

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Full-Information Optimal-Stopping Problems: Providing People with the Optimal Policy Does not Improve Performance

Permalink

<https://escholarship.org/uc/item/6cs782x5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Singmann, Henrik

Xiong, Yunxi

SONG, Yue

et al.

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Full-Information Optimal-Stopping Problems: Providing People with the Optimal Policy Does not Improve Performance

Henrik Singmann
(singmann@gmail.com)

UCL, Department of Experimental Psychology, 26 Bedford Way, London, WC1H 0AP, UK

Yunxi Xiong
(yunxi.xiong.22@ucl.ac.uk)

Yue Song
(yuesong00@outlook.com)

Mia Breen
(mia.breen@hotmail.com)
UCL, Experimental Psychology

Christiane Baumann
(cbaumann@hks.harvard.edu)
Harvard University, USA

Abstract

In optimal-stopping problems, people encounter options sequentially with the goal of finding the best one; once it is rejected, it is no longer available. Previous research indicates that people often do not make optimal choices in these tasks. We examined whether additional information about the task's environment enhances choices, aligning people's behaviour closer to the optimal policy. Our study implemented two additional-information conditions: (1) a transparent presentation of the underlying distribution and (2) a provision of the optimal policy. Our results indicated that while choice patterns varied weakly with additional information when providing the optimal policy, it did not significantly enhance participants' performance. This finding suggests that the challenge in following the optimal strategy is not only due to its computational complexity; even with access to the optimal policy, participants often chose suboptimal options. These results align with other studies showing people's reluctance to rely on algorithmic or AI-generated advice.

Keywords: optimal stopping; sequential decision-making; optimality, algorithmic advice

Introduction

Many decision in the real world are sequential in nature. For instance, when looking for an apartment, a job, or a partner, people typically consider their options one at a time. In this process, they must decide to accept or reject each option without the knowledge of whether future opportunities might be better. Decision-making problems of this nature – having to balance the value of the current option with the unknown value of future options and the inability to return to a previously rejected option – are known as *optimal stopping problems* (Gilbert & Mosteller, 1966; Hill, 2009).

Past research indicated that people struggle with sequential decisions, tending to settle on choices too soon when compared to the optimal solution (Bhatia et al., 2021; Goldstein et al., 2020; Lee et al., 2004; von Helversen & Mata, 2012; von Helversen et al., 2011). This study explores whether providing additional information to the decision maker, such as a complete overview of potential values or an indicator on the best time to make decisions, could assist people in making better decisions. The wider aim of this research is to assess people's willingness to integrate environmental cues or algorithmic advice when they need to make choices in a step-by-step manner.

Several studies have focused on two approaches of the optimal stopping problem: the rank-order (or no-information) and the full-information version. In the rank-order version

(Bearden et al., 2006; Seale & Rapoport, 1997) only the relative rank of an option compared to previous ones is provided. In contrast, the full-information approach (Lee, 2006; Lee et al., 2004) involves presenting the actual value of each option.

This paper focuses on the full-information version of the optimal stopping task. To illustrate, consider someone searching for the most affordable flight, where prices change daily. This person reviews the current price each day and must decide whether to buy or wait, knowing they can't return to a previously rejected offer and must purchase before their trip starts. Assuming the distribution of tickets is known, the best strategy in this scenario is to calculate the expected value of future options. From this, a decision threshold is established for each day. This threshold either steadily increases when seeking the lowest price or decreases when looking for the highest value. The decision rule is simple: if the current option's value falls below (for minimum) or exceeds (for maximum) this threshold, it is the right choice to make.

Previous research indicates that participants often stop their search too early (Guan & Lee, 2018; Lee, 2006; Lee & Courey, 2021), employing more relaxed decision thresholds than those suggested by the optimal policy. Furthermore, Baumann et al. (2020, 2023) found a significant shift in this pattern over the course of the decision sequence. While participants initially use decision thresholds that are too relaxed, surpassing a certain point in the sequence leads them to adopt stricter decision thresholds compared to the optimal ones, leading to an acceptance rate that is lower than the optimal level (as illustrated in Figure 1A, Baumann et al., 2020).

