

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Rapoport, Amnon; Seale, Darryl A.; Spiliopoulos, Leonidas

Article — Published Version Progressive stopping heuristics that excel in individual and competitive sequential search

Theory and Decision

Provided in Cooperation with: Springer Nature

Suggested Citation: Rapoport, Amnon; Seale, Darryl A.; Spiliopoulos, Leonidas (2022) : Progressive stopping heuristics that excel in individual and competitive sequential search, Theory and Decision, ISSN 1573-7187, Springer US, New York, NY, Vol. 94, Iss. 1, pp. 135-165, https://doi.org/10.1007/s11238-022-09881-0

This Version is available at: https://hdl.handle.net/10419/308734

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.





Progressive stopping heuristics that excel in individual and competitive sequential search

Amnon Rapoport¹ · Darryl A. Seale² · Leonidas Spiliopoulos³ ()

Accepted: 4 February 2022 / Published online: 10 March 2022 $\ensuremath{\mathbb{C}}$ The Author(s) 2022

Abstract

We study the performance of heuristics relative to the performance of optimal solutions in the rich domain of sequential search, where the decision to stop the search depends only on the applicant's relative rank. Considering multiple variants of the secretary problem, that vary from one another in their formulation and method of solution, we find that descriptive heuristics perform well only when the optimal solution prescribes a single threshold value. We show that a computational heuristic originally proposed as an approximate solution to a single variant of the secretary problem performs equally well in many other variants where the optimal solution prescribes multiple threshold values that gradually relax the criterion for stopping the search. Finally, we propose a new heuristic with near optimal performance in a competitive or strategic variant of the secretary problem with one another to hire job applicants. Both heuristics share a simple computational component: the ratio of the number of interviewed applicants to the number of those remaining to be searched. We present the subgame-perfect Nash equilibrium for this competitive variant and an algorithm for its computation.

Keywords Simple heuristics · Individual and competitive secretary problems · Computational models · Descriptive models

Leonidas Spiliopoulos spiliopoulos@mpib-berlin.mpg.de

> Amnon Rapoport amnon@arizona.edu

Darryl A. Seale darryl.seale@unlv.edu

- ¹ University of Arizona, Tucson, USA
- ² University of Nevada Las Vegas, Las Vegas, USA
- ³ Max Planck Institute for Human Development, Berlin, Germany

1 Introduction

The heuristics literature in cognitive psychology is concerned with two different but related issues. First, are heuristics predictive of human behavior? Second, how does heuristic performance compare to the optimal solutions of the decision-making tasks? When testing models in the controlled environment of the laboratory, the focus of the literature typically is on the former question, their predictiveness, namely, how well the predictions of the model match systematic and replicable patterns of behavior that are observed in the raw data. In contrast, our manuscript focuses on the second question, namely, the relative performance of heuristics (simple rule-of-thumb models) vis-à-vis an optimal benchmark.

One viewpoint emphasizes the performance shortcomings of descriptive heuristics, arguing that they often lead to biases and sub-optimal behavior (e.g., Kahneman, 2003a, b, 2011; Tversky & Kahneman, 1974). An alternative viewpoint contends that heuristics can be fast and frugal, exhibiting excellent performance and even outperforming normative models in environments of irreducible uncertainty arising from nature (Gigerenzer et al., 1999, 2011; Hertwig et al., 2013; Todd et al., 2012) or imperfect knowledge of opponents' strategic behavior and payoffs in games (Spiliopoulos & Hertwig, 2020). This latter strand of the literature emphasizes that heuristics perform best when their processes match relevant characteristics of the environment, thereby exploiting them efficiently in the spirit of Simon's 'scissors' metaphor, particularly in the face of uncertainty—see Hertwig et al. (2019) for examples of this across a wide range of different decision-making tasks.

Ultimately, moving away from a binary antagonistic stance regarding heuristic performance, research should be directed at demarcating the boundary between environments where heuristics perform well and environments where they may be inappropriate. Where does this boundary lie? Are heuristics restricted to performing well only in relatively simple environments or do they excel in exactly the opposite type, i.e., complex environments? In choosing a new domain of complex problems to test heuristics, one restriction is that some problems are intractable within reasonable time and computational limits (e.g., NP-hard problems such as the Traveling Salesman problem). Consequently, the optimal solutions are unknown, rendering relative performance an undefinable property.

We have opted to study the taxing, yet numerically solvable, domain of sequential search problems, which sometimes are dubbed as optimal stopping problems (Ferguson, 2002), where the decision maker (DM) interviews the choice alternatives (items, applicants) sequentially. There is a very large body of theoretical research on the standard secretary problem, which originated about 60 years ago (Chow et al., 1964; Lindley, 1961). We narrow down this domain to a class of no-information sequential search problems (Ferguson, 2002) in which the DM has no prior knowledge about the probability distribution of the choice alternatives and, therefore, cannot acquire information from experience about its parameter values. Rather, in deciding whether to accept or reject a given item, the DM is only informed of the relative rank of the present item in comparison to all the

items preceding it in the sequence. The optimal stopping rules for these problems may, in principle, be numerically calculated by dynamic programming (see, e.g., Chow et al., 1964; Gilbert & Mosteller, 1966). As Simon (1990) has noted, optimal strategies provide important insights into the nature of the decision-making environment under investigation, but the mathematical methods for discovering them often use a formal language which is alien to most DMs and at times incomprehensible. Consequently, optimal strategies are inaccessible to most humans and organizations, both in the laboratory and in practice, because of their formal language, computational costs, and the constraints imposed by the users' cognitive abilities. This begs the question whether there exist simple heuristics that are both accessible to DMs and capable of excellent performance in this domain.

Our paper presents three contributions. First, we investigate the performance of heuristics in a non-competitive decision-making domain, consisting of four problem variants (presented in Sect. 2), that is considerably richer and more complex than the domains already studied. One set consists of three heuristics proposed by Seale and Rapoport (1997, 2000) and further investigated by Stein et al. (2003)-these three heuristics have been found to be used by subjects in experiments. Two of these heuristics perform extremely well in variants of sequential search problems for which the optimal decision rule consists of a single threshold. However, they perform poorly in problem variants whose optimal decision rules call for multiple (rather than a single) threshold. Second, in response to this finding, we investigate for the first time the performance of a heuristic originally proposed only for a specific sequential search problem (the expected rank minimization problem; Krieger & Samuel-Cahn, 2009) in three other variants with multiple thresholds-we call this the Progressive Stopping (PS) heuristic (Sect. 3). We establish that this heuristic achieves more than 95% of the optimal performance across all these variants and for sequences with different numbers of items (Sect. 4). Third, we investigate a competitive secretary problem where multiple employers compete with one another to hire the best job applicant (Sect. 5). We then construct an optimal solution (a subgame-perfect Nash equilibrium) for this game following (and correcting) the proof of Karlin and Lei (2015). We propose a new heuristic related to the PS heuristic—the Inverse Progressive Stopping heuristic (Sect. 5.2)—and show that it exhibits remarkably high performance for the whole set of employers. Readers conversant with the literature on sequential search problems and/or heuristics may skip the following Sects. 1.1 and 1.2, respectively, and continue from Sect. 2.

1.1 A brief literature review of sequential search problems

The class of no-information sequential search decision problems has been studied theoretically in applied probability and operations research (see literature reviews by Ferguson, 1989, 2002; Freeman, 1983; Samuels, 1991) and experimentally in the disciplines of psychology and behavioral economics (see, e.g., Corbin, 1980; Bearden & Rapoport, 2005; Lee, 2006; Mak et al., 2019; Palley & Kremer, 2014; Seale, 1996; Seale & Rapoport, 1997, 2000). Non-distributional models of sequential search can be divided into two streams depending on the objective of

the DM conducting the search. In one stream (e.g., Lindley, 1961), commonly referred to as the probability maximization problem (PMP), or the best choice secretary problem, the DM's objective is to maximize the probability of selecting the best item. This may appear very restrictive; however, in the business world, returns to many investments are increasingly described by an all-or-nothing distribution of returns in the long run, e.g., due to first mover advantage or a superior technology dominating the industry. Or at the very least, there are strongly increasing returns to choosing higher relative ranks; consider a sigmoidal relationship, which approximates an all-or-nothing objective function.

Alternatively, the second stream of research (Chow et al., 1964), which is called the expected rank minimization problem (ERMP), relaxes the all-or-nothing assumption as the DM's objective is assumed to minimize the expected absolute rank of the item selected. Under this objective, the payoffs are either the same (equal weight) or they decrease monotonically in the absolute rank of the selected items; the smaller the rank, the higher the payoff. These two objectives differ from each other in their assumptions about the DM's objective, and so do the methods for computing their optimal solutions.

At first sight, the assumption of knowledge of relative ranks—and ignorance of the distribution of items-may seem overly restrictive, particularly to economists. Traditionally, the economics literature deals with distributional models, which allow for normative solutions derived from applications of Bayesian updating, e.g., work on incomplete information in game theory. By contrast, research in Operations Research and Applied Statistics often consists of non-distributional models, where the utilities of items and their associated probability distribution are unknown-see Bearden and Rapoport (2005) for a more detailed comparison of distributional and non-distributional models of search. Consequently, neither computation of utilities and probabilities is required nor complex Bayesian updating after the presentation and inspection of each item. There are three major problems that severely restrict the applicability of distributional models to sequential search in the field. First, fullinformation solutions are very sensitive to the right extreme tail of the distribution of item valuations leading to non-robustness if the distribution is not known perfectly. In the wild, distributions are almost never presented by description to a decision-maker, at best s/he may approximately learn the distribution through observation. However, the tail ends of distributions consisting of rare events are virtually impossible to learn with any reasonable degree of precision even with large number of observations. Second, Knightian uncertainty is considerable in large worlds, therefore attaching cardinal valuations to items even before the sequential search commences is difficult. Third, and most importantly, distributional models are restricted to a single attribute, whereas non-distributional models of sequential search (e.g., searching sequentially for a date; searching sequentially for an apartment after moving to a new location; attempting to choose the best k, k > 1, proposals which are evaluated sequentially) do not have this restriction.

