

Opportunity costs in the game of best choice

Madeline Crews* Brant Jones[†] Kaitlyn Myers[‡] Laura Taalman[†]
Michael Urbanski[§] Breeann Wilson[¶]

Submitted: Nov 17, 2018; Accepted: Feb 19, 2019; Published: Mar 22, 2019

© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

The game of best choice, also known as the secretary problem, is a model for sequential decision making with many variations in the literature. Notably, the classical setup assumes that the sequence of candidate rankings is uniformly distributed over time and that there is no expense associated with the candidate interviews. Here, we weight each ranking permutation according to the position of the best candidate in order to model costs incurred from conducting interviews with candidates that are ultimately not hired. We compare our weighted model with the classical (uniform) model via a limiting process. It turns out that imposing even infinitesimal costs on the interviews results in a probability of success that is about 28%, as opposed to $1/e \approx 37\%$ in the classical case.

Mathematics Subject Classifications: 05A05, 05A16

1 Introduction

The game of best choice, or secretary problem, is a model for sequential decision making. In the simplest variant, an **interviewer** evaluates a pool of N **candidates** one by one. After each interview, the interviewer ranks the current candidate against all of the candidates interviewed so far, and decides whether to **accept** the current current candidate (ending the game) or to **reject** the current candidate (in which case, they cannot be recalled later). The goal of the game is to hire the best candidate out of N . It turns out that the optimal strategy for large N is to reject an initial set of N/e candidates

*Rose-Hulman Institute of Technology, Terre Haute, IN., USA

[†]Department of Mathematics and Statistics, James Madison University, Harrisonburg, VA., USA
[jones3bc, taalma1a]@jmu.edu

[‡]Rowan University, Glassboro, NJ., USA

[§]Western New England University, Springfield, MA., USA

[¶]Whitworth University, Spokane, WA., USA

and hire the next candidate who is better than all of them (or the last candidate if no subsequent candidate is better). The probability of hiring the best candidate out of N with this strategy also approaches $1/e$. See [GM66] for an introduction to these results. Many other variations and some history have been given in [Fer89] and [Fre83].

We model interview orderings as permutations. The permutation π of N is expressed in one-line notation as $[\pi_1\pi_2\cdots\pi_N]$ where the π_i consist of the elements $1, 2, \dots, N$ (so each element appears exactly once). In the best choice game, π_i is the rank of the i th candidate interviewed in reality, where rank N is best and 1 is worst. What the player sees at each step, however, are relative rankings. For example, corresponding to the interview order $\pi = [2516374]$, the player sees the sequence of permutations $1, 12, 231, 2314, 24153, 241536, 2516374$ and must use only this information to determine when to accept a candidate, thereby ending the game.

Let \mathfrak{S}_N be the set of all permutations of size N . Given some statistic $c : \mathfrak{S}_N \rightarrow \mathbb{N}$ and a positive real number θ , we define a discrete probability distribution on \mathfrak{S}_N via

$$f(\pi) = \frac{\theta^{c(\pi)}}{\sum_{\pi \in \mathfrak{S}_N} \theta^{c(\pi)}}.$$

Given a sequence of i distinct integers, we define its **flattening** to be the unique permutation of $\{1, 2, \dots, i\}$ having the same relative order as the sequence. Given a permutation π , define the **i th prefix flattening**, denoted $\pi|_{[i]}$, to be the permutation obtained by flattening the sequence $\pi_1, \pi_2, \dots, \pi_i$. In the **weighted game of best choice**, introduced in [Jon19], some $\pi \in \mathfrak{S}_N$ is chosen randomly, with probability $f(\pi)$, and each prefix flattening $\pi|_{[1]}, \pi|_{[2]}, \dots$ is presented sequentially to the player. If the player stops at value N , they win; otherwise, they lose. We are interested in calculating the win probability, under optimal play, for finite N as well as in the limit as $N \rightarrow \infty$.

In this note, we follow a suggestion by the first author to let $c(\pi)$ be the position of the largest element in π , indexed starting from 0; that is, $c(\pi) = \pi^{-1}(N) - 1$. Equivalently, this is the number of “wasted” interviews required before we can hire the best candidate. Setting $\theta < 1$ has the effect of imposing a multiplicative cost of θ on each wasted interview. For example, the best candidate being hired immediately will contribute $1 = \theta^0$ (before normalization) to the win probability, whereas each failed interview reduces the contribution of an eventually successful hire by a factor of θ . This weighted model is relevant when the interviews themselves are costly, or if time spent interviewing detracts from the time spent working productively such as when the position being filled is only for a limited term or requires a substantial training investment. Also, observe that when $\theta = 1$, we recover the complete uniform distribution on \mathfrak{S}_N , corresponding to the classical model.