Given the computational complexity of the optimal policy for full-information problems, the failure of people to follow this policy is perhaps unsurprising. Importantly, not only is the calculation of the optimal threshold rather challenging, it also requires perfect knowledge of the probability distribution from which options are drawn. However, it remains unclear whether these are the specific aspects of these sequential decisions that people find challenging.

Our research aims to determine whether giving participants specific information about the statistical properties of the task can help them to more closely follow the optimal policy and ultimately make better decisions. To this end, we introduced two extra-information conditions. Firstly, we investigated whether showing the value distribution at each decision point leads to improved choices. If people's decisions align more

closely with the optimal solution when they have this information, it would suggest that a lack of knowledge into the search environment is an important reason for their deviation from the optimal strategy. Secondly, our study investigates whether clearly indicating the optimal threshold at each step enables participants to use this information for more effective choices. If providing people with the optimal thresholds would lead people to behave more optimally, it would suggest that the computational complexity of the optimal policy is a reason why people diverge from optimality. To foreshadow our results, while we observed some minor variations in the patterns of choices, there were no significant differences in overall performance between the conditions. This indicates that providing additional information about the search environment, such as a description of the underlying distribution or even the optimal threshold, may affect people's choices to some extent, but it does not necessarily aid them in making better decisions.

Optimal Stopping Task and Optimal Policy

In an optimal stopping task, the goal is to identify the most attractive value within a sequence of options. We focus on the full-information version, in which the distribution of ticket prices is assumed to be known to the decision-maker, and the task of purchasing a plane ticket, as previously mentioned, serves as an illustrative example of this type of problem. We now provide a more formal explanation of this task and its optimal policy.

We consider a decision-maker who encounters a sequence of 1 to N options with rewards denoted by x_1, \dots, x_N drawn from a known distribution X , with probability density function $f(x)$ and cumulative distribution function $F(x)$. Their goal is to find the minimum value in the sequence. If the decision-maker accepts option i , then the sequence terminates and they receive reward x_i ; otherwise, they continue to the next option. When the last option N is reached, it must be accepted. The optimal policy is to choose option i when it goes below a position-dependent threshold t_i . As shown by Gilbert and Mosteller (1966), the optimal policy depends only on the distribution of rewards and the number of remaining choices.

The goal of the optimal policy for full-information problems is to maximise the expected payoff. To do so, a position specific optimal threshold is calculated starting from the last position of the sequence. If the last position is reached, the corresponding option x_N has to be accepted – which corresponds to a random draw from distribution X . Hence, the expected payoff on the last position is the expected value of the distribution $E_N[X] = \int_{-\infty}^{\infty} xf(x) dx$ (i.e., the distribution's mean). At the second-to-last option, x_{N-1} , the expected payoff can be better than the expected value of the distribution. More specifically, because the expected payoff at the last position N is $E_N[X]$, one should only accept option x_{N-1} if it is lower than $E_N[X]$ – the threshold for the second-to-last position is $t_{N-1} = E_N[X]$. Based on this threshold, we can now also calculate the expected payoff for the second-to-last position; it is (a) the probability of obtaining an option better than

the threshold times (b) the conditional expected value of such an option plus (c) the probability of obtaining a value worse than the threshold times (d) the expected payoff of the last position (i.e., the current threshold),

$$E_{N-1}[X] = \underbrace{F(t_{N-1})}_{(a)} \times \underbrace{\int_{-\infty}^{t_{N-1}} xf(x) dx}_{(b)} + \underbrace{(1 - F(t_{N-1}))}_{(c)} \times \underbrace{t_{N-1}}_{(d)}.$$

Following the same logic as above, the expected payoff of the second-to-last position is then the optimal threshold for the third-to-last position, $t_{N-2} = E_{N-1}[X]$. Using this threshold we can calculate the expected payoff for the third-to-last position, $t_{N-3} = E_{N-2}[X]$, using the equation above and replacing t_{N-1} with t_{N-2} . To calculate the optimal thresholds for all positions we need to repeat this process until reaching the threshold for the first position. Put succinctly, calculating the optimal threshold of the optimal policy requires backwards induction and repeatedly integrating and evaluating the probability distribution.

