# Sequential search behavior changes according to distribution shape despite having a rank-based goal

**Tsunhin John Wong (jwong@mpib-berlin.mpg.de) [1]**

**Jonathan D. Nelson (jonathan.d.nelson@gmail.com) [1,2]** **Lael J. Schooler (lschoole@syr.edu) [1,3]**

[1] Max Planck Institute for Human Development, Lentzeallee 94, Berlin 12167, Germany

[2] School of Psychology, University of Surrey, Guildford, GU2 &XH, The United Kingdom

[3] Department of Psychology, 509 Huntington Hall, Syracuse University, NY 13244, USA

## Abstract

In the area of sequential choice, the 'Secretary Problem' has been a prominent paradigm within the study of optimal stopping for sequential search tasks. Most recent studies of the Secretary Problem present decision makers with the relative ranks of options. A recurring finding is that decision makers tend to end their search earlier than optimal decision strategies (e.g. Helversen, Wilke, Johnson, & Schmid, 2011; Seale & Rapoport, 1997, 2000). By revealing only relative ranks of options or items, issues of learning and incomplete knowledge are avoided; however, this leaves open the question of how sensible human decision makers are when they know more about the distribution of items. Rather than presenting merely ranks to decision makers, we presented numerical values drawn from three distinct distributions in which relatively high value items were scarce, evenly distributed, or abundant. We found that they selected their items earlier than they would if they utilized the optimal selection rule. More importantly, in contrast to the conclusion of Kahan, et al. (1967), we found the selection points of decision makers were sensitive to the underlying distribution. In contrast, the optimal strategy is totally based on quantile ranks regardless of the type of distributions.

**Keywords:** Sequential choice, Secretary Problem, Heuristics

## Introduction

In everyday life, there are many situations in which we need to choose from options presented sequentially. The decision makers may need to choose the best option out of a randomized sequence and may not have the chance to choose an option they have previously rejected. One version of the problem with the goal to find the largest option in the sequence appeared in the mathematical games column by Gardner (1960a, 1960b) in Scientific American. This problem is also known as the *Secretary Problem*.

In the Secretary Problem, there is a reward only if the best item (an interchangeable term for 'option' in our paper) in the sample is chosen. This scenario does not occur too often in daily life, as every option usually has its own value. However, these scenarios do exist, for example, if you are going to strategically sponsor a presidential candidate during their elections for the future benefit of your company, and you probably have to choose only one out of many. At the end, there will only be a single president. In another example, when companies compete to become the contractor of projects, at the end, in most cases, there is only one contractor per project; as an investor or collaborator, you want to choose the one and only winner. Basically, this scenario holds true for winner-takes-all games.

### Previous studies with no-information problems

Since the 1960s, many mathematical and behavioral studies have investigated various aspects of the Secretary Problem and its variants that similarly share the goal of choosing a desirable option based on a single attribute of quality. The mathematical studies usually aimed at finding out the optimal choice strategies in the targeted sequential choice problems. Many mathematical analyses assume that the options are drawn from a distribution fully known to the decision maker, also known as *full-information* problems. The behavioral studies typically compare human behavior to an optimal strategy and attempt to explain whether and why human decision makers are optimal or not. Some behavioral studies are based on *partial-information* problems, in which the decision maker knows some (perhaps distribution family and some parameters), but not all, aspects of the distribution from which the options are drawn. The relative rank-based problems, also called *no-information* problems, are where only the relative ranks of options are presented; the ranks of previous options are updated as new options appear.

In recent years, studies of the Secretary Problem have mainly considered no-information problems (e.g. Helversen et al., 2011; Lee, 2006; Lee, O'Connor, & Welsh, 2004; Rapoport & Tversky,1970; Seale & Rapoport, 1997, 2000). No-information problems present only relative ranks of items and make the Secretary Problem more tractable because complexities such as how decision makers learn the underlying distribution and their individual differences in learning and understanding are altogether avoided. However, in daily life, much of the time decision makers judge an option with some degree of knowledge or prior belief about the distribution it is coming from. Therefore, partial-information problems are closer to most of the sequential choice scenarios in daily life. In this study, we tried to approach this classic sequential problem by presenting values to the decision makers instead of relative ranks. By using real values, it becomes possible to manipulate the distribution environment and investigate the choice behavior of decision makers in different underlying distribution shapes.