We obtain some interesting behavior vis-à-vis the classical model. The optimal strategy is still positional, for which we reject about $0.435/(1 - \theta)$ initial candidates and select the next best candidate. As $N \rightarrow \infty$ and $\theta \rightarrow 1$, however, this strategy succeeds about 28% of the time even though we have a $1/e \approx 37\%$ success rate at $\theta = 1$. That is, the asymptotically optimal strategy does not vary continuously with the parameter θ which seems to limit the durability of any “policy advice” derived from the classical model

(such as e.g. [SV99]). We found a similar discontinuity in the *optimal strategy* for the Mallows model in [Jon19], although the *success probability* there still approached $1/e$. In the present model, both the strategy and probability of success are discontinuous. Evidently, there is a “price” of about 8.6% in the asymptotic success rate for imposing any wasted interview penalty, no matter how small.

Although there is an established “full-information” version of the game in which the player observes values from a given distribution, it seems that only a few papers have considered nonuniform rank distributions for the secretary problem. Pfeifer [Pfe89] considers the case where interview ranks are independent but have cumulative distribution functions containing parameters determined by the interview positions. The paper [RF88] considers an explicit continuous probability distribution that allows for dependencies between nearby arrival ranks via a single parameter. Inspired by approximation theory, the paper [KKN15] studies some general properties of non-uniform rank distributions in the secretary problem. Our work also fits into a recent stream of asymptotic results for random permutations by researchers in algebraic combinatorics such as [MP14, CDE18, ABNP16].

2 The Model

The **left-to-right maxima** in a permutation π consist of elements π_j that are larger in value than every element π_i to the left (i.e. for $i < j$). In the game of best choice, it is never optimal to select a candidate that is not a left-to-right maximum. A **positional** strategy for the game of best choice is one in which the interviewer transitions from rejection to hiring based only on the position of the interview (as opposed to adjusting the transition based on the prefix flattenings that are encountered). More precisely, the interviewer may play the **r -positional** strategy on a permutation π by rejecting candidates $\pi_1, \pi_2, \dots, \pi_r$ and then accepting the next left-to-right maximum thereafter. We say that a particular interview rank order is **r -winnable** if transitioning from rejection to hiring after the r th interview captures the best candidate. For example, 574239618 is r -winnable for $r = 2, 3, 4$, and 5. It is straightforward to verify that a permutation π is r -winnable precisely when position r lies between the last two left-to-right maxima in π .

It follows from the results in [Jon19, Section 3] that the optimal strategy in our game of best choice is positional ¹, and we let

$$W_N(r) = \sum_{r\text{-winnable } \pi \in \mathfrak{S}_N} \theta^{\pi^{-1}(N)-1}.$$

Theorem 1. *We have the recurrence*

$$W_N(r) = (N - 1)W_{N-1}(r) + r(N - 2)! \theta^{N-1}$$

with initial conditions $W_1(0) = 1$ and $W_1(r) = 0$ for all $r \geq 1$.

¹Briefly, our statistic $c(\pi) = \pi^{-1}(N) - 1$ is essentially prefix equivariant [Jon19, Definition 3.2] in the sense that $c(\pi) - c(\sigma_q \cdot \pi) = 0 = c(12 \cdots k) - c(q)$ for all eligible prefixes q . This is enough to obtain the results in [Jon19, Theorem 3.4] and subsequently [Jon19, Theorem 3.7].

Proof. There are two cases for the r -winnable permutations π of N . If N does not lie in the last position, then we may view the initial segment of π uniquely as an r -winnable permutation of $N - 1$ by flattening. Since there are $N - 1$ possible values for the last position, this case contributes $(N - 1)W_{N-1}(r)$ to $W_N(r)$. If N lies in the last position, then π will be winnable if and only if $N - 1$ lies in one of the first r positions of π . For each of these choices, we may permute the remaining entries in $(N - 2)!$ ways, so these contribute $r(N - 2)!\theta^{N-1}$ all together. \square

Corollary 2. *We have*

$$W_N(r) = \begin{cases} (N - 1)! & \text{if } r = 0 \\ (N - 1)! r \sum_{i=r}^{N-1} \frac{\theta^i}{i} & \text{if } 1 \leq r \leq N - 1 \end{cases}$$

Proof. This follows from Theorem 1 by induction. \square

Theorem 3. *Fix some positive $\theta \neq 1$. The probability of winning the game of best choice using the strategy that rejects r initial candidates is*

$$P_r(N, \theta) = \frac{r(1 - \theta) \sum_{i=r}^{N-1} \frac{\theta^i}{i}}{1 - \theta^N}$$

if $r > 0$ and is $\frac{1-\theta}{1-\theta^N}$ if $r = 0$.