The Current Study

The aim of this study is to investigate if giving participants more information about the underlying distribution and the optimal threshold can improve their decision-making, bringing it closer to the optimal solution. To this end, we conduct a comparison between three conditions. The first is the *Baseline condition*, representing the classical optimal stopping task as previously implemented. In the *Distribution condition*, we introduced a visual representation to show the distribution from which the options (tickets) are drawn. Alongside this, we displayed the current option's placement within that distribution. In the *Optimal condition*, we went a step further by showing participants the optimal thresholds, and indicating whether the current option is more favorable or less favorable than these thresholds.

Our study uses the *Ticket Shopping Task* introduced by Baumann et al. (2020). In the Ticket Shopping Task, participants' task is to find the cheapest plane ticket within a sequence of a fixed number of tickets, presented one after the other. For each ticket, participants must decide whether to accept that ticket and end their search, or to reject it and continue looking for a better option. As in every optimal stopping problem, participants cannot return to previously rejected tickets, they have to accept the last ticket in a sequence, and – once a ticket is accepted – the current sequence ends and the next sequence starts. The task is performed in a fully incentive compatible manner using a proportional payoff function; the cheaper the ticket participants get in every sequence, the higher their final payout.

Individual tickets in the Ticket Shopping Task are drawn from a fixed distribution that stays the same for each ticket in every sequence. In order to avoid learning only during the task, an inherent feature of the full-information optimal stopping task, participants are exposed to samples from this distribution beforehand during a training phase.

Figure 1: Screenshots of the Test Phase of the Three Experimental Conditions



Note. A) Baseline condition. The information available to participants is the current ticket price and the position in the current sequence. B) Distribution condition. In addition to the information in the Baseline condition, participants are shown the distribution of ticket prices as a histogram and the position of the current ticket price in the distribution (in red). C) Optimal condition. In addition to the ticket price distribution and the current ticket price, participants are also shown the position dependent optimal threshold in black. If the current price is better than the threshold, it was shown in green (as here), otherwise it was shown in red. *Not shown:* The current trial number was displayed in the top left throughout in all conditions: “Sequence x of 200”.

Methods

Participants

A total of 117 participants were recruited through Prolific. All participants were included in the final analysis, because none failed the inclusion criterion of Baumann et al. (2020), not accepting the first option in a sequence in more than 95% of trials. Participants were randomly assigned to the three conditions, 48 participants to the Baseline condition, 34 to the Distribution condition, and 35 to the Optimal condition.

Participants’ compensation was determined completely by their performance. The total payout (in \$0.01) was calculated based on their mean accepted ticket price at the end of the experiment, $\text{payout} = 10 \times (240 - \text{the mean accepted ticket price})$. This formula ensured that with a mean accepted ticket price equal to the distribution mean of \$180, their payout would be \$6.00, which was our advertised minimum payment for the experiment. The mean payout was \$8.23, ranging from \$6.51 to \$8.67. Participants spent on average 37.5 minutes on this task. Our study received full ethical approval from the UCL Experimental Psychology ethics committee.

Design and Procedure

The experiment was comprised of two separate phases, a training and a test phase. The training phase served to familiarise participants with the distribution of ticket prices. In the test phase, participants performed the optimal stopping task. All participants started with the training phase and subsequently proceeded to the test phase.