We argue that the majority of important managerial decisions are based on incomplete information due to the considerable state of flux in economic conditions, the unpredictability of innovation, etc., leading to limited managerial control over organizational processes (March & Simon, 1993). A relevant example would be a venture capital firm, sifting through start-ups with considerably uncertain valuations and uncertain likelihoods of the future states affecting the valuations in an attempt to decide which one to invest in. In short, we believe that non-distributional models are more frequently applicable in practice and more realistically capture the nature of sequential searches than distributional models. The important characteristics of search problems are preserved in non-distributional models that allow for the sequential search of multi-attribute items while by-passing the problem of integrating them into a single value and still allowing for a rich set of variants in terms of alternative objective functions, probabilistic knowledge of the number of items and other assumptions—this will be apparent in Sect. 2 where we describe the different variants that we investigate.

1.2 A brief literature review of heuristics

One viewpoint of the descriptive heuristics literature argues that they often lead to biases and sub-optimal behavior (e.g., Kahneman, 2003a, 2003b, 2011; Tversky & Kahneman, 1974) as they are subject to an effort-performance tradeoff (Payne et al., 1988). While the reduction in effort may outweigh the loss in performance for some applications, the contention is that heuristics necessarily suffer a considerable degradation in performance. This viewpoint has been challenged by the fast and frugal heuristics literature, which argues that heuristics can achieve excellent performance and even outperform normative models in environments of irreducible uncertainty (Gigerenzer et al., 1999). Consequently, an effort-performance tradeoff is not a given; quite the opposite, less can be more. Earlier work investigated fast and frugal heuristics in considerably simpler decision-making tasks with a focus on tasks of inference (Gigerenzer et al., 1999, 2011; Todd et al., 2012), individual decisions under risk and uncertainty (Hertwig et al., 2019; Payne et al., 1988; Thorngate, 1980) and social environments (Hertwig et al., 2013). In parallel research, economists have turned their attention to the axiomatization of heuristics (and, more generally, principles of bounded rationality) from the psychology literature Manzini and Mariotti (2007, 2012a, 2012b, 2014); see also Mandler et al., (2012).

More recent work has extended the domain of inquiry from individual to strategic decision making. Theoretical work by Spiliopoulos and Hertwig (2020) shows that bounded-rational heuristics are more robust than sophisticated decision rules—including the normative Nash equilibrium—to both strategic and payoff uncertainty in one-shot strategic interactions. A burgeoning management literature concludes that managers often employ heuristics (Bingham & Eisenhardt, 2011) and that they can be effective decision-making tools in the face of uncertainty (Artinger et al., 2015). For example, Åstebro and Elhedhli (2006) show that simple heuristics can be more effective than complex regression models in predicting the success of risky ventures. That is, even in complex real world managerial domains, heuristics are not necessarily a second-best solution, particularly in large world environments were uncertainty and noise reign supreme.

2 The set of optimal stopping problems

It is generally agreed that the first statement of the PMP appeared in the 1960 column of Scientific American (Gardner, 1960). Lindley (1961) seems to be the first to have solved the PMP in a scientific publication; his work has been extended by many others (e.g., Freeman, 1983). As stated above, the problem is to establish a stopping rule that determines whether to choose item *i* based only on its relative rank R_i , i.e., its rank among items 1 through *i* (*i* = 1, 2, ..., *n*). The PMP is stated in terms of the assumptions that underlie the sequential search for the best applicant:

Assumption 1 (number of applicants). The number of applicants for employment, n, is finite and known.

Assumption 2 (no. of positions). A single position is available.

Assumption 3 (random order of arrival). The applicants are interviewed one at a time in a random order, where the n! orderings are equally likely.

Assumption 4 (no ties). The decision maker (DM) can rank order all the *n* applicants in terms of their absolute rank, A_i , from best ($A_i = 1$) to worst ($A_i = n$) with no ties.

Assumption 5 (decision rule). On each stage i of the search, the DM is only informed of the relative rank, R_i , of applicant i. Based on this information, the DM either accepts the applicant for the position and thereby ends the search or rejects it and then interviews the next one. If no selection is made prior to the *nth* applicant, then the last applicant must be selected.

Assumption 6 (no recall). Once rejected, an applicant may not be recalled.

Assumption 7 (no refusal). An offer of selection is accepted with certainty by the applicant.

Assumption 8 (payoff function). The DM's objective is to maximize his/her expected payoff: 1, if the best applicant $(A_i = 1)$ is selected, and 0, otherwise $(A_i \neq 1)$.

A note on notation and terminology. The absolute and relative ranks of applicant *i* depend on *n*. Because the value of *n* in sequential search experiments is fixed, and the results do not necessarily generalize to other values of *n*, the superscript *n* is suppressed. In the rest of this section, we use the terms candidate for an item with relative rank $R_i = 1$ and applicant for all relative ranks.

Most of the assumptions stated above have been relaxed in one way or another thereby giving rise to multiple variants of the best choice problem—see early review papers by Gilbert and Mosteller (1966) and Freeman (1983), and subsequently Chun (1998, 2000), Bearden et al. (2005) and Bearden et al. (2006). For example, Presman and Sonin (1972) and Bruss and Samuels (1987) replace Assumption 1 by:

Assumption 1'. The DM is only informed of the probability distribution of the number of items, n.

We consider four typical variants of this standard specification. In the first variant, called PMP-UN (Bruss & Samuels, 1987; Presman & Sonin, 1972), the DM only knows the probability distribution of n. The second variant (Gilbert & Mosteller, 1966; Woryna, 2017; Yeo & Yeo, 1994), called PMP-k, assumes that the

DM's objective is to maximize his/her expected payoff: 1, if any of the *k* best applicants is selected ($A_i = 1, ..., A_i = k$), and 0, otherwise. We will examine the two most-researched cases, namely, k = 2 and k = 5. The standard PMP problem is essentially the special case where k = 1. In the third variant, called ERMP (Chow et al., 1964), the DM's objective is to choose an item that minimizes the expected value of the absolute rank of the item selected (Assumption 8'). In the fourth variant (Gilbert & Mosteller, 1966, Problem 2b), which we term problem PMP-DC, at each stage *i* of the search, the DM is informed of the relative rank R_i of applicant *i*. She is allowed *two* (rather than a single) choices, termed r_1 and r_2 ($r_1 < r_2$); if either of them selects the overall best applicant, then the search stops with a win (success). The first choice is used on the first candidate starting with item r_1 . If it fails, then the second choice is used on the first candidate starting with item r_2 . A detailed comparison of the assumptions of these non-competitive PMP problems can be found in Table 1. Appendix 1 documents the optimal solutions for all these variants.

We investigate one final variant, the strategic PMP-COMP (for competitive), where multiple employers are competing to hire the best out of n applicants—this will be presented in detail later in the text. To the best of our knowledge, this is the first time that heuristic decision rules are examined in the context of competitive secretary problems. This is an important extension, as arguably, many real-world optimal stopping problems exhibit such game theoretic characteristics arising from the competition between employers pursuing a limited number of applicants.

3 The progressive stopping heuristic

The computation of the optimal decision rule problems with multiple thresholds, such as the ERMP, becomes tedious, costly, or highly time consuming (particularly if n is very large). This led KS-C to consider simpler rules: "We consider finding simple stopping rules that perform well in minimizing the sum of the expected value of the absolute ranks of the items selected, when one or more items are desired" (2009, p. 1042). For the case where a *s*ingle item is to be selected, KS-C proposed the following heuristic rule for the ERMP, which we henceforth refer to as the progressive stopping (PS) heuristic:

$$t_n(c) = \inf\{i : R_i \le c\delta\}, \delta = \frac{i}{n+1-i},\tag{1}$$

where $t_n(c)$ is the threshold stopping rule and $c \ge 1$ is a constant. This rule stops and chooses the first applicant with a relative rank of R_i satisfying the above constraint, i.e., meeting the threshold $t_n(c)$, and guarantees that some applicant is always chosen as $Pr(t_n(c) \le n) = 1$. This rule performed very well in the ERMP that KS-C focused on (see Table 13) especially for very high *n*—for which the optimal solution is even more computationally intensive—choosing an expected rank smaller than 4 and achieving 98.5% of the optimal (maximum) performance for n > 100,000. KS-C write: "Hence, when the number of items becomes large, the case where it is hard to implement dynamic programming, is when the simple rule performs almost as well as the optimal rule" (2009, p. 1053). The performance of this simple rule is striking.