### How distribution shapes affect strategies on the full-information Secretary Problem

Some studies have investigated the effect of distribution shape in sequential problems. Rapoport and Tversky (1966, 1970) trained their subjects for a few weeks on distributions of item quality and concluded that their subjects performed optimally on two-thirds of the tasks. However, Rapoport and Tversky only used a uniform distribution for item quality. Distribution shapes, for instance positive skew, negative skew, or uniform, have been manipulated during experimental investigations of the Secretary Problem by only a few studies (e.g. Guan & Lee, 2014; Kahan, Rapoport, & Jones, 1967). One early study was conducted by Kahan et al. (1967); they trained their participants for over 3 weeks in the experiments. They concluded that there was not sufficient evidence to say their participants used different stopping points in environments with the different underlying distribution shapes. Guan & Lee (2014) tested their participants in a slightly different setting: without the benefit of the extensive training that Kahan et al's participants underwent, Guan & Lee's participants worked on randomized sequences of five items drawn from one of two distributions derived from the Beta distribution. They concluded that their participants could have used multiple-thresholds with decreasing values towards the end of the sequence, and that these thresholds are not affected by the value of preceding items.

Gilbert and Mosteller (1966, Section 3) explain how to derive the optimal strategy for the Secretary Problem under different distributions, when the goal is to find the highest item in the sequence and the distribution shape is fully known. The optimal strategy is in the form of a multi-threshold rule, a sequence of nonincreasing thresholds, usually monotonically decreasing. The largest and first threshold is for deciding whether to accept the first item; provided that the first item was not accepted, the second threshold is used to decide whether to accept the second item; and so on. The rules of the game require that the last item must be accepted if no previous item has been accepted. In the case of full knowledge problem of the Secretary Problem, the optimal strategy is based on a distribution-related percentile-based multithreshold rule (See Gilbert & Mosteller, Section 3). The percentile-based thresholds vary with the number of items in the sequence, however, the strategy is basically the same across distributions. What will be different is the exact numerical threshold values for different distributions. For a number of continuous and discrete distribution families, we derived the corresponding multi-threshold optimal decision strategies, which numerically vary according to the underlying distribution shapes.

Consider the behavior of the optimal strategy given three Beta distributions: a positively-skewed distribution ($\beta(1,3.7)$; skew = 1.00), a negatively-skewed distribution ($\beta(3.7,1)$; skew = -1.00), and the uniform distribution ($\beta(1,1)$; skew = 0). If we consider the underlying distribution about the quality of items to be an 'environment', we can characterize a positively-skewed distribution as a scarce environment with a lot of low quality items and only a few high quality items; and correspondingly a negatively-skewed distribution can be characterized as an abundant environment with mostly high quality items. These distributions are normalized (Z-scored) so that irrespective of the underlying distribution, a random response strategy has expected payoff of zero, and standard deviation of 1. The optimal thresholds are highest for the positively-skewed distribution, and then for the uniform distribution, while the negatively-skewed distribution has the lowest thresholds.. We wanted to see (1) how well the optimal models account for the decision makers' performance and (2) to what extent the decision makers are sensitive to the different underlying distributions.

When the optimal strategy for the full-information Secretary Problem is used (Gilbert & Mosteller, 1966), although the numeric values of the multithreshold rule vary according to the underlying distribution, the probability of stopping a search follows a fixed set of probabilities, and does not vary with the underlying distribution. However, it is unknown whether search behavior will be affected by the underlying distributions in the Secretary Problem, even when decision makers are familiar with the distributions. If people are insensitive to the underlying distribution shapes, their search lengths (how long they reach the randomized sequence before selecting an item) and success rates will remain the same in different distributions, as predicted by the optimal models.