The distribution of ticket prices used throughout the whole experiment was a normal distribution with a mean of \$180 and standard deviation of \$20, truncated at a minimum of \$120 and a maximum of \$240 (i.e., $\text{mean} \pm 3 \text{ SD}$).¹ Each ticket presented to participants was independently selected

¹For the current design, optimal thresholds were practically indistinguishable for truncated and unconstrained normal distribution.

from the same distribution, meaning the distribution of ticket prices remained consistent throughout all positions in a sequence and between the learning and test phases. This information was clearly communicated to the participants.

Training Phase. The training phase was divided into two stages. In the first stage, participants’ encountered samples from the specified distribution and their task was to estimate their average price. More specifically, they were shown five series of airplane ticket prices and each series contained ten tickets. Participants had to click a ‘Next’ button to move from one ticket to the next. After viewing one set, they had to estimate the average price of the tickets and were then provided with the correct answer along with the difference between their answer and the actual average price.

In the second stage, participants were asked to estimate the distribution of a sample of 100 ticket prices, drawn from the same distribution they encountered in the first stage of the training phase. The task required them to determine how many ticket prices fell into each of seven predefined price ranges. These ranges spanned from \$120 to \$240 in seven equally sized intervals. To make their estimates, participants could either adjust the height of bars on a histogram or enter numbers directly into designated input fields, each corresponding to a specific price range. After distributing all tickets and submitting their answer, participants received feedback in the form of the correct distribution, which was superimposed over their estimate (Goldstein & Rothschild, 2014).

Test Phase. In the test phase, participants completed a total of 200 trials. In each trial, they were presented with a sequence of ten tickets with a price independently drawn from the ticket price distribution. Participants’ goal was to select the cheapest ticket in each trial. At each position in the sequence, participants could either accept the current ticket or proceed to the next ticket (see examples in Figure 1). Once

a ticket was rejected, participants could not return to it. If the final 10th position in a sequence was reached, participants had to accept the ticket. After accepting a ticket in a sequence, participants received feedback regarding whether the chosen ticket was the cheapest one in the sequence. Additionally, after every 50 trials, participants were updated on their progress (the fraction of sequences they had completed) and earnings. The current trial number and the total trial number were shown throughout the whole experiment.

Whereas the training phase was identical for all experimental condition, the test phase differed across conditions as shown in Figure 1. Participants also received condition specific instructions prior to the test phase.

In the Baseline condition, instruction reminded participants that tickets were drawn randomly from the distribution of ticket prices and that this entailed that ticket prices within a sequence were unrelated. During the test trials, participants only knew the current ticket price and the position in the sequence (Figure 1A).

In the Distribution condition, participants were additionally informed that “to make the task easier” they would see the distribution of ticket prices as a histogram as well as the current ticket price. As in the Baseline condition, they were reminded: “Whereas the current ticket price changes for every ticket, the distribution always stays the same”. The histogram as well as the current ticket price in the distribution were shown throughout (Figure 1B).

In the Optimal condition, participants were informed that they would receive “the strategy of an optimal agent (which you can imagine as a smart AI that knows how to play the task)”. They were also told that the optimal strategy uses an optimal threshold that changes with every position and how to use the threshold (accept if better, otherwise reject). The optimal threshold, the current ticket price, and the corresponding decision were shown on a histogram of the ticket price distribution (Figure 1C). Participants were furthermore informed that: “Using the optimal threshold does not guarantee that the cheapest ticket is found in each sequence.” And: “It is not necessary for you to follow the guidance of the optimal threshold. Feel free to use your own judgement.”

Results

Global Summary Statistics

To compare performance across condition, we first considered two summary statistics describing overall performance, the mean accepted ticket price and the mean search length. Results of these two statistics are shown in Figures 2A & 2B. The black points show participants’ overall averages and look quite similar across conditions. Furthermore, by comparing black and red data points we can compare participants’ actual performance with the performance they could have had had they followed the optimal policy. This suggests that, as in previous studies, participants stop too early (Figure 2B) resulting in a higher (i.e., worse) mean accepted ticket price than when following the optimal strategy (Figure 2A). How-

ever, the difference between observed behaviour and optimal policy for search length appeared to be minimal for the Optimal condition.