		PMP	PMP- UN	PMP- <i>k</i> (2, 5)	ERMP	PMP- DC
Solution: # of	f thresholds	1	1	k	Multiple	2
# of	<i>n</i> is finite	~	~	~	~	~
applicants	Only the probability distribution of <i>n</i> is <i>known</i>	-	~	-	-	-
# of choices		1	1	1	1	2
Payoff function	DM maximizes expected payoff: 1, if the best applicant is selected, and 0, otherwise	~	~	_	~	~
	DM maximizes expected payoff: 1, if any of the k best applicants is selected $(A_i = 1,, A_i = k)$, and 0, otherwise	-	-	~	-	-
	DM minimizes the expected value of the absolute rank of the selected applicant	-	-	-	~	-
Decision rule		*1	*1	*1	*1	*2
# of positions		1	1	1	1	1
Arrivals	Applicants are interviewed one at a time in equally likely	a rand	om orde	r, where the	n! orderin	gs are
No ties	The decision maker (DM) can rank order a rank, A_i , from best ($A_i = 1$) to worst ($A_i = 1$)	ll the <i>n</i> = <i>n</i>) wit	applicat h no ties	nts in terms	of their ab	solute
Decision rule	On each stage <i>i</i> , the DM is only informed o this information, the DM either accepts the the search or rejects it and then interview the <i>n</i> th applicant, then the last applicant of	f the rel ne appli s the ne must be	lative ran cant for xt one. I selected	nk, <i>R_i</i> , of ap the position f no selection	plicant <i>i</i> . B and thereb on is made	ased on by ends prior to
No recall	Once rejected, an applicant may not be rec	alled				

Table 1 The set of individual optimal stopping problems and their assumptions

*1 On each stage *i* of the search, the DM is only informed of the relative rank, R_i , of applicant *i*. Based on this information, the DM either accepts the applicant for the position and thereby ends the search or rejects it and then interviews the next one. If no selection is made prior to the *n*th applicant, then the last applicant must be selected

An offer of selection is accepted with certainty by the applicant

*2 On each stage *i* of the search, the DM is informed of the relative rank R_i of applicant *i*. She is allowed two choices, termed r_1 and r_2 ($r_1 < r_2$); if either of them selects the overall best applicant (with absolute rank 1), then the search stops with a win (success). The first choice is to be used on the first candidate starting with item r_1 . If it fails, then the second choice is to be used on the first candidate starting with item r_2

We conjecture, and subsequently test the hypothesis that the PS heuristic rule performs well in other variants of the secretary problem where the optimal solutions call for multiple decision thresholds.

How likely is it that this heuristic—which was originally proposed as a computational device, not as a heuristic descriptive of human behavior—may be discovered and adopted by decision makers? While the exact quantitative optimal solution is difficult to deduce without formal mathematical training, we posit that qualitative aspects of the optimal solution may be accessible to inexperienced DMs; they may first attempt to deduce basic qualitative features of an optimal solution to a

No refusal

problem and then construct a heuristic out of basic building blocks that satisfies these qualitative requirements. Alternatively, they may learn inductively what works as long as the heuristic is relatively simple. Consider the experimental economics literature on eureka or epiphany moments, where after repeated exposure participants suddenly gain an insight into complicated games, with the effect observable both in choices and response times (Dufwenberg et al., 2010; McKinney & Huyck, 2013; Schotter & Trevino, 2020).

What are the building blocks and insights related to optimal stopping problems? Dealing with insights, it is often easier to first consider the border-cases of a problem—this can be construed as a type of decrease-and-conquer approach (Levitin, 2012). For these problems, they are stopping the search immediately and choosing the first candidate or continuing search until the last candidate. It is immediately apparent that in the first case the probability of success is quite low as no knowledge has been accumulated through search with which to compare and improve on the first candidate-it is in essence, equivalent to a random uninformed choice. In the latter case, since not choosing a candidate leads to a loss, the probability of choosing the last candidate conditional on search continuing that long, should be 1. These two observations imply high thresholds for the beginning of the search and zero thresholds at its end. Alternatively, the DM may realize that there are two opposing forces at work, which imply the same qualitative conclusions. The longer the search, the more likely it is to come across the desired candidates; but continuing the search risks losing the most desirable candidates by passing up the opportunity to hire.

What is a simple way to construct a heuristic satisfying this insight and what building blocks are required? Tracking the ratio of the number of applicants already interviewed to the number of the remaining to be interviewed applicants satisfies these requirements in a frugal and intuitive manner. Specifically, the PS heuristic considers the ratio of two integers as combined in the search depth δ : the number of the present stage (how far you are away from the beginning of the search) divided by the number of stages (plus 1) that remain in the search: $\frac{i}{n+1-i}$. This ratio of two integers is multiplied by the scaling parameter c, which is the only parameter of the heuristic. Because R_i is an integer assuming the values 1, 2, ..., n, the heuristic rule consists of a sequence of *n* thresholds and has the same form as the optimal rule for multiple-threshold problems. As the search progresses (i), the numerator in the ratio $\frac{i}{n+1-i}$ increases and the denominator decreases, leading to an increase in the search depth variable δ , and a gradual relaxation of the criterion for stopping the search. Figure 1 plots the progression of R_i as *i* increases for various values of *c*. Higher values of c lead to a quicker relaxation of the stopping threshold criterion. The curves are convex, i.e., increase at an increasing rate with *i*, as delaying choice quickly increases the probability of passing up the best prospects, and not choosing an item before all *n* items are observed necessarily leads to failure. This conforms with the qualitative aspects of the solutions at the border-cases described above. Note that larger values of *n* reveal the same general patterns but on a different scale. The PS heuristic may thus be also viewed as a satisficing/aspiration rule in the spirit of Simon (1957, p. 263):



"(a) When performance falls short of the level of aspiration, search behavior (particularly search for new alternatives of action) is induced. (b) At the same time, the level of aspiration begins to adjust itself downward until goals reach levels that are practically attainable."

The PS heuristic has all the properties of a fast and frugal heuristic as it ignores a large proportion of information that could be gleaned from extensive search, and has obvious search, stopping, and decision components, as described above.

4 Results for non-competitive optimal stopping problems

The heuristic performance is measured by an approximation score, called the *a*-score, which is defined below in terms of the probability of achieving the objective, called the *probability of success*:

$$a$$
-score = $\frac{\Pr(\text{success achieved by the heuristic rule})}{\Pr(\text{success achieved by the optimal solution})}$.

Before proceeding with an in-depth analysis of the PS heuristic in multiplethreshold sequential search problems, we first summarize our findings on existing heuristics for the PMP, PMP-UN, and PMP-2 problems. The cutoff heuristic, where the DM rejects the first r - 1 applicants, and then selects the first candidate thereafter achieves perfect performance (*a*-score = 1), for the optimal r^* value, in the standard PMP and PMP-UN. However, it (and two other heuristics) suffers a significant degradation in performance in the PMP-2, which calls for two rather than a single threshold. This failure leads us to consider the PS heuristic as an alternative for the class of multiple-threshold search problems. A more detailed discussion of these heuristics and their performance in these problems can be found in Appendix E of the Supplemental Online Material.

Table 2 The approximationscores of the PS heuristic in	Problem	n = 20	<i>n</i> = 50	<i>n</i> = 100
multiple threshold problems	PMP-DC	1	1	1
	PMP-2	0.978	0.981	0.988
	PMP-5	0.999	0.986	0.985
	ERMP	0.995	0.974	0.963

4.1 The progressive stopping heuristic in multiple-threshold search problems

4.1.1 PMP-DC, PMP-5, PMP-2 and ERMP

In assessing the PS heuristic's performance, we searched for the best value of the single parameter (*c*), i.e., the one that maximizes the probability of success. The *a*-score for problem PMP-5 was computed by simulation, but the *a*-scores for problems PMP-2 and PMP-DC were computed directly from Eqs. (4) and (6) in Appendix 1, respectively, as these two equations apply to any values of the thresholds r_1 and r_2 and not only for the optimal values r_1^* and r_2^* .

The results for problems PMP-DC, PMP-2, PMP-5 and ERMP for n = 20, 50, and 100 are summarized in Table 2—more details, such as the optimal *c* values, threshold values, and the probability of winning can be found in Tables 9, 10, 11, 12 and 13 in Appendix 2. A key finding of the present paper is that the PS heuristic performs extremely well across these variants. In PMP-DC, the PS heuristic emulates the optimal solution across all *n* (achieving an *a*-score = 1) for the appropriate values of *c*. In PMP-2 and PMP-5, while not achieving perfection, the PS heuristic achieves exceptionally high *a*-scores ranging from 0.978 to 0.999. In the ERMP, the heuristic achieves *a*-scores ranging from the near-perfect 0.995 for n = 20 to 0.963 for n = 100 and as KS-C showed, still performs well for higher values of *n*: 0.963 for n = 100, 0.981 for n = 1000, 0.984 for n = 10,000 and 0.985 for n > 100,000. In some of these cases, the PS heuristic improves in performance as *n* increases, that is as the complexity of the optimal solution increases this heuristic becomes even more effective—problem complexity is tamed by a simple decision process as embodied in the PS heuristic.

The PS heuristic outperforms other approximate solutions that have been suggested for these problems. Dietz et al. (2011) propose two policies approximating the optimal solution to PMP-2, one of which consists of a single threshold and the other of two thresholds. The former has two free parameters (the relative ranking cutoff to be applied after a position threshold), whereas the latter has four parameters (two rankings and two positions). In this sense, they are considerably more complex that the PS heuristic's single free parameter. Furthermore, deriving the optimal parameters requires maximization over complex functional forms involving obscure—to the uninitiated—combinatorics (Dietz et al., 2011, Eq. 2, p. 160, Eq. 6, p. 164). From their Table 2 (p. 167), we calculate the approximation

score for n = 100 and k = 5 of the single-level policy as 0.894 and of the double-level policy as 0.967.¹ Recall from Table 2 that the relevant *a*-score of the PS heuristic is 0.985; it is superior to the single- and double-level policies both in performance and parsimony.