We used a between-subjects manipulation of distribution shape in an experiment with real-money payoffs, and a scarce distribution, a neutral distribution and an abundant distribution, to test these predictions with human subjects.

## Method

We used an Internet card game with real money payoffs to implement the Secretary Problem. In a between subjects design, each subject was randomly assigned to the positively-skewed, negatively-skewed, or uniform distribution. There was a training phase of at least 10 rounds of the card game, and a test phase with exactly 10 rounds of the same card game. The value of each card that was shown to the subjects was formed by taking the underlying Z-scored distribution value, adding 4, and then multiplying it by 1200, such that all cards would have positive point values, between approximately 2500 and 9500.

### Stimulus

In each round, subjects were asked to obtain the card with the largest number out of a 25-card sequence. Among all 10 test-phase rounds of the experiment, a successful round required selection of the largest card, and this added a fixed bonus to the participation payment. The conversion to real money was 2000 points to $1 US; this was disclosed at the outset. Subjects were encouraged to achieve as high of a score as they could.

### Subjects

We recruited subjects through the Amazon M-Turk platform, and allowed only US subjects with consistently good reviews under the Amazon M-Turk monitoring system ('M-Turk Master workers') to participate. Bonus payments proportional to subjects' performance were rewarded through the M-Turk system, on top of a one US dollar participation fee. Informed consent was obtained from all subjects.

**Training** After the instructions and a practice example, there was a training phase. A subject had to correctly select the largest card for at least 4 rounds in their 10 rounds of training; otherwise, the phase would begin all over again with the count of their successful rounds reset. No bonus payment was offered for the training phase. If the training phase was successfully completed, then the subject had to complete a test phase consisting of 10 rounds with a real-money (U.S. dollars) bonus payment of 5 cents for each time successfully selecting the largest card in the sequence of the 10 test rounds.

## Result

208 recruited subjects were randomly assigned among the conditions: 68 subjects to the positive skew distribution condition, 70 to the uniform condition, and 70 to the negative skew condition. 155 subjects completed the experiment with 54 in the positive skew condition, 52 in the uniform condition, and 49 in the negative skew condition.

### Conscientiousness of subjects

In the full-information Secretary Problem, the optimal strategy sometimes (about 3.4% of the time) selects the very first card. However, selecting a card that is smaller than any previously viewed card guarantees failure. For each distribution and phase (Figure 1), we checked the rate of these behaviors in our subjects. These behaviors were rare, consistent with conscientious behavior of the subjects.

### Average success rate

Among the number of rounds of Secretary Problem attempted, the percentage of rounds in which the best item in the sequence was selected – success rates (Figure 2, thick horizontal lines) of these three distributions are similar. The success rate of random responding (4% in all conditions, denoted by the dashed grey lines) lies well beyond the interquartile range (IQR) for all conditions and beyond the whiskers (first quantile – 1.5 x IQR, and third quantile + 1.5 x IQR) for most of the rounds. When using the Wilcoxon one-sample tests to the success rates of training rounds and

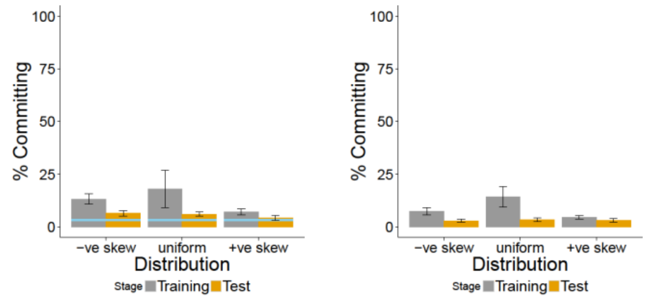


1a) % choosing 1st card 1b) % choosing non-max

Figure 1a & b: Left – Percentage of rounds choosing the first card. Error bars cover the range one standard error above and below the mean. The full-knowledge optimal strategies also select the first item out of 25 with a 3.24% chance for all the distributions (see the blue line); Right – Percentage of rounds not choosing the largest card. Error bars cover the range one standard error above and below the mean. The graph indicates how often subjects selected cards that were not the largest card seen in the sequence so far.

test rounds to the 4% random success rate, p-values are all < 0.001 for the 3 distributions in the training rounds, and for the test rounds; therefore the possibility that subjects were responding randomly can be ruled out. Subjects from the positively-skewed distribution condition had the highest success rates, closely followed by the subjects from the uniform distribution, and then from the negatively skewed distribution (Figure 2).