To further corroborate this picture, we performed two ANOVAs, one for each of the two summary statistic, each with a single independent variable, condition, with three levels. In line with the visual impression, in neither of these ANOVAs the effect of condition reached significance (ticket price: $F(2, 114) = 0.99$, $p = .374$; search length: $F(2, 114) = 1.51$, $p = .226$). The same result was also observed after removing one participant from the Control condition with relatively poor performance (i.e., mean accepted ticket price ≈ 175 ; $p = .248$ and $p = .204$, respectively).


In a next analysis steps, we checked whether diverging from the optimal policy actually led to a significant decrease in performance. For this, we ran a second set of ANOVAs, one for each summary statistic, that also included the performance of an agent that perfectly follows the optimal policy. Each of these ANOVAs now had two independent variables, condition (with three levels, between-subjects), agent (human vs. optimal, within-subjects), as well the condition \times agent interaction. These ANOVAs revealed a highly significant effect of agent (ticket price: $F(1, 114) = 131.81$, $p < .001$; search length: $F(1, 114) = 27.35$, $p < .001$). Had participants followed the optimal policy their mean accepted ticket price would be 155 compared to the actual 158, which would have resulted in an on average \$0.29 higher payout. Likewise, participants would have searched for an additional 0.47 positions under the optimal policy. Importantly, there were no significant condition \times agent interactions (ticket price: $F(2, 114) = 0.33$, $p = .719$; search length: $F(1, 114) = 1.84$, $p = .164$), suggesting that this pattern held across all conditions. Furthermore, there were again no main effects of condition ($ps > .2$). As before, the results pattern remained the same after removing the one participant with poor performance from the Baseline condition.

Taken together, the analysis of the summary statistics showed no discernible differences across conditions. Providing participants with extra information about the distribution or even showing them the optimal strategy did not improve their overall search behaviour or increased their payout. Furthermore, participants in all conditions diverged significantly from the optimal policy and this was to their financial detriment. Finally, there was no statistical evidence corroborating the impression that the difference between observed and optimal behaviour in the Optimal condition was attenuated compared to the other two conditions.

Behaviour Across Positions

To further explore whether the extra-information conditions really did not improve performance, we also looked at participants’ probability to accept at each position of the sequence. These results are shown in Figure 2C separately for each condition. In each condition we can see a pattern already observed by Baumann et al. (2020). Participants’ increase in acceptance probabilities across positions is linear, which con-

Figure 2: Participants' Performance Across Conditions and Comparison with Optimal Policy



Note. A) Mean accepted ticket price across conditions. Each grey point shows the mean accepted ticket price of one participant, the black point shows the overall mean with associated 95% confidence interval. The red \times s shows the mean performance if participants had followed the optimal policy. B) Mean search length across conditions. As in panel A, grey points are by-participant averages, black points are the overall average, and red \times s are the optimal policy. C) Mean probability to accept a ticket at each position per condition. Black lines show participant averages (with 95% CI) and red lines show the performance had participants followed the optimal policy (with 95% CI). The grey line shows the asymptotic (i.e., analytical) optimal policy.

trasts with the clearly non-linear increase prescribed by the optimal policy. This linear pattern in the observed behaviour results in participants both searching too little and too much. At early positions (i.e., before position 6), participants are generally risk averse, they search too little and stop too early. However, once participants arrive at the end of the sequence their behaviour switches and they become risk seeking, they now search too much and stop too late. Whereas this pattern is present in all conditions it appears to be attenuated in the Optimal condition, especially for early positions.

To further corroborate this pattern we analysed participants by-position acceptance probabilities with an ANOVA with two factors, condition (with three levels, between-subjects), position (9 levels, within-subjects), as well as the condition \times position interaction. As before, we did not see a significant main effect of condition, $F(2, 114) = 1.44$, $p = .240$. We also saw an unsurprising main effect of position $F(5.07, 578.44) = 60.64$, $p < .001$.² Furthermore, we

found a just significant condition \times position interaction, $F(10.15, 578.44) = 1.90$, $p = .040$, indicating that the acceptance behaviour differed across positions.