We are not aware of such close approximations of the optimal solutions by the same simple, single-parameter heuristic to such a wide range of complex decisionmaking problems requiring formal mathematical techniques to solve, which for most people are intractable and incomprehensible. In the next section, we propose a variant of the PS heuristic and show that the impressive performance of this class of heuristics is not limited to individual or non-competitive secretary problems as it also applies to competitive problems.

How robust is the PS heuristic to different values of c and n? Fig. 2 presents the approximation scores associated with different values of these two variables for each of the four problems. A striking result is that for all problems there is very little variation in the approximation score with respect to the number of items *n* conditional on values of *c*. For PMP-DC, the *a*-score is quite insensitive, or robust, to values of c, as the curve is quite flat for a large range of values. This is not the case for the other three problems, where the curve reveals greater sensitivity to c. Let us examine the range of c values for which the a-score is greater than 0.95. For the sake of exposition, we report here the ranges for n = 50 as there is little variation across n. The ranges for problems PMP-DC, ERMP, PMP-2 and PMP-5 are [2.0-6.2], [1.7-2.7], [1.1-2.4] and [1.2-2.6], respectively. Note that c values between 2.0 and 2.4 guarantee *a*-scores greater than 0.95 for all four problems. Consequently, the PS heuristic generalizes admirably, that is, a decision-maker who has learned appropriate c values from experience in one of these problems can transfer this knowledge over to a different type of problems and immediately achieve excellent performance. Furthermore, due to the discreteness of the possible (integer) threshold stopping values, there exist ranges of c values that produce the same optimal thresholds and perfect *a*-scores in problem PMP-DC. In conclusion, the PS heuristic exhibits considerable robustness across different problems and their features-this is a valuable trait for a fast and frugal heuristic.

5 The competitive secretary problem: PMP-COMP

Immorlica et al. (2006) and subsequently Karlin and Lei (2015) consider a new variant of the PMP in which the sequential search for the best applicant is conducted in a competitive setting by multiple DMs. This opens a new horizon for future applications of the secretary problem as a general model of sequential search. As a motivating example, consider the case of k ($k \ge 2$) academic departments sending their faculty representatives, one per department, to a national academic conference with the instructions of interviewing and subsequently hiring a single job applicant

¹ While not formally defined in their paper, we assumed from the name of their metric that the relative (%) errors are computed as $RE = \frac{100|Pr(optimalsuccess) - Pr(policysuccess)|}{Pr(optimalsuccess)}$. Consequently, the a-score can be computed as $1 - \frac{RE}{100}$.



Fig. 2 Robustness of the PS heuristic to values of c and n

for a junior position in their department. Exactly n job applicants attending the conference are ranked from 1 (best) to n (worst) with no ties (same ranking for all the k employers) and arrive at the interview, one at a time, in a random order. The interviewers (hereafter called employers) are also ranked from 1 (best) to k (worst) with no ties in terms of the "quality" of their job offers (e.g., starting salary, teaching load, academic prestige, location of the school, or some combination of the above). Employers are instructed by their institutions to hire the best applicant (Karlin & Lei, 2015).

The assumptions underlying the sequential search for the best applicant in the competitive variant of the PMP problem (hereafter called PMP-COMP) are stated below.

Assumption 1 (no. of applicants). The number of applicants, n, is finite and commonly known.

Assumption 2 (no. of employers). The number of employers, k, is finite and commonly known.

Assumption 3 (priorities). (a) The *n* applicants are ordered in terms of their absolute ranks A_i from best ($A_i = 1$) to worst ($A_i = n$) with no ties. (b) The employers are ranked by the applicants from best (1) to worst (*k*) with no ties. The employers' ordering is publicly known whereas the applicants' ordering is not revealed to the employers.

Assumption 4 (interview). The n applicants are interviewed independently by the k employers, one at a time, in a random order (all n! orderings are equally likely).

Assumption 5 (applicant's decision rule). If applicant i receives multiple job offers, then she accepts the offer from the highest ranked employer among those making the offer. Refusing an offer is not an option, and the terms of the job offer are not negotiable.

Assumption 6 (employer's decision rule). At each stage *i* of the search, at the end of the interview all the *k* employers are informed of the relative rank, R_i , of applicant *i* (*i* = 1, 2, ... *n*). Based on this information, any employer may either accept an applicant *i* (i.e., make her a job offer) or reject her. The *k* (binary) decisions are made independently and are irrevocable. Once an employer hires an applicant, he may submit no further offers. Note, that an employer may make more than one offer, as prior offers may be rejected by applicants accepting offers from higher ranked employers.

Assumption 7 (no recall). Once rejected by all the k employers, the applicant may not be recalled.

Assumption 8 (objective). The objective of each employer is to maximize the probability of hiring the best applicant (payoff = 1 if A_i = 1 and 0, otherwise).

5.1 The optimal solution

Using backwards induction to compute employer i's best response, Karlin and Lei (K&L) constructed the subgame-perfect Nash equilibrium solution to the PMP-COMP problem. The solution also yields the probability of success for each of the k employers. The equilibrium solution has the form of a multi-threshold strategy (t_1, t_2) t_2, \ldots, t_k , where for each value of n and each value of j between 1 and k, there is a unique integer t_i , called the optimal threshold value, such that employer j accepts applicant i if (1) applicant i is a candidate $(R_i = 1)$ and (2) $i \ge t_i$. Thus, the first t_i applicants are rejected by all the employers, in the interval $(t_1 + 1, t_2)$ an applicant is accepted by employer i = 1 if it is a candidate, in interval $(t_2 + 1, t_3)$ employer j = 1 (if he is still active in the game) and employer j = 2 submit simultaneously two job offers to applicant *i* if it is a candidate, and so on. To remain consistent with the K&L notation, the thresholds are defined by the number of applicants to reject, whereas in the previous sections the thresholds, r, were defined as the applicant with which an employer should start submitting offers, i.e., rejecting the first r - 1applicants. In Appendix 3 we provide counter examples for which the algorithm in K&L fails to solve for the true optimal thresholds, as those are reported in Table 1 of K&L.

The structure of the equilibrium proposed by K&L is correct, only the specific values of the thresholds $(t_1, t_2, ..., t_k)$ are problematic. We constructed an algorithm that resolves this inconsistency and returns the true optimal thresholds for all *n*—the reasoning for our proof is identical to that given in K&L with a single modification to their Eq. 3 (2015, p. 946). K&L define the optimal risk, $R_k(i)$, as that belonging to a set of rules *ignoring* the first *i* applicants. Consequently, their Eq. 3 should read $R_k(i-1)$ instead of $R_k(i)$. In Appendix 3, we present the pseudo-code for the algorithm that calculates the subgame-perfect Nash equilibrium. Table 3 reports optimal threshold values using our algorithm for the parameter values $n = \{10, 20, 50, 100, 1, 000, 10, 000\}$ where k = 10. We report identical thresholds as

148

Empl.	<i>n</i> =	20	<i>n</i> =	50	<i>n</i> =	100	n = 1	,000	<i>n</i> = 10	,000
Ĵ	t	Pr(win)	t	Pr(win)	t	Pr(win)	t	Pr(win)	t	Pr(win)
1	7	0.389	18	0.378	37	0.371	368	0.368	3679	0.367
2	4	0.235	11	0.224	22	0.224	223	0.224	2231	0.224
3	2	0.143	7	0.142	14	0.143	141	0.141	1411	0.142
4	1	0.092	4	0.094	9	0.092	91	0.091	910	0.091
5	1	0.063	2	0.059	5	0.060	59	0.060	594	0.060
6	0	0.051	1	0.040	3	0.040	39	0.039	391	0.039
7	0	0.022	1	0.027	2	0.027	25	0.026	259	0.026
8	0	0.005	0	0.020	1	0.017	17	0.017	172	0.017
9	0	0.001	0	0.011	1	0.010	11	0.012	115	0.011
10	0	$0.000^{\rm a}$	0	0.003	0	0.010	7	0.007	77	0.008

Table 3 Optimal threshold values for k employers competing with one another for hiring the best applicant: n = 10, 50, 100, 1,000 and 10,000, and k = 10

^aThe exact value is 0.0002

Table 4 Optimal threshold values (as a proportion of *n*) for employers competing with one another for hiring the best applicant as $n \to \infty$

Employer	Threshold (proportion)	Employer	Threshold (proportion)	Employer	Threshold (proportion)
1	0.367879	13	0.002350	25	0.000022
2	0.223130	14	0.001586	26	0.000015
3	0.141093	15	0.001073	27	0.000010
4	0.091017	16	0.000726	28	0.000007
5	0.059429	17	0.000492	29	0.000004
6	0.039124	18	0.000334	30	0.000003
7	0.025913	19	0.000226	31	0.000002
8	0.017242	20	0.000154	32	0.000001
9	0.011515	21	0.000104	33	0.000000
10	0.007713	22	0.000071	34	0.000000
11	0.005180	23	0.000048	35	0.000000
12	0.003485	24	0.000033		

K&L for the case $n \to \infty$, which we approximated with n = 100,000—see Table 4 for the results as $n \to \infty$ for large number of employers, k. For j = 1, 2, 3, and 4 the top four employers—K&L report the optimal threshold values $e^{-1} \approx 0.368$ (same as the asymptotic threshold value in the basic PMP problem), $e^{-3/2} \approx 0.223$, $e^{-47/24} \approx 0.141$, and $e^{-2761/1152} \approx 0.091$, respectively. In the limit, the thresholds (as proportions of n) are equal to the limiting probabilities of success (Matsui & Ano, 2016).