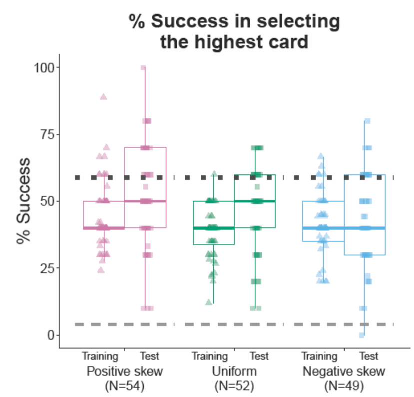


Figure 2: Percentage of success in selecting the highest card. The lower and upper hinges of the boxplots correspond to the 25th and 75th percentile of the data. The whiskers extend to 1.5 times beyond the bounds of the interquartile range. Dotted horizontal dotted lines denote the average score of the optimal strategy for each distribution. Dashed grey lines denote the success rate of random responding (4%, in all conditions). Scores have a slight ordering from positive-skew to uniform and then to negative-skew distribution.

### Values of chosen cards

Values of cards our subjects selected at every position under each distribution were plotted, and compared to the theoretical threshold values that the optimal strategies predict: conditional that the card values are higher than the threshold for a given position and that no card previously shown is higher than the selected card (Figure 3). The ordering of card values obtained was highest for the positively-skewed distribution, second highest for the uniform distribution, and lowest for the negatively-skewed distribution, resembling the ordering in the optimal strategies. Subjects initially would select cards with lower values that showed up in the first few positions during their training (Figure 3, top); however, they might have already adjusted and increased their thresholds when they reached the test rounds. In the test rounds, the accepted values are closer to the optimal strategies in the earlier positions, and the accepted values, unlike the pattern exhibited by the optimal strategies, have a stretched plateau shape, generally extending from 1$^{st}$ position until to almost the 23$^{rd}$ or 24$^{th}$ position out of 25. Data seem to get closer to the patterns of chosen card values of the optimal multithreshold rules from training rounds to test rounds, nevertheless, their behavior is not optimal. Given the similarity between the patterns of the values of chosen cards and the conditional expected values of chosen cards using optimal strategies with non-increasing multithreshold rules, the ordering of underlying thresholds utilized by our subjects among the conditions is probably consistent with the ordering of the non-increasing multithreshold rules.

### Efficiency

The optimal strategies suggest that in the positively-skewed distribution, the thresholds used should be the highest, followed by the uniform distribution, and then by the negatively-skewed distribution. This trend was demonstrated in the percentile ranks of cards selected by our participants. The optimal strategies suggested that in order to select the best item with the optimal thresholds, search length would be the same for all three distributions regardless of their characteristics and skewness; this trend in search length was not exhibited by our subjects (Figure 4). One-sample Wilcoxon tests showed that the search lengths of all distributions from both training rounds and test rounds are shorter than that of the optimal strategies (all p-values $< 0.001$). There were also differences in search length as a function of distribution shape (Kruskal-Wallis rank sum test, $p = 0.004$); it appears that they searched longer in the positively-skewed distribution than in the other two distributions. In addition to the above observations, the subjects also selected cards with higher values in the test rounds than in the training rounds (Wilcoxon Signed-Rank Sum Tests, $W=268$, $p <0.001$, $d= -0.58$ for the negatively skewed distribution; $W=418$, $p = 0.017$, $d= -0.25$ for the uniform distribution; and $W=373$, $p <0.001$, $d= -0.44$ for positively skewed distribution), despite having similar search length in training and test rounds. It seems that our subjects have learnt about the distributions and had their performance improved towards the performance of the optimal strategies.