To get at the source of this interaction we compared the linear trends of positions across conditions using model-based follow-up tests of the interaction. These comparisons showed that the linear trend in the Optimal condition was steeper than in the Baseline condition, $t(114) = 2.14$, $p = .035$. However, there were no significant difference in the linear trend between the Optimal and the Distribution condition, $t(114) = 1.71$, $p = .090$, and the Distribution and the Baseline condition, $t(114) = 0.28$, $p = .778$. These follow-up tests were performed without controlling for multiple testing.

A different way to look at this interaction is to look at the pairwise comparisons across conditions for each position separately. When doing so (again without controlling for multiple testing), we find that for position 3 the Distribution condition has a significantly larger acceptance probability than both the Baseline condition, $t(114) = 2.31$, $p = .023$, as well as the Optimal condition, $t(114) = 2.94$, $p = .004$. For position

²Degrees of freedoms of effects involving within-subjects factors with more than two levels are Greenhouse-Geisser corrected.

4, the Optimal condition has a significantly lower acceptance probability than both the Baseline condition, $t(114) = 3.16$, $p = .002$, as well as the Distribution condition, $t(114) = 2.73$, $p = .007$. None of the remaining pairwise comparisons reached significance, $ps > .08$. This suggests that around positions 3 and 4, the Optimal condition is more similar to the optimal policy than the other two conditions (Figure 2C).

Taken together these results suggest that there is no discernible difference in the probability to accept across positions between the Baseline condition and the Distribution condition. However, there is some evidence for a difference between the Optimal condition and the Baseline condition. Comparing the behaviour of participants in the Optimal condition to the optimal policy we can see that participants appear to be almost in line with the optimal policy for earlier positions (i.e., appear less risk averse than the other two conditions). However, for later positions they show the same risk seeking behaviour also exhibited by the other two conditions.

Discussion

The aim of this study was to test whether providing participants with additional information in full-information optimal stopping tasks with proportional payoff function and known time horizon affects their search behaviour. We provided additional information in two levels. In the Distribution condition participants were shown the distribution from which options were sampled throughout as well as value of the current option in this distribution. In the Optimal condition, we additionally provided participants with the optimal threshold and the decision based on the optimal policy.

The results showed that the behaviour in the Distribution condition, where the distribution was shown in each search step, was essentially indistinguishable from the Baseline condition, both in terms of overall summary statistics as well as in terms of the behaviour across positions. This suggests that in optimal stopping problems in which participants have learned the distribution of options beforehand, additionally providing participants with a graphical representation of the distribution and the position of the current options in this distribution does not change their behaviour. In other words, the reason people diverge from the optimal policy in such optimal stopping problem is not because they do not know the distribution of options well enough.

There was some evidence to suggest that participants in the Optimal Condition behaved more in line with the optimal policy than participants in the other two conditions. When looking at the probability to accept an option across positions, the Optimal condition was more similar to the optimal policy for early positions – people in the Optimal condition were less risk averse than people in the other two conditions. However, for later positions there was no difference between conditions and we saw the same risk seeking behaviour in all conditions. This suggests that directly providing participants with the optimal policy helps people a bit in improving their

performance, but they are still not fully adopting the optimal policy and behave worse than they could.

Given these (admittedly small) differences in search behaviour it is surprising that we did not see any global differences in overall performance. The payoff and mean search length did not differ significantly between participants in the Optimal condition compared to participants in the other two conditions. We believe the most likely explanation for this pattern is the finding that while participants improve their search behaviour for early positions, they still exhibit risk seeking for later positions. By searching more in early position they increased their chance of finding better tickets in later positions, but by failing to adapt their search strategies in later position they failed to benefit from this change in search behaviour. Instead, just like in the other conditions they too often reached the final position. A focused test of this possible explanation might be a fruitful avenue for future research.