The algorithm for the optimal solution is calculated recursively starting from the highest-ranked employer, who behaves as if there were no competition whatsoever; that is, the solution is identical to that of the standard best choice problem. Once this is fixed, the employer in the second rank best responds to the behavior of the top-ranked employer but disregards the behavior of all lower ranked employers. Therefore, the solution for each employer *j* is independent of all lower-ranked employers, so that the results for k < 10 are simply the same threshold values in Table 3 truncated at *k*. Similarly, if the optimal threshold is zero for the *j*th employer, then it must also be zero for all employers > *j*.

Before moving on we present an intriguing finding. Consider the extension of the PMP problem where the DM can make *r* choices (r < n), which can be found in Sect. 2c of the paper by Gilbert and Mosteller (1966). Recall that earlier we examined the PMP-DC problem, for which r = 2. The numerical solution for the generalized *r*-choice problem appears in the fifth column of Table 4 by Gilbert and Mosteller (the column titled "P(win)" for the solution with *r* Starting Numbers). Note that when $n \rightarrow \infty$, the solution to the PMP-COMP is identical to the *r*-choice problem solution. This finding can be obtained from our Table 4 by simply adding the probabilities of winning for all employers up to rank *j*—that is, calculating the probability of any employer hiring the best candidate. Consequently, the subgame perfect equilibrium solution for the strategy of each of the competitors has a single threshold, is identical to the strategy of the single player in the *r*-choice problem, which calls for using multiple threshold values in the sequential search for selecting the best applicant.

5.2 The heuristic solution to problem PMP-COMP

How can DMs intuit important insights into this competitive problem? Recall that the optimal solution is structured in terms of multiple threshold values similar to the problems we examined above where we applied the PS heuristic with one important difference: whereas the PS heuristic produces increasing thresholds, the PMP-COMP optimal solution requires thresholds that are decreasing in the (numerical) rank of the employer *j*. For lower-ranked employers to have a chance of hiring they must—on average—make offers earlier than higher-ranked employers to avoid directly competing with them (recall that multiple offers are always resolved in favor of the highest-ranked employer). Armed with the insight that thresholds are decreasing in the (numerical) rank of the employer provides the necessary qualitative characteristics of a solution.

We propose the Inverse Progressive Stopping heuristic (IPS) (Eq. 2), inspired by the original PS heuristic, due to its dependence on easily accessible information to the employer, i.e., the number of applicants already interviewed and the number remaining to be searched:

$$t_j(c) = \begin{cases} \sup\left\{i: R_j \le \frac{1}{c\delta}\right\} & \text{if } \exists i \text{ s.t. } R_j \le \frac{1}{c\delta}, \quad \delta = \frac{i}{n+1-i} & \text{for } c > 0. \end{cases}$$
(2)

In contrast to the original PS heuristic, the right-hand side fraction is now inverted so that it is decreasing in the search depth δ (and by extension, *i*) instead of increasing, and we seek the supremum of *i* for which the inequality holds rather than the infimum. Also, note that the left-hand side of the inequality is the relative rank of employer *j*. Finally, as the indexing of applicants *i* has a lower bound of 1, we added a second possibility which allows for thresholds of zero, i.e., ignoring no applicants and choosing the first one that appears, as is often the case in the optimal solution. If an *i* does not exist for satisfying the first possibility, then the threshold is set to zero.

In deriving the optimal value for the scaling parameter c, we chose a criterion more relevant to the multiple players in the competitive secretary problem than to the individual DM in the non-competitive problems. Let p_j^* be the optimal probability of winning for employer j, and p_j' be the probability of wining according to the heuristic. The latter assumes that all employers are using the heuristic and the same value of c. The optimal value of c minimizes the mean absolute deviation (D) in the optimal and heuristic win probabilities: $D = 1/k \sum_{j=1}^{k} |p_j^* - p_j'|$. The optimization was performed via a two-stage grid search to ensure a global maximum with the first stage in increments of 0.1 for $0.1 \le c \le 10$ and the second with finer increments of 0.05 for $1 \le c \le 5$.

As can been seen in Table 5 (associated optimal *c* values and corresponding threshold can be found in Table D.1 of the online supplemental material), the performance of the IPS heuristic is quite impressive, as *D* varies between 0.0035 and 0.006 for different *n* values. That is, measuring this as the percentage of games an employer is expected to win, the difference between the optimal and heuristic solution leads to less than a 0.6%-point change on average to the win rate. Furthermore, the lowest deviation score *D* (0.0035) is observed for n = 1000, which would be the most computationally expensive for determining the optimal solution. For n = 10, both the optimal and heuristic solutions exhibit a total win probability of 1 (the best applicant will surely be hired by one of the employers). For *n* other than

Table 5 Optimal and IPS heuristic threshold values for		<i>n</i> =				
problem PMP-COMP		10	20	50	100	1000
	Heuristic					
	D	0.00591	0.00601	0.00503	0.00383	0.00348
	P(Twin)	1	0.9534	0.9615	0.9603	0.9594
	Optimal					
	P(Twin)	1	1	0.9989	0.9942	0.9851

10, using the heuristic instead of the optimal solution leads to a small decrease in the total win probability corresponding to a 4-5% point increase in the chances that the best applicant will not be hired at all. However, note that our deviation score criterion *D* that was used to choose the optimal *c* did not specifically maximize the total probability of any employer winning. Directly optimizing for the latter leads to a significant reduction in the gap between the heuristic and the optimal solution—see Table D.2 in Appendix D of the online supplemental material.

Another significant observation is that the optimal c values across different values of n all lie in a narrow range from 1.65 to 2.65 (see also Table 12 in Appendix 2 for problem ERMP)-recall that this overlaps significantly with the range of c [2.0–2.4] that guaranteed high performance in the non-competitive problems in Sect. 4. Furthermore, as is evident from Fig. 3, which shows how the deviation score D varies by c and n, the deviation score is quite insensitive around the optimal values of the scaling factor c = 2. We conclude that the heuristic exhibits excellent performance with similar c values across varying orders of magnitude of n and that it is quite robust to deviations from c around the optimal value. That is, the same heuristic with identical c value could be used by any employer from 1 to k for virtually any range of n with very small performance degradation in terms of the probability of winning. Even without prior experience, it is likely that DMs will still attain high performance, which can be fine-tuned with experience. Note that the parameter c is not some opaque parameter embedded in a complex functional form making it difficult to interpret and adjust; rather, it is a multiplicate scaling factor that adapts the rate of change in the threshold values as the search progresses.

6 Discussion

More than 60 years ago, Simon (1957, 1959, 1982) suggested that while most people strive to make rational choices, their decisions are often subjected to cognitive limitations. People are constrained, he posited, by the amount of information they have at their disposal, the amount of time they need for deliberation before making their decisions, and their previous experience. Inspired in part by Simon, a large body of research on rules-of-thumb, or heuristics, has been growing exponentially in psychology and behavioral economics (e.g., Gigerenzer





et al., 1999, 2011). They conclude that cognitive constraints need not be detrimental to decision-making performance if they are attuned to environmental characteristics, i.e., heuristics may be ecologically rational. In general, behavioral models are often solely judged on their predictiveness, how well they match behavioral data from experiments. The implicit assumption is that we should expect heuristics to be inferior in performance to the normative solutions; therefore, heuristics are valuable only as descriptive models. Consequently, less attention has been directed to the model's relative performance: how close is the performance of the heuristic to the optimal solution. If the insight that the simplicity of the rule-of-thumb and its accuracy are not mutually exclusive, then the relative performance of the rule is a critical construct for validating heuristics not only as descriptive, but also as prescriptive models.

Our results support previous claims in the heuristic literature that there exist simple heuristics that can provide straightforward explanations for solving astonishingly complex decision-making tasks. A synthesis of our findings reveals that existing descriptive heuristics, such as the cutoff and successive non-candidate heuristics (Seale & Rapoport, 1997), should be employed by decision-makers whenever the optimal solution of the problem requires a single threshold. However, the PS and the inverse PS heuristics should be used in multiple-threshold problems, as they achieved near optimal performance across five variants of sequential search problems, including a more complex competitive problem where the search for the best job applicant is conducted sequentially by a group of k interviewers. As in Spiliopoulos and Hertwig (2020), this result pushes back against the philosophical arguments (e.g., StereIny, 2003) that heuristics would not perform well in strategic environments against other humans in contrast to individual decision-making tasks against nature.

Armed with a toolbox of such heuristics, a DM may achieve near optimal solutions with minimal computational demands across a wide domain of problems with varying assumptions, including different objective functions. Our proposed heuristics are "simple" in the sense that they are restricted to include only a single parameter, and they require only counting of items and elementary mathematical operations like maximization, minimization, summation, and division. Furthermore, they do not compute probabilities and utilities, their cost of implementation is negligible, and they employ a minimum of processing time. Notably, both heuristics are built around a simple and intuitive ratio that DMs can track in real timenamely, the ratio of the number of interviewed candidates to the number of those remaining in the sequence. This contrasts sharply with the necessary computations for the optimal solutions (that involve recursive dynamic programming), which suffer from the curse of dimensionality. Furthermore, we have shown that within each type of problem, the optimal value of the only free parameter of these heuristics, c, is relatively invariant to the number of items, n. In the terminology of Gigerenzer et al. (1999), the PS and IPS heuristics can be classified as "fast and frugal" in terms of the information processing they require, storage space, and mental computation.

