} + \frac{1}{2ey^2} - \frac{1}{ey}\right) \end{aligned}$$

which is positive when $y \in [0, r^*]$ and negative when $y > r^*$. Thus the best response for player 2 is

$$y^* = \max\{r^*, x\} = \begin{cases} r^* & \text{if } x \in (0, r^*], \\ x & \text{if } x \in [r^*, 1/e]. \end{cases}$$

□

With the best response at hand, one can easily compute the unique Nash equilibrium in this two-player game, which will be identical to the one found in Proposition 3.9.

In addition, if we define the social welfare of a strategy profile in this secretary game as the overall probability that the best candidate is hired (by any player), then we can also compute the social welfare loss of the Nash equilibrium compare to the maximum social welfare of any strategy profile. This is usually termed the *price of anarchy* of the game.

Note that the social welfare of a strategy profile $(1/e, x, 1/e, y)$ is

$$\begin{aligned} &\Pr[\text{Player 1 wins}] + \Pr[\text{Player 2 wins}] \\ &= (1-e)xy + \left(1 - \frac{1}{e}\right)x\left(2 + \ln\frac{y}{x}\right) + \frac{1}{e} \end{aligned}$$

which has the maximum value of $\frac{2}{e} - \frac{1}{e^2}$, achieved at $x = y = 1/e$. The social welfare of Nash equilibrium $x = y = r^*$ is $2r^* \approx 0.59148$, thus the price of anarchy is $\frac{2r^*}{\frac{2}{e} - \frac{1}{e^2}} = \frac{2e(\sqrt{1-e+e^2}-1)}{(e-1)(2e-1)} \approx 0.98511$.

4 Non-Adaptive Strategies

In this section we turn our focus to agents with non-adaptive strategies. That is, we assume that players are unaware whether and when other players have made their hiring decisions and have to each stick to a single threshold strategy throughout the game. We begin with the 2-player case.

4.1 Two Players: Winning Probability and Nash Equilibria

Assume that player 1 and 2 use the non-adaptive strategy with thresholds s and r, respectively. We denote by $f_i(s, r)$ the winning probability of player i. We first characterize the winning probability of the two players.

Lemma 4.1. *Assume player 1 and 2 use non-adaptive strategies s and r, respectively. Then the winning probability of player 1 is*

$$f_1(s, r) = \begin{cases} s\ln\frac{1}{s} - \frac{1}{2}sr\ln^2\frac{1}{r}, & \textit{if } s \le r, \\ \frac{1}{2}sr\ln^2\frac{1}{s} + \left(1 - r\ln\frac{1}{r}\right)s\ln\frac{1}{s}, & \textit{if } s \ge r. \end{cases} \tag{5}$$

Note that because the players' identity do not matter, we can also use this lemma to compute player 2's winning probability as $f_2(s, r) = f_1(r, s)$.

From Lemma 4.1 we can pin down all the Nash equilibria in this two-player game, and it turns out that there is a unique Nash equilibrium, which is symmetric.

Theorem 4.2. *The two-player secretary game has only one non-adaptive Nash equilibrium (s^*, r^*), in which $s^* = r^* = 0.29533...$ is the unique solution to the equation*

$$-\frac{x}{2}\ln^2\frac{1}{x} + \ln\frac{1}{x} - 1 = 0.$$

Proof. Without loss of generality let us assume that $s^* \le r^*$. Then by Lemma 4.1, we have

$$f_1(s, r^*) = s\ln\frac{1}{s} - \frac{1}{2}sr^*\ln^2\frac{1}{r^*}, \text{ when } s \le r^*, \tag{6}$$

$$f_2(s^*, r) = f_1(r, s^*) = \frac{1}{2}s^*r\ln^2\frac{1}{r} + \left(1 - s^*\ln\frac{1}{s^*}\right)r\ln\frac{1}{r}, \text{ when } r \ge s^*. \tag{7}$$

Since s^* is a best response to r^*, we know that $s^* \in \text{argmax}_{0 \le s \le r^*} f_1(s, r^*)$. As f_1 is smooth in its domain, s^* is either at one of the two ends (i.e. 0 and r^*), or $\frac{\partial}{\partial s}f_1(s, r^*) = 0$. Note that $f_1(0, r^*) = 0$, thus s^* cannot be 0. A similar argument shows that r^* is either s^* or satisfies $\frac{\partial}{\partial r}f_2(s^*, r) = 0$.