Another possibility is that with a more focused design and more statistical power one might still find a (small) difference in payoff between the Optimal and other conditions. However, even if this were the case this would not vindicate the performance of the Optimal condition. Our results show a significant difference between the payoff in the Optimal condition and an agent that perfectly follows the optimal policy. Thus, any possible effect of directly providing participants with the optimal policy is small compared to the effect of participants voluntarily diverging from the optimal policy.

Even though we found significant differences between participants' behaviour and an optimal agent when providing the optimal policy, these differences were overall small. Compared to the optimal policy participants lost on average \$0.29. In contrast, an agent who responds randomly in our task, resulting in a mean accepted price of 180, would lose \$2.50 compared to the optimal solution. This contrasts with other decision-making tasks where not following the optimal policy results in dramatic earning losses (Camerer, 2011; Goeree & Holt, 2001; Tversky & Kahneman, 1992). This suggests that the heuristics used by people to solve optimal-stopping tasks appear to be clearly resource rational (Lieder & Griffiths, 2020), at least in the Baseline and Distribution condition.

Our main result, showing that people often ignore the provided optimal policy, are in line with findings showing that people tend not to accept algorithmic or AI advice (e.g., Dawes et al., 1989; Dietvorst et al., 2015). Interestingly, the strongest evidence for people's unwillingness to accept algorithmic advice comes from the medical domain (Longoni et al., 2019), in less critical domains (e.g., music recommendations) people in contrast have a tendency for appreciating algorithmic advice (Logg et al., 2019). An interesting avenue for future research is whether people's unwillingness to accept advice in optimal-stopping problems is domain-specific or whether this generally applies to sequential decision-making tasks.

Acknowledgments

We are thankful to Maggie Xie for her help performing the research reported here. Henrik Singmann was supported by Grant 100014.179121 from the Swiss National Science Foundation (SNSF).

References

- Baumann, C., Schlegelmilch, R., & von Helversen, B. (2023). Adaptive behavior in optimal sequential search. *Journal of Experimental Psychology: General*, *152*(3), 657–692.
- Baumann, C., Singmann, H., Gershman, S. J., & Helversen, B. v. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755.
- Bearden, J. N., Rapoport, A., & Murphy, R. O. (2006). Experimental studies of sequential selection and assignment with relative ranks. *Journal of Behavioral Decision Making*, *19*(3), 229–250.
- Bhatia, S., He, L., Zhao, W. J., & Analytis, P. P. (2021). Cognitive models of optimal sequential search with recall. *Cognition*, *210*, 104595.
- Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton university press.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, *243*(4899), 1668–1674.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114–126.
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, *61*(313), 35–73.
- Goeree, J. K., & Holt, C. A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, *91*(5), 1402–1422.
- Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, *66*(3), 1375–1394.
- Goldstein, D. G., & Rothschild, D. (2014). Lay understanding of probability distributions. *Judgment and Decision Making*, *9*(1), 1–14.
- Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361.
- Hill, T. (2009). Knowing when to stop. *American Scientist*, *97*(2), 126.
- 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., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*, *4*(1), 1–17.
- Lee, M. D., O'Connor, T., & Welsh, M. (2004). Threshold models of human decision making on optimal stopping problems in different environments. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 819–824).
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103.
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, *46*(4), 629–650.
- 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.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, *5*, 297–323.
- von Helversen, B., & Mata, R. (2012). Losing a dime with a satisfied mind: Positive affect predicts less search in sequential decision making. *Psychology and aging*, *27*(4), 825.
- von Helversen, B., Wilke, A., Johnson, T., Schmid, G., & Klapp, B. (2011). Performance benefits of depression: Sequential decision making in a healthy sample and a clinically depressed sample. *J Abnorm Psychol*, *120*(4), 962–968.