The performance of these heuristics might be evaluated in comparison to the prolong amount of time and considerable effort in deriving the optimal solutions to these variants of the secretary problem. Experimental studies in which the sequential search problems that we have selected are iterated in time to allow for learning might provide important additional information. In prior studies, Seale and Rapoport (1997) report no evidence of learning, whereas Goldstein et al. (2020) do; note that in the latter study, but not in the former, the distribution of applicants' ability could be learned with experience. While the experimental evidence on learning in repeated secretary games is meager with mixed results, one of the goals of the present paper is to further instigate this line of research for a broader class of problems and heuristics. We have shown that a set of simple heuristics with excellent performance exists. Now, significant efforts must be directed in understanding how decision makers facing optimal sequential search problems arrive at and construct specific heuristics, how successful they are in doing so, and how they learn with repeated exposure to improve upon these or adjust them.

Appendix 1: Optimal solutions to variants of the secretary problem

Problem PMP: the standard secretary problem.

The optimal policy for the standard PMP has the form of a single threshold rule: there exists an integer r^* depending on n. To maximize the probability of selecting the best applicant, ignore the first $r^* - 1$ applicants in the sequence and then choose the first candidate. If no such candidate appears starting with item r^* then the DM incurs the reward 0. The probability t(r, n) of choosing the best applicant out of n, if one passes the first $r^* - 1$ applicants and then chooses the first candidate, if any, is given by

$$t(r, n) = \frac{r-1}{n} \sum_{k=r-1}^{n-1} \frac{1}{k}, 1 < r < n,$$

and the optimum value of r, denoted by r^* , is the smallest value of r that satisfies the inequality

$$\frac{1}{r} + \frac{1}{r+1} + \dots + \frac{1}{n-1} < 1 < \frac{1}{r-1} + \frac{1}{r} + \frac{1}{r+1} + \dots + \frac{1}{n-1}.$$
 (3)

Table 1 (Gilbert & Mosteller, 1966) presents the optimal threshold values r^* and the associated probabilities of choosing the best applicant for n = 10, 20, ..., 100, and 1,000. In the limit, as $n \to \infty$, $t(r, n) = 1/e \approx 0.368$ and $r^* \approx n/e$.

Stein et al. (2003) note that the DM does not need to know the value of r^* in order to perform well. This presumes that the DM uses a single-value threshold policy for some value of r, which includes r^* as a special case (Table 6). In the section on heuristics, we study how the cutoff rule behaves for all values of r, not only the optimal value of r^* .

Table 6	Optimum value	s of r^* for pro	oblem PMP a	nd the associs	ted probabili	ity of choosin	g the best apl	plicant for sel	ected values o	of <i>n</i>		
и	10	20	30	40	50	60	70	80	06	100	1000	8
r^*	4	8	12	16	19	23	27	30	34	38	369	n/e
Pr(win)	0.399	0.384	0.379	0.376	0.374	0.373	0.372	0.372	0.371	0.369	0.368	0.368

Progressive stopping heuristics that excel in individual and...

Problem PMP-UN: the secretary problem with random population size

In this subsection, we consider a variant of the PMP, referred to as PMP-UN, which satisfies Assumptions 1', 2, 3, 4, 5, 6, 7, and 8. In the standard PMP, the DM faces the risk that if she rejects a candidate, then she may discover later that it was the best applicant. In the PMP-UN, the DM is facing an additional risk, namely, that if she rejects a candidate then she may later discover that it was the last applicant in the sequence, in which case her payoff is zero.

Presman and Sonin (1972) show that, in general, the PMP-UN does not have a single-value threshold solution. It does have an optimal solution in three cases where the probability distribution of *n* is either uniform, geometric, or Poisson. The optimal policy in these three special cases assumes the single-value threshold form, where the first $r^* - 1$ applicants are rejected and then the first candidate thereafter, if any, is accepted. In the PMP-UN, if the number of applicants is distributed randomly over the integers in [1, *N*], then the cutoff proportion r^*/n assumes the values 0.20, 0.15, 0.15, and 0.1375 for N = 10, 20, 40, and 80, respectively (see Rasmussen & Robbins, 1975). The respective probabilities of choosing the best applicant out of *n*, if the DM passes the first $r^* - 1$ applicants and chooses the first candidate thereafter, are 0.351, 0.308, 0.289, and 0.279. As $n \to \infty$, $r^*/n \to 1/e^2 = 0.135$, and the probability of success converges to 2/e = 0.2707 (see Table 7).

Problem PMP-2: the secretary problem with either the highest or second highest rank

Gilbert and Mosteller (1966) reported yet another variant of the PMP, referred to as the PMP-2, where the objective of the DM is to choose the applicant with either the highest or second highest absolute rank. It turns out that the optimal policy no longer has the form of a *single-value* threshold rule. Rather, it has the form of a threshold rule with *two* values: reject the first $r_1 - 1$ applicants, then choose the first candidate thereafter, but starting with applicant r_2 choose the first applicant in the sequence with a relative rank of either 1 or 2. This policy relaxes the stopping rule as the search progresses.

The probability of selecting one of the two best applicants (referred to as the probability of success) consists of three parts:

Table 7 Optimum values of r^* for problem PMP-UN and the associated probability of choosing the best applicant for selected values of n

10	20	40	80	1000	∞
0.2	0.15	0.15	0.1375	0.136	$1/e^2 = 0.1353$
0.351	0.308	0.289	0.279	0.271	0.2707
	10 0.2 0.351	10 20 0.2 0.15 0.351 0.308	10 20 40 0.2 0.15 0.15 0.351 0.308 0.289	10 20 40 80 0.2 0.15 0.15 0.1375 0.351 0.308 0.289 0.279	10 20 40 80 1000 0.2 0.15 0.15 0.1375 0.136 0.351 0.308 0.289 0.279 0.271

$$\Pr(\operatorname{success}|r_1, r_2; n) = \frac{r_1 - 1}{n(n-1)} \sum_{i=r_1}^{r_2} \left[\frac{2(n-1)}{i-1} - 1 \right] + \frac{(r_1 - 1)(r_2 - 2)}{n(n-1)} \cdot \sum_{i=r_2+1}^n \left(\frac{1}{i-2} \right) \left[\frac{2(n-1)}{i-1} - 1 \right]$$
(4)

The optimal values of the two threshold values r_1^* and r_2^* ($r_1^* \le r_2^*$) are calculated numerically for relatively small values of *n* by scanning an r_1 by r_2 grid. As $n \to \infty$, Pr(success) $\to 0.574$.

Table 8 presents the optimal rules and the associated probability of success for selected values of n (Gilbert & Mosteller, 1966).

Problem PMP-5: the secretary problem for selecting one of the k best applicants

In this problem, the DM has a single choice for selecting one of the *k* best applicants. We consider the case k = 5. The optimal solution has a threshold form where the *k* values r_1 , r_2 , r_3 , r_4 , and r_5 are computed for each value of *c* in a *k*-dimensional grid search. The optimal values r_1^* , r_2^* , r_3^* , r_4^* , and r_5^* are the ones that maximize the probability of success. The equation of the probability of success by selecting the *a*th best applicant out of *n* appears in Yeo and Yeo (1994, Eq. 4):

$$\Pr\left(\operatorname{success}|r_1^*, r_2^*, r_3^*, r_4^*, r_5^*, n\right) = \frac{1}{(n)\binom{n-1}{a-1}} \sum_{d=1}^k \sum_{j=r_j}^{r_{d+1}-1} \prod_{i=1}^d \frac{r_i - i}{j-1} \sum_{s=1}^{d \wedge a} \binom{j-1}{s-1} \binom{n-j}{a-s}$$
(5)

where $d \wedge a = \min(d, a)$. Table 10 presents the values r_1^* , r_2^* , r_3^* , r_4^* , and r_5^* and the Pr(success) for n = 20, 50, and 100 and k = 5 in all cases.

Problem PMP-DC: the standard secretary problem with two choices.

As noted by Gilbert and Mosteller (1966), there are three mutually exclusive ways of winning in this problem with the two-threshold strategy (r_1, r_2) :

(i) win with the first choice (never use the second choice),

Table 8 Optimal values of r_1^* and r_2^* for problem PMP-2 and the associated probability of choosing either the best or second-best applicant for selected values of n

n	10	20	40	80	1000	∞
r_1^*	4	8	15	35	348	0.347 <i>n</i>
r_2^*	7	14	27	67	667	2n/3
Pr(win)	0.637	0.605	0.589	0.580	0.574	0.574

- (ii) win with the second choice where no choice is used before r_2 ,
- (iii) win with the second choice, where the first choice is used in one of the rounds r_1 , $r_1 + 1$, $r_1 + 2$, ... $r_2 1$.

The respective probabilities of success are

$$Pr(a) = \frac{r_1 - 1}{n} \left(\frac{1}{r_1 - 1} + \frac{1}{r_1} + \frac{1}{r_1 + 1} + \dots + \frac{1}{n - 1} \right), \quad \text{if } r_1 > 1$$

$$Pr(b) = \frac{r_1 - 1}{n} \sum_{v=r_2+1}^{n} \sum_{u=r_2}^{v-1} \frac{1}{(u - 1)(v - 1)}, \quad \text{if } r_2 > r_1 > 1 \quad (6)$$

$$Pr(c) = \frac{r_2 - r_1}{n} \sum_{v=r_2}^{n} \frac{1}{v - 1}, \quad \text{if } r_2 > r_1 > 1.$$

Therefore, the probability of success (win) is computed from Pr(success with (r_1, r_2)) = Pr(a) + Pr(b) + Pr(c). The values of r_1^*, r_2^* , and Pr(success with (r_1^*, r_2^*)) are determined numerically—see Table 11 for these values and a comparison with the PS heuristic where n = 20, 50, and 100.