If $s^* \ne r^*$, then we have $\frac{\partial}{\partial s}f_1(s, r^*) = \frac{\partial}{\partial r}f_2(s^*, r) = 0$. That is, by Equation (6) and (7),

$$\begin{cases} -\frac{r^*}{2}\ln^2\frac{1}{r^*} + \ln\frac{1}{s^*} - 1 = 0 \\ \frac{s^*}{2}\ln^2 r^* + (s^* - 1 - s^*\ln s^*)\ln r^* - (1 + s^*\ln s^*) = 0 \end{cases}$$

Simple manipulations of this system of equations gives

$$s^* = r^* \quad \text{and} \quad -\frac{s^*}{2}\ln^2\frac{1}{s^*} + \ln\frac{1}{s^*} - 1 = 0,$$

and further solving this gives $s^* = r^* = 0.29533....$ This gives a necessary condition for (s^*, r^*) being a Nash equilibrium. Finally, it is not hard to check that this particular (s^*, r^*) is indeed a Nash equilibrium. Thus it is the only Nash equilibrium. □

4.2 More Than Two Players: First Threshold and Symmetric Equilibria

Next we turn to the case with more than two players. We will first show an interesting fixed-point property of the smallest threshold in any Nash equilibrium. This is a generalization of the phenomena in single-player case that both the optimal threshold and the corresponding winning probability are $1/e$. This phenomenon appears in the vertex-arrival setting in (Immorlica, Kleinberg, and Mahdian 2006) as well (Theorem 1 of that paper), and it is interesting to see that it holds in our edge-arrival setting.

Lemma 4.3. *In any non-adaptive strategy Nash equilibrium $(r_1^*, \ldots, r_k^*)$ with k players, assume that r_i^* is the smallest among $r_1^*, \ldots, r_k^*$. Then the winning probability of player i in this equilibrium is exactly r_i^*.*

We have seen how to compute Nash equilibria in a two-player game in Section 4.1. With $k \geq 3$ players, the computation to find all Nash equilibria becomes too complicated and tedious. Therefore, we only focus on symmetric Nash equilibria, i.e., equilibria in which all players have the same threshold r. Using Lemma 4.3 one can easily characterize the equilibrium threshold value r, as shown in the following theorem.

Theorem 4.4. *In a symmetric Nash equilibrium in the non-adaptive secretary game with k players, the equilibrium threshold $r^*(k)$ is the unique root of the following equation in range $(0, 1/e)$*

$$1 - \left(1 - r \ln \frac{1}{r}\right)^k = k \cdot r. \tag{8}$$

Proof. The equality holds because the two sides are the two ways of calculating the probability that the best candidate is hired by *any* player from this process.

For the left hand side: We compute the hiring probability directly.

$$\begin{aligned} & \Pr\left[\text{the best candidate is hired by any player}\right] \\ = \;& 1 - \Pr\left[\text{the best candidate is hired by no one}\right] \\ = \;& 1 - \prod_{i=1}^{k} \Pr\left[\text{the best candidate is not hired by player } i\right] \\ = \;& 1 - \left(1 - r^*(k) \ln \frac{1}{r^*(k)}\right)^k. \end{aligned}$$

where in the last step we used Lemma 2.2.

For the right hand side: In a symmetric equilibrium, all players win the game with the same probability. Thus by Lemma 4.3, we know that each player wins with probability exactly $r^*(k)$. Note that the events of player i winning, for $i = 1, \ldots, k$, are exclusive. Therefore, the probability that there exists some player successfully hiring the best candidate equals to $k \cdot r^*(k)$. This shows Equation (8).

Next we show the uniqueness of the equation to complete the proof. First $r^*(k) \leq 1/e$ as when there are more players the winning probability of each player drops, thus by Lemma 4.3, $r^*(k)$ also drops with k, and thus is smaller than $r^*(1) = 1/e$. Next we argue that there is a unique solution to Equation (8) in the range $(0, 1/e)$.