Proof. By definition, the probability of winning is

$$\left(\sum_{r\text{-winnable } \pi \in \mathfrak{S}_N} \theta^{c(\pi)} \right) / \left(\sum_{\pi \in \mathfrak{S}_N} \theta^{c(\pi)} \right) = \frac{W_N(r)}{(N - 1)!(1 + \theta + \theta^2 + \dots + \theta^{N-1})}.$$

The result then follows from the previous corollary. \square

3 Results

Suppose now that $\theta < 1$, and take the limit as $N \rightarrow \infty$. Then the probability of success for the strategy that initially rejects r candidates becomes

$$P_r(\theta) = r(1 - \theta) \sum_{i=r}^{\infty} \frac{\theta^i}{i}.$$

We obtain a curve for each nonnegative value of r (interpreting P_0 as $1 - \theta$), the first several of which we have plotted in Figure 1. For each value of θ , one of the curves is maximal, yielding the optimal strategy and probability of success. For example,

$r =$	0	1	2	3
is optimal for $\theta \in$	$(0, 0.6321]$	$[0.6321, 0.7968]$	$[0.7968, 0.8609]$	$[0.8609, 0.8945]$




Figure 1: The first few asymptotic curves P_0, P_1, \dots

Lemma 4. For each i , the intersection of P_{i-1} and P_i coincides with the maximum value of P_i .

Proof. To see this, the derivative of P_r with respect to θ is

$$r(1-\theta) \sum_{i=r}^{\infty} \theta^{i-1} - r \sum_{i=r}^{\infty} \frac{\theta^i}{i} = r(1-\theta) \frac{\theta^{r-1}}{1-\theta} - r \sum_{i=r}^{\infty} \frac{\theta^i}{i} = r \left(\theta^{r-1} - \sum_{i=r}^{\infty} \frac{\theta^i}{i} \right).$$

whereas the successive differences $P_{r-1} - P_r$ are

$$(1-\theta) \left(((r-1) - r) \sum_{i=r}^{\infty} \frac{\theta^i}{i} \right) + (r-1)(1-\theta) \frac{\theta^{r-1}}{r-1} = (1-\theta) \left(\theta^{r-1} - \sum_{i=r}^{\infty} \frac{\theta^i}{i} \right).$$

Hence,

$$P_{r-1} - P_r = \frac{1-\theta}{r} \frac{dP_r}{d\theta}$$

so the successive differences and derivatives have the same zeros. \square

The first intersection occurs at $\theta = 1 - 1/e$ with value $P = 1/e$. Subsequent intersections can be estimated numerically but have no elementary closed form:

r	$\frac{d}{d\theta} P_r$	solution for $\frac{d}{d\theta} P_r = 0$	value of P_r
1	$\ln(-\theta + 1) + 1$	$\theta = 1 - 1/e \approx 0.63212$	$1/e \approx 0.36788$
2	$4\theta + 2 \ln(-\theta + 1)$	$\theta \approx 0.796812$	0.323805
3	$9/2\theta^2 + 3\theta + 3 \ln(-\theta + 1)$	$\theta \approx 0.860917$	0.309256
4	$16/3\theta^3 + 2\theta^2 + 4\theta + 4 \ln(-\theta + 1)$	$\theta \approx 0.894457$	0.302113
5	$25/4\theta^4 + 5/3\theta^3 + 5/2\theta^2 + 5\theta + 5 \ln(-\theta + 1)$	$\theta \approx 0.915009$	0.297883




Figure 2: The maximum (α, β) for $x E_1(x)$

Thus, the optimal strategy and probability of success is given by $P(\theta) = \max_{r \geq 0} P_r(\theta)$, the maximum P_r function in the regime determined by θ . By Lemma 4, $P(\theta)$ is monotonically decreasing and bounded below. Hence, there is a limiting value as $\theta \rightarrow 1$. However, the limit is clearly bounded away from $1/e$, which is the value at $\theta = 1$ according to the classical analysis. Our goal in this section is to determine $\lim_{\theta \rightarrow 1} P(\theta)$ more precisely.