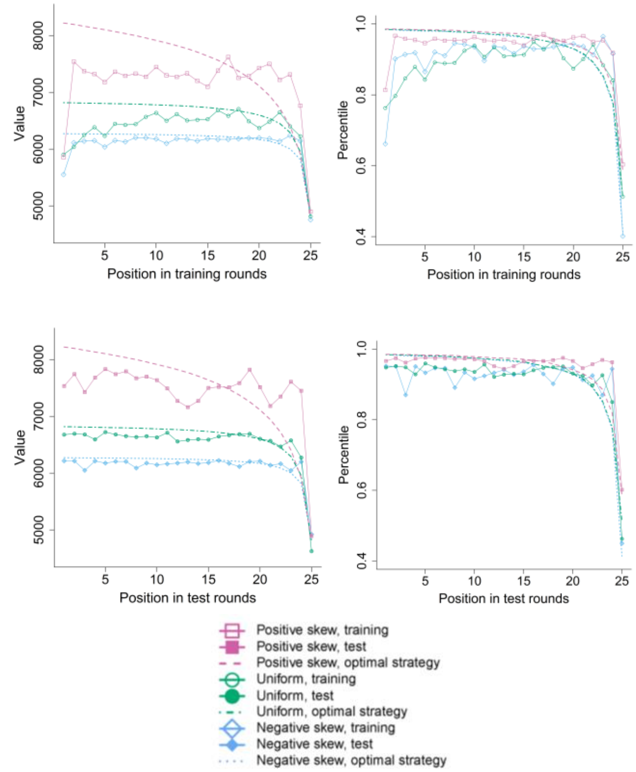


Figure 3: Values and percentiles of chosen cards at positions where they are chosen. We grouped the results by their decision points in the sequence and found out the mean values for each position on condition that the subjects stopped there. The dashed lines are the values of card at a certain position when the decision maker makes use of the optimal (full-knowledge) strategy from the very beginning of the sequential choice problem, on condition that the stop happens to be at that position

### Search length

Besides looking at the means of search length, we also looked at the cumulative distribution of how frequently a selection had already been made by a particular position in the sequence (Figure 5). Optimal strategies for the full-information Secretary Problem lead to the same pattern of search length (Figure 5, solid black line without marker) irrespective of the underlying distribution (Gilbert & Mosteller, 1966). For our subjects, when combined as a group, search length was closer to the optimal strategies in the positive skew condition, and not as close in the negatively skewed distribution or in the uniform distribution. The search lengths of the subjects follow a pattern of concave downward until the 24$^{th}$ item in the sequence; shapes of the cumulative distribution curves suggest a tendency for the subjects to select an item sooner than the optimal strategies: the curves of all conditions are concave downward whereas the curve of the optimal strategy is slightly concave upward – in other words, they all have the tendency to stop too soon in both training rounds and test rounds. Although the outcome of optimal strategies does not predict that, there are some obvious differences between the curves of the positively skewed distribution and that from the other two distributions in both training rounds and test rounds. For the comparison of these curves, we conducted the Kolmogorov-Smirnov tests to examine their differences: in particular, the difference of search lengths among the three underlying distributions; however, none of the Kolmogorov-Smirnov tests was significant, including the comparison between the curves from subjects' data to the curves from the optimal strategies. We suspect the Kolmogorov-Smirnov test for distributions may not be for right test for our purpose. The sudden increases from item 24 to item 25 reflect the rule of the task that participants have to accept the last item as long as it is reached. For the cumulative selection probability, learning was reflected only a bit more obviously in the positive skew distribution as a shift closer to the optimal strategies from training to test (Figure 5, top to bottom).