Problem ERMP

The optimal decision rule for retaining a single item that minimizes the expected value of the absolute rank of the item selected can be computed by dynamic programming (Chow et al., 1964). This rule consists of a sequence of *n* integers (thresholds) $r_n(j)$, j = 0, 1, 2, ..., n, that satisfy the inequalities

$$1 < r_n(0) < r_n(1) \dots < r_n(n) = n.$$

Denote the index of the item by i (i = 1, 2, ..., n). Then, if $i \le r_n(0)$, the item is never accepted. If $i \le r_n(1)$, then the item with relative rank $R_i = 1$ is accepted. In general, item i is accepted if its relative rank satisfies $R_i \le j$ and $i \le r_n(j)$, provided that no item was retained earlier. If $n \to \infty$, then the optimal decision rule has the following form (Ferguson, 2002). Do not choose any applicant until 25.8% of all the applicants have been interviewed and discarded; then select any applicant i with R_i = 1. After interviewing 44.8% of the applicants, select any applicant i with relative rank of either 1 or 2. After 56.4% of the applicants have been interviewed, select any applicant with relative rank of either 1, 2, or 3, and so on.

Let V(n) denote the expected value of the absolute rank of the applicant selected by the optimal rule described above when there are *n* applicants. Chow et al. (1964) shows that if $n \to \infty$, then

$$V(n) = V(\infty) = \prod_{j=1}^{\infty} \left(1 + \frac{2}{j}\right)^{1/(j+1)} = 3.869.$$

Krieger and Samuel-Cahn write: "This is an astonishing result since it shows, for example, that from one million items we can sequentially select an item with expected rank less than 4" (2009, p. 1042). We concur.

Appendix 2 Further results comparing the optimal solutions to the PS heuristic

PS in problem PMP-2 (Table 9)

Model	п	С	r_{l}	r_2	Pr(candidate selection)	Pr(win)	Approximation score
Optimal	20	N/A	8	14	0.773	0.6046	1.000
Heuristic	20	1.65-1.80	8	12	0.815	0.5913	0.978
Optimal	50	N/A	18	34	0.780	0.5858	1.0
Heuristic	50	1.45-1.50	21	30	0.772	0.5749	0.981
Optimal	100	N/A	35	67	0.776	0.5796	1.000
Heuristic	100	1.80	39	57	0.764	0.5724	0.988

Table 9 Optimal and PS heuristic values for problem PMP-2 for n = 20, 50, and 100

Values of the a-score > 0.95 appear in bold

PS in problem PMP-5 (Table 10)

Model	п	с	<i>r</i> ₁	<i>r</i> ₂	<i>r</i> ₃	<i>r</i> ₄	<i>r</i> ₅	Pr(candidate selection)	Pr(win)	Approximation score
Optimal	20	N/A	7	11	14	16	18	0.9502	0.9051	1.000
Heuristic	20	1.90	8	11	13	15	16	0.9630	0.9040	0.999
Optimal	50	N/A	17	26	33	39	44	0.9382	0.8781	1.000
Heuristic	50	1.70	19	28	33	36	39	0.9344	0.8654	0.986
Optimal	100	N/A	33	52	66	77	87	0.9332	0.8692	1.000
Heuristic	100	1.70	38	55	65	71	76	0.9379	0.8559	0.985

Table 10	Optimal	and PS	heuristic	for	PMP-5	selecting	for 1	i = 20,	50,	and 1	00
----------	---------	--------	-----------	-----	-------	-----------	-------	---------	-----	-------	----

Values of the a-score > 0.95 appear in bold

PS in problem PMP-DC (Table 11)

PS in problem ERMP

Table 12 presents threshold values for n = 100, 1000, and 10,000 (column 1). The top three rows present the proportional threshold values, namely $r_i(n)/n$, to allow

Model	п	с	r_1	r_2	Pr(candidate selection)	Pr(win)	Approximation score
Optimal	20	N/A	5	8	0.802	0.6178	1.0
Heuristic	20	3.30 to 4.00	5	8	0.802	0.6178	1.0
Optimal	50	N/A	12	19	0.780	0.6014	1.0
Heuristic	50	3.40 to 3.60	12	19	0.780	0.6014	1.0
Optimal	100	N/A	23	38	0.780	0.5962	1.0
Heuristic	100	3.40 to 3.45	23	38	0.780	0.5962	1.0

Table 11 Optimal and PS heuristic values for problem PMP-DC with n = 20, 50, and 100

Values of the a-score > 0.95 appear in bold

Table 12 Proportional threshold values by the PS heuristic rule for problem ERMP

n	с	$r_0(n)/n$	$r_1(n)/n$	$r_2(n)/n$	$r_3(n)/n$	$r_4(n)/n$	$r_5(n)/n$	$r_{6}(n)/n$
100	2.30	0.31	0.46	0.57	0.64	0.69	0.73	0.76
1,000	2.36	0.297	0.459	0.560	0.629	0.680	0.718	0.748
10,000	2.41	0.2932	0.4535	0.5545	0.6240	0.6748	0.7135	0.7438
$n \rightarrow \infty$	2.42	0.2584	0.4476	0.5640				

comparison across different values of *n*. The bottom row presents the proportional threshold values for the optimal rule computed by dynamic programming for j = 1, 2, and 3 (reported by Ferguson, 2002). Column 2 presents the corresponding values of the parameter *c*. It shows that the value of *c* is quite stable, and that it converges to c = 2.42. Table 12 also shows that as the value of *n* increases, the proportions $r_j(n)/n$ decrease gradually.

Table 13 (Table 5 in KS-C) displays the expected values of the absolute rank of the item selected under the optimal rule (column 2) and under the heuristic rule (column 3) for selected values of n. Column 4 exhibits the corresponding values of

		DG 1 1		Approximation score	
n	Optimal expected rank	PS expected rank	Optimal value of c		
20	3.002	3.017	2.1	0.995	
50	3.412	3.503	2.2	0.974	
100	3.6032	3.7348	2.30	0.963	
1000	3.8324	3.9062	2.36	0.981	
10,000	3.8649	3.9258	2.41	0.984	
100,000	3.8690	3.9279	2.41	0.985	
1,000,000	3.8695	3.9281	2.42	0.985	

Table 13 Optimal and PS heuristic for selecting a single item in problem ERMP

Values of the *a*-score > 0.95 appear in bold

the parameter c. The right-hand column displays the approximation scores of the heuristic rule. Note that for the values of n reported in Table 13, one may sequentially select an item with expected rank less than 4 and that the approximation score increases in n from 96.3 percent for n = 100 to 98.5% for n > 100,000.

Appendix 3: The optimal subgame perfect equilibrium solution to problem PMP-COMP

Table 1 (second column) in K&L presents the optimal thresholds for n = 10. Given the reported thresholds, we estimate the probability of employers 4, 5, 6, 7, and 8 winning (hiring the best candidate), as 0.075, 0.017, 0.002, 1.7×10^{-4} and 1×10^{-5} , respectfully. The thresholds for employers 9 and 10 are reported as zero; recall that the highest-ranked employer with a threshold of zero (in this case 9) is guaranteed of winning with a probability of 1/n = 0.1. As we have estimated the probability of winning for employers 4 through 8 to be less than 0.1, each of these employers independently has an incentive to deviate from playing the reported threshold of one to playing zero. This is because deviating would make them the most highly ranked employer with a zero threshold, increasing their probability of winning to 0.1. However, this incentive to deviate is inconsistent with the notion of a Nash equilibrium, i.e., that no employer has an incentive to unilaterally deviate from his action. Consequently, the threshold values for n = 10 reported in Table 1 of K&L cannot be a subgame-perfect Nash equilibrium of the game. It can be shown that for n = 50, a similar incentive for unilateral deviation also exists for employers ranked 6 through to 10 given the thresholds reported in K&L (2015, Table 1).

Our solution below follows the reasoning in K&L with two changes. The first is the use of 1-based indexing, rather than 0-based indexing—this is merely a matter of convenience for programming languages based on the former rather than the latter. The second is defining the optimal risk, $R_k(i)$ as that belonging to a set of rules that make offers to candidates starting with the *i*th applicant onwards (instead of ignoring the first *i* applicants)—this definition is more convenient when combined with 1-based indexing.