Recall the exponential integral

$$E_1(x) = \int_x^\infty \frac{e^{-t}}{t} dt$$

which we view as a function of a positive real variable x (see e.g. [OLBC10]). This is a standard special function implemented in many mathematical software systems.

For our main result, we consider the maximum value attained by the related function $F(x) = x E_1(x)$ on $(0, \infty)$; see Figure 2 for a plot. Although there is no elementary form for this maximum, it occurs where $E_1(x) = e^{-x}$ so can be estimated numerically to arbitrary precision. Let α and β be defined by $F'(\alpha) = 0$ and $F(\alpha) = \beta$. Then, $\alpha \approx 0.43481821500399293$ and $\beta \approx 0.28149362995691674$.

We are now in a position to give our main result.

Theorem 5. *As θ approaches 1 from the left, the optimal strategy in our asymptotic weighted game of best choice approaches a positional strategy that rejects $\frac{\alpha}{1-\theta}$ initial candidates and selects the next candidate better than all of them. This strategy has a success probability of β .*

Proof. We would like to optimize $P_r(\theta) = r(1-\theta) \sum_{i=r}^\infty \frac{\theta^i}{i}$ for large r and θ chosen appropriately close to 1. We estimate the series by viewing it as a left or right sum for the corresponding integrals:

$$\int_{t=r}^\infty \frac{\theta^t}{t} dt < \sum_{i=r}^\infty \frac{\theta^i}{i} < \int_{t=r}^\infty \frac{\theta^{t-1}}{t-1} dt = \int_{t=r-1}^\infty \frac{\theta^t}{t} dt.$$

Hence, we may approximate $P_r(\theta)$ by $\tilde{P}_r(\theta) = r(1 - \theta) \int_{t=r}^{\infty} \frac{\theta^t}{t} dt$ with error less than

$$r(1 - \theta) \frac{\theta^{r-1}}{r-1} < 4(1 - \theta)\theta^r$$

since the integrand is decreasing, $r - 1 > r/2$, and $\theta > 1/2$.

Next, we change variables from r to $c = (1 - \theta)r$, and from t to $u = (1 - \theta)t$ in the integral. We obtain $du = (1 - \theta) dt$ so

$$\tilde{P}_c(\theta) = c \int_{u=c}^{\infty} \frac{(\theta^{1/(1-\theta)})^u}{u} du.$$

and our error estimate for $|P - \tilde{P}|$ becomes $4\theta^{c/(1-\theta)}(1 - \theta)$.

Now, we are in a position to take the limit as $\theta \rightarrow 1$, using $\lim_{\theta \rightarrow 1} \theta^{1/(1-\theta)} = 1/e$. This forces $P \rightarrow \tilde{P}$ by our error estimate, and

$$\tilde{P}_c \rightarrow c \int_{u=c}^{\infty} \frac{e^{-u}}{u} du = cE_1(c).$$

Optimizing this function for $c \in (0, \infty)$ then determines the asymptotically optimal positional strategy (where we reject $r = \frac{c}{1-\theta}$ initial candidates) and probability of success. \square

We can also solve the model when $\theta > 1$. One interpretation here is that there is some “trend” in the candidate pool (e.g. due to changes in general economic conditions such as unemployment or interest rates) that is amplifying the probability of seeing the best candidate later. Once again, we find that including even an infinitesimal trend completely changes the optimal asymptotic strategy.

Theorem 6. *If $\theta > 1$, the probability of success for the strategy that initially rejects $r = \frac{N}{\lambda}$ candidates approaches $\frac{1}{\lambda}$, as $N \rightarrow \infty$. Hence, the asymptotic model does not depend on θ .*

Proof. Recall Theorem 3; we claim

$$\lim_{N \rightarrow \infty} P_{N/\lambda}(N, \theta) = \lim_{N \rightarrow \infty} \frac{N(1 - \theta)}{\lambda(1 - \theta^N)} \sum_{i=N/\lambda}^{N-1} \frac{\theta^i}{i} = \frac{1}{\lambda}.$$

To see why, consider the “almost telescoping” sum

$$(1 - \theta) \sum_{i=N/\lambda}^{N-1} \frac{\theta^i}{i} = \frac{1}{N/\lambda} \theta^{N/\lambda} - \frac{1}{N-1} \theta^N - \sum_{i=(N/\lambda)+1}^{N-1} \frac{1}{i(i-1)} \theta^i.$$