It is easy to see that the function

$$g(r) := 1 - \left(1 - r \ln \frac{1}{r}\right)^k - kr \to 0 \text{ when } r \to 0.$$

It is also easily seen that $g(1/e) = 1 - (1-1/e)^k - k/e < 0$, because $(1-1/e)^k > 1 - k/e$. Next we calculate the second derivative and see that

$$\begin{aligned} & \frac{d^2}{dr^2} g(r) \\ & = -(k-1)k(1+\ln r)^2(1+r\ln r)^{k-2} - \frac{k}{r}(1+r\ln r)^{k-1} \\ & < 0 \end{aligned}$$

where we used $1 + r \ln r > 0$ when $r \in (0, 1/e)$. Putting these three properties of $g(r)$ together, we see that there is a unique solution $g(r) = 0$ in range $r \in (0, 1/e)$. □

Properties of the symmetric Nash equilibrium Equation (8) unfortunately does not have a closed-form solution, but we can estimate it to a second-order precision.

Lemma 4.5. *The unique solution $r^*(k)$ of Equation (8) satisfies*

$$\frac{1}{k} - \frac{1}{k^2} < r^*(k) < \frac{1}{k}.$$

From this theorem, the following corollary directly follows.

Corollary 4.6. *In the symmetric Nash equilibrium of the k-player non-adaptive secretary game, the equilibrium threshold $r^*(k)$ approaches to 0 in the order of $\frac{1}{k}$.*

In addition, the probability that the best candidate is not hired by any player in the end is no more than $\frac{1}{k}$.

5 Concluding Remarks and Future Directions

In this paper we study a variant of the online secretary game where a set of employers compete to hire the best candidate from a common pool of candidates. Different from previous works, in our model, the arriving order of the candidates are uniformly random but independent across different employers. We analyze and characterize the best response behavior and Nash equilibria in both the adaptive version and non-adaptive version of the game.

We list some potential questions and directions for future exploration.

1. For two-player nonadaptive strategies, we showed in Theorem 4.2 that there is a unique Nash equilibrium and it is symmetric. Does this generalize to more players?
2. We analyzed the order of r_k^*, the best candidate's unemployment probability, and the Price of Anarchy in the non-adaptive case. Can we also make similar analysis for the adaptive case?

Acknowledgments

This research is supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (MOE2019-T2-1-045).

References

Ajtai, M.; Megiddo, N.; and Waarts, O. 2001. Improved algorithms and analysis for secretary problems and generalizations. *SIAM Journal on Discrete Mathematics*, 14(1): 1–27.

Cownden, D.; and Steinsaltz, D. 2014. Effects of competition in a secretary problem. *Operations Research*, 62(1): 104–113.

Dynkin, E. 1967. Game variant of a problem on optimal stopping. In *Soviet Math. Dokl.*, volume 10, 270–274.

Dynkin, E. B. 1963. The optimum choice of the instant for stopping a Markov process. *Soviet Mathematics*, 4: 627–629.

Ezra, T.; Feldman, M.; and Kupfer, R. 2020. On a Competitive Secretary Problem with Deferred Selections. *arXiv preprint arXiv:2007.07216*.

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

Freeman, P. 1983. The secretary problem and its extensions: A review. *International Statistical Review/Revue Internationale de Statistique*, 189–206.

Fushimi, M. 1981. The secretary problem in a competitive situation. *Journal of the Operations Research Society of Japan*, 24(4): 350–359.

Immorlica, N.; Kleinberg, R.; and Mahdian, M. 2006. Secretary problems with competing employers. In *International Workshop on Internet and Network Economics*, 389–400. Springer.

Karlin, A.; and Lei, E. 2015. On a competitive secretary problem. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 944–950.

Sakaguchi, M. 1995. Optimal Stopping Games: A Review. *Mathematica japonicae*, 42: 343–351.

Szajowski, K. 1993. Double stopping by two decision-makers. *Advances in applied probability*, 438–452.

Szajowski, K. 1995. Optimal stopping of a discrete Markov process by two decision makers. *SIAM journal on control and optimization*, 33(5): 1392–1410.