Pseudo-code for calculating the optimal thresholds in the PMP-COMP problem

T(1) = n;for j = 1:k $R_j(n + 1) = 1;$ for i = n: -1:1if $i \le T(j)$ $R_j(i) = \frac{1}{i} \cdot \min\left\{R_j(i + 1), 1 - \frac{i}{n}\right\} + \left[1 - \frac{1}{i}\right] \cdot R_j(i + 1)$ else $R_j(i) = \frac{1}{i} \cdot R_{j-1}(i + 1) + \left[1 - \frac{1}{i}\right] \cdot R_j(i + 1)$ end end $T(j + 1) = \inf\left\{i: R_j(i) \ge \left[1 - \frac{1}{i}\right]\right\} - 1$ end $t_k = T(k + 1)$

Funding Open Access funding enabled and organized by Projekt DEAL.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11238-022-09881-0.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicate otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

- Artinger, F., Petersen, M., Gigerenzer, G., & Weibler, J. (2015). Heuristics as adaptive decision strategies in management. *Journal of Organizational Behavior*, 36(S1), S33-52. https://doi.org/10.1002/job. 1950
- Åstebro, T., & Elhedhli, S. (2006). The effectiveness of simple decision heuristics: Forecasting commercial success for early-stage ventures. *Management Science*, 52(3), 395–409. https://doi.org/ 10.1287/mnsc.1050.0468
- Bearden, J. N., Murphy, R. O., & Rapoport, A. (2005). A multi-attribute extension of the secretary problem: theory and experiments. *Journal of Mathematical Psychology*, 49(5), 41–422. https://doi. org/10.1016/j/jmp.2005.08.002
- Bearden, J. N., & Rapoport, A. (2005). Operations research in experimental psychology. In C. Smith (Ed.), *Tutorial in operations research: Emerging theory, methods, and applications.* INFORM.

- Bearden, J. N., Rapoport, A., & Murphy, R. O. (2006). Sequential observation and selection with rankdependent payoffs: An experimental study. *Management Science*, 52(9), 1437–1449.
- Bingham, C. B., & Eisenhardt, K. M. (2011). Rational heuristics: The 'simple rules' that strategists learn from process experience. *Strategic Management Journal*, 32(13), 1437–1464. https://doi.org/10. 1002/smj.965
- Bruss, F. T., & Samuels, S. M. (1987). A unified approach to a class of optimal selection problems with an unknown number of options. *Annals of Probability*, 15(2), 824–830.
- Chow, Y. S., Moriguti, S., Robins, H., & Samuels, S. M. (1964). Optimum selection based on relative rank (the "Secretary Problem"). *Israel Journal of Mathematics*, 2, 81–90.
- Chun, Y. H. (1998). Selecting the best choice in the full information group interview problem. *European Journal of Operations Research*, 119, 635–651.
- Chun, Y. H. (2000). Sequential search and selection problem under uncertainty. *Decision Sciences*, 31, 627–648.
- Corbin, R. M. (1980). The secretary problem as a model of choice. *Journal of Mathematical Psychology*, 21, 1–29.
- Dietz, C., van der Laan, D., & Ridder, A. (2011). Approximate results for a generalized secretary problem. Probability in the Engineering and Informational Sciences, 25(2), 157–169. https://doi. org/10.1017/s026996481000032x
- Dufwenberg, M., Sundaram, R., & Butler, D. J. (2010). Epiphany in the game of 21. Journal of Economic Behavior and Organization, 75(2), 132–143. https://doi.org/10.1016/j.jebo.2010.03.025
- Ferguson, T. S. (1989). Who solved the secretary problem? Statistical Science, 4, 282-289.
- Ferguson, T. S. (2002). *Optimal stopping and applications*. Dept. of Mathematics, UCLA. http://math.ucla.edu/~tom/stopping/contents.html
- Freeman, P. R. (1983). The secretary problem and its extensions: A review. International Statistical Review, 51, 189–206.
- Gardner, M. (1960). Mathematical games. Scientific American, 202(2), 150-156.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The foundations of adaptive behavior*. Oxford University Press.
- Gigerenzer, G., Todd, P. M., the ABC Research group. (1999). Simple heuristics that make us smart. Oxford University Press.
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the America Statistical Association*, 61, 35–73.
- Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, 6(3), 1375–1394.
- Hertwig, R., Hoffrage, U., the ABC Group. (2013). *Simple heuristics in a social world*. Oxford University Press.
- Hertwig, R., Woike, J. K., Pachur, T., Brandstätter, E., & the Center for Adaptive Rationality (eds.). (2019). The robust beauty of heuristics in choice under uncertainty. In *Taming uncertainty* (pp. 29–50). MIT Press. https://doi.org/10.7551/mitpress/11114.001.0001
- Immorlica, N., Kleinberg, R., & Mahdian, M. (2006). Secretary problems with competing employers. WINE'06: Proceedings of the Second international conference on Internet and Network Economics, 389–400.
- Kahneman, D. (2003a). Maps of bounded rationality: psychology for behavioral economics. *American Economic Review*, 93(5), 1449–1475.
- Kahneman, D. (2003b). A perspective on judgment and choice: mapping bounded rationality. American Psychologist, 58(9), 697–720.
- Kahneman, D. (2011). Thinking. Penguin.
- Karlin, A. & Lei, E. (2015). On a competitive secretary problem. Proceedings of the 29th conference on artificial intelligence (pp. 944–950).
- Krieger, A. M., & Samuel-Cahn, E. (2009). The secretary problem of minimizing the expected rank: a simple suboptimal approach with generalizations. Advances in Applied Probability, 41, 1041–1058.
- Lee, M. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26.
- Levitin, A. (2012). Introduction to the design and analysis of algorithms (3rd ed.). Pearson/Addison-Wesley.
- Lindley, D. V. (1961). Dynamic programming and decision theory. Applied Statistics, 10, 39-51.

- Mak, V., Seale, D. A., Rapoport, A., & Gisches, E. J. (2019). Voting rules in sequential search by committees: theory and experiments. *Management Science*, 65(9), 4349–4364. https://doi.org/10. 1287/mnsc.2018.3146
- Mandler, M., Manzini, P., & Mariotti, M. (2012). A million answers to twenty questions: choosing by checklist. *Journal of Economic Theory*, 147(1), 71–92.
- Manzini, P., & Mariotti, M. (2007). Sequentially rationalizable choice. American Economic Review, 97(5), 1824–1839.
- Manzini, P., & Mariotti, M. (2012a). Categorize then choose: boundedly rational choice and welfare. Journal of the European Economic Association, 10(5), 1141–1165.
- Manzini, P., & Mariotti, M. (2012b). Choice by Lexicographic semi-orders. *Theoretical Economics*, 7(1), 1–23.
- Manzini, P., & Mariotti, M. (2014). Stochastic choice and consideration sets. *Econometrica*, 82(3), 1153–1176.
- March, J. G., & Simon, H. A. (1993). Organizations revisited. Industrial and Corporate Change, 2(1), 299–316. https://doi.org/10.1093/icc/2.1.299
- Matsui, T., & Ano, K. (2016). Lower bounds for Bruss' odds problem with multiple stoppings. Mathematics of Operations Research, 41(2), 700–714.
- McKinney, C. N., Jr., & Huyck, J. B. V. (2013). Eureka learning: heuristics and response time in perfect information games. *Games and Economic Behavior*, 79, 223–232. https://doi.org/10.1016/j.geb. 2013.02.003
- Palley, A. B., & Kremer, M. (2014). Sequential search and learning from rank feedback: Theory and experimental evidence. *Management Science*, 60(10), 2525–2542. https://doi.org/10.1287/mnsc. 2014.1902
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(3), 534–552.
- Presman, E. L., & Sonin, I. M. (1972). The best choice problem for a number of objects. *Theory of Probability and Its Applications*, 17, 657–668.
- Rasmussen, W. T., & Robbins, H. (1975). The candidate problem with an unknown population size. *Journal of Applied Probability*, 12, 692–701.
- Samuels, S. M. (1991). Secretary problems. In B. K. Ghosh & P. K. Se (Eds.), Handbook of sequential analysis (pp. 381–405). Marcel Dekker.
- Schotter, A., & Trevino, I. (2020). Is response time predictive of choice? An experimental study of threshold strategies. *Experimental Economics*. https://doi.org/10.1007/s10683-020-09651-1
- Seale, D. A. (1996). Sequential observation and selection with relative ranks: An empirical investigation of the secretary problem. University of Arizona: Unpublished Ph.D. dissertation.
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making and relative ranks: an experimental investigation of the "Secretary Problem." Organizational Behavior and Human Decision Processes, 69(3), 221–236.
- Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: the secretary problem with unknown population size. *Journal of Behavioral Decision Making*, 13, 391–411.
- Simon, H. A. (1957). Models of man: Social and national. Wiley.
- Simon, H. A. (1959). Theories of decision-making in economics and behavioral science. American Economic Review, 49, 253–283.
- Simon, H. A. (1982). Models of bounded rationality. MIT Press.
- Simon, H. A. (1990). Invariants of human behavior. Annual Review of Psychology, 41(1), 1-20.
- Spiliopoulos, L., & Hertwig, R. (2020). A map of ecologically rational heuristics for uncertain strategic worlds. *Psychological Review*, 127(2), 245–280. https://doi.org/10.1037/rev0000171
- Stein, W. E., Seale, D. A., & Rapoport, A. (2003). Analysis of heuristic solutions to the best choice problem. *European Journal of Operations Research*, 151, 140–152.
- Sterelny, K. (2003). Thought in a hostile world. Wiley-Blackwell.
- Thorngate, W. (1980). Efficient decision heuristics. Behavioral Science, 25(3), 219–225. https://doi.org/ 10.1002/bs.3830250306
- Todd, P. M., Gigerenzer, G., the ABC Group. (2012). *Ecological rationality: Intelligence in the world*. Oxford: Oxford University Press.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131.
- Woryna, A. (2017). The solution of a generalized secretary problem via analytic expressions. *Journal of Combinatorial Optimization*, 33(4), 1469–1491. https://doi.org/10.1007/s10878-016-0050-8

Yeo, A. J., & Yeo, G. F. (1994). Selecting satisfactory secretaries. Australian Journal of Statistics, 36(2), 185–198.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.