If we divide by $-\theta^N$ and take the limit $N \rightarrow \infty$ term by term with $\theta > 1$, we find that only the leading (i.e. middle) term survives. Hence, our limit is $\lim_{N \rightarrow \infty} \frac{N}{\lambda(1 - \theta^N)} \left(-\frac{1}{N-1} \theta^N\right) = \frac{1}{\lambda}$. \square

Thus, for large N we find that it is optimal to choose the last candidate (and win almost all of the time!), obtaining another discontinuity with the $\theta < 1$ and classical models.

Note added in proof

Although we were unaware of it until after our work was accepted for publication, the paper [Ras75] solves a very similar problem to the one we are considering. More specifically, we compute in Theorem 3 the probability of winning the game of best choice under a non-uniform distribution whereas Rasmussen–Pliska compute in their Equation (2.6) the expected value of a random variable representing the non-uniform payoff for the game played on a uniform distribution. For any particular N and θ , these problems are dual to each other in the sense that their corresponding formulas are off by the multiplicative constant $\frac{\theta + \theta^2 + \theta^3 + \dots + \theta^N}{N}$. Since our weights form a probability distribution, we believe our model facilitates a clearer comparison with the classical secretary problem.

Rasmussen–Pliska obtain an asymptotic estimate for the optimal strategy that agrees with ours, using different methods. They also note that, for fixed θ , their expected payoff tends to 0 as N tends to infinity (because their denominator does not scale with θ). By contrast, our model has nonzero probabilities, as N tends to infinity, given by the value of $P_r(\theta)$ where r is optimal for the fixed θ .

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Number 1560151. It was initiated during the summer of 2018 at a research experience for undergraduates (REU) site at James Madison University mentored by the second and fourth authors.

References

- [ABNP16] Nicolas Auger, Mathilde Bouvel, Cyril Nicaud, and Carine Pivoteau. Analysis of algorithms for permutations biased by their number of records. In *Proceedings of the 27th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms—AofA’16*, page 12. Jagiellonian Univ., Dep. Theor. Comput. Sci., Kraków, 2016.
- [CDE18] Harry Crane, Stephen DeSalvo, and Sergi Elizalde. The probability of avoiding consecutive patterns in the Mallows distribution. *Random Structures Algorithms*, 53(3):417–447, 2018.
- [Fer89] Thomas S. Ferguson. Who solved the secretary problem? *Statist. Sci.*, 4(3):282–296, 1989. With comments and a rejoinder by the author.
- [Fre83] P. R. Freeman. The secretary problem and its extensions: a review. *Internat. Statist. Rev.*, 51(2):189–206, 1983.
- [GM66] John P. Gilbert and Frederick Mosteller. Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.*, 61:35–73, 1966.
- [Jon19] Brant Jones. Weighted games of best choice. [arXiv:1902.10163](https://arxiv.org/abs/1902.10163), 2019.

- [KKN15] Thomas Kesselheim, Robert Kleinberg, and Rad Niazadeh. Secretary problems with non-uniform arrival order. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 879–888. ACM, New York, 2015.
- [MP14] Sam Miner and Igor Pak. The shape of random pattern-avoiding permutations. *Adv. in Appl. Math.*, 55:86–130, 2014.
- [OLBC10] Frank W. J. Olver, Daniel W. Lozier, Ronald F. Boisvert, and Charles W. Clark, editors. *NIST handbook of mathematical functions*. U.S. Department of Commerce, National Institute of Standards and Technology, Washington, DC; Cambridge University Press, Cambridge, 2010. With 1 CD-ROM (Windows, Macintosh and UNIX).
- [Pfe89] Dietmar Pfeifer. Extremal processes, secretary problems and the $1/e$ law. *J. Appl. Probab.*, 26(4):722–733, 1989.
- [Ras75] W. T. Rasmussen and S. R. Pliska. Choosing the maximum from a sequence with a discount function. *Appl. Math. Optimization*, 2:279–289, 1975.
- [RF88] J. H. Reeves and V. F. Flack. A generalization of the classical secretary problem: dependent arrival sequences. *J. Appl. Probab.*, 25(1):97–105, 1988.
- [SV99] Dimitris A. Sardelis and Theodoros M. Valahas. Decision making: a golden rule. *Amer. Math. Monthly*, 106(3):215–226, 1999.