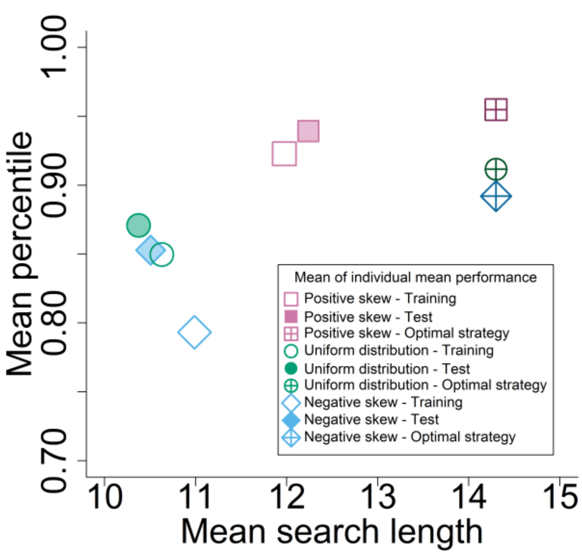


Figure 4: Values of chosen cards vs search length. Means of individual scores are plotted as a function of means of individual search length (number of items searched up until an item is accepted). The aggregate means for all subjects in each distribution are plotted with different markers (unfilled for training, filled for test). The optimal strategy is plotted with markers with internal plus signs. The upward switch of locations of these markers from when subjects were during training rounds to when they were in the test rounds shows how performance improved from training to test for each distribution, despite similar search length.

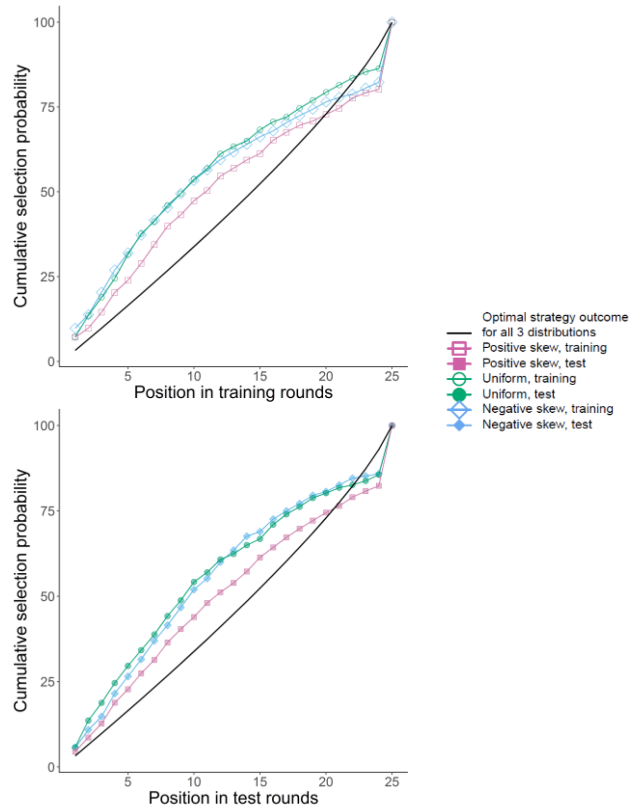


Figure 5: The patterns of search length are similar between training rounds and test rounds, having similar concave downward patterns. All patterns showed they stopped earlier than the optimal strategy does. Subjects in the positively skewed distribution seemed to stop early with a lesser tendency. The jumps towards the last position were due to the rule of the Secretary Problem that the last item when reached must be selected.

## Discussion

There have been experiments on the Secretary Problem (Gardner, 1960a), in which a payoff is obtained only if the largest item is selected (e.g. Bearden, Rapoport, & Murphy, 2006; Lee et al., 2004; Seale & Rapoport, 1997, 2000), as well as related tasks in which the goals are to obtain an item in the top 10% or 25% (Todd & Miller, 1999). Although many real-life situations involve real value options, Secretary Problems with real values have received little attention in psychological experiments. We conducted a study of the Secretary Problem with information, to look at how human sequential choice behavior may vary as a function of distribution shape. Our subjects were sensitive to the underlying distribution shape (as shown in Figure 4 and Figure 5).

The age-old claim continues: our subjects stopped their searches early, i.e. earlier than the optimal strategies do (Figure 4 and Figure 5), and it was also not likely for them to keep searching until the last few items of the sequence. There are a few plausible explanations. The tendency to stop earlier than predicted could be obtained by an enhanced model that incorporates a small intrinsic sampling cost (Seale & Rapoport,1997). Another possibility is that in addition to search cost, subjects were affected by a goal that is more natural for them, for example a goal that involves satisficing (Todd & Miller, 1999), that is, an implicit goal to be happy about obtaining a high value item, instead of waiting out for the highest item to show up.

After training, the performance of our subjects got closer to the optimal strategy, for all of the distributions tested (Figures 3 and 4). Moreover, our subjects searched for cards slightly longer in positively-skewed environment than in the other two environments. This behavior contrasts with the optimal strategies, which have exactly the same search length pattern, irrespective of the underlying distribution of item quality. It is not totally clear why decision makers exhibit these behavior patterns. It is possible that our subjects exhibited these behaviors because of imperfect knowledge of the underlying distributions, and there was noise in sampling when they sampled to set and adjust their thresholds. A further direction to explore is plausible simple strategies that decision makers may employ (Gigerenzer, Hertwig, & Pachur, 2011) to select cards in a manner with monotonic decreasing thresholds. Such a strategy in theory could potentially come quite close to the optimal strategies in terms of success rate and efficiency.

## Acknowledgments

Part of this work was presented at the 2013 Mathematical Psychology Conference in Potsdam, Germany. We thank Gerd Gigerenzer, the ABC Research Group, and Peter Todd for very helpful feedback and ideas. This research was supported by German Academic Exchange Service (DAAD) grant 91540389 to TJW, and by Deutsche Forschungsgemeinschaft (DFG) grant 1713/2 to JDN as part of the priority program "New Frameworks of Rationality" (SPP 1516).

## References

Bearden, J. N., Rapoport, A., & Murphy, R. O. (2006). Sequential Observation and Selection with Rank-Dependent Payoffs: *An Experimental Study. Management Science*, *52*(9), 1437–1449.

Gardner, M. (1960a). Mathematical Games: A fifth collection of brain teasers. *Scientific American*, 150–155.

Gardner, M. (1960b). Mathematical Games: The games and puzzles of Lewis Carroll, and the answers to February's problems. *Scientific American*, 172–182.

Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The Foundations of Adaptive Behavior*.

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73.

Guan, M., & Lee, M. D. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the cognitive science society* (pp. 553–558). Austin, TX: Cognitive Science Society.

Helversen, B. V., Wilke, A., Johnson, T., & Schmid, G. (2011). Performance Benefits of Depression: Sequential Decision Making in a Healthy Sample and a Clinically Depressed Sample. *Journal of Abnormal Psychology, 120*(4), 962–968.

Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception & Psychophysics, 2*(8), 374–376.

Lee, M. D. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive science, 30*(3), 1–26.

Lee, M. D., O'Connor, T. A., & Welsh, M. B. (2004). Decision-Making on the Full Information Secretary Problem. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 819–824). Mahwah, NJ: Erlbaum.

Rapoport, A., & Tversky, A. (1966). Cost and accessibility of offers as determinants of optional stopping. *Psychonomic Science, 4*(1), 145–146.

Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance, 5*(2), 105–120.

Seale, D. A., & Rapoport, A. (1997). Sequential Decision Making with 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.

Todd, P. M., & Miller, G. F. (1999). From Pride and Prejudice to Persuasion: Satisficing in Mate Search. In G. Gigerenzer & P. M. Todd (Eds.),*Simple heuristics that make us smart* (pp. 287–308). Oxford University Press.

Wong, T. J., Nelson, J. D., & Schooler, L. J. (2016). Wait for the best, or settle for less? Search adapts to distribution shape in a value-based Secretary Problem. Manuscript submitted for publication.