Experimental Economics (2024) 27:854–873
https://doi.org/10.1007/s10683-024-09832-2

ORIGINAL PAPER

Strategies people use buying airline tickets: a cognitive modeling analysis of optimal stopping in a changing environment

Michael D. Lee[1] · Sara Chong[1]

Received: 3 December 2023 / Revised: 17 May 2024 / Accepted: 27 May 2024 /
Published online: 12 June 2024

Abstract
We study how people solve the optimal stopping problem of buying an airline ticket. Over a set of problems, people were given 12 opportunities to buy a ticket ranging from 12 months before travel to 1 day before. The distributions from which prices were sampled changed over time, following patterns observed in industry analysis of flight ticket pricing. We characterize the optimal decision process in terms of a set of thresholds that set the maximum purchase price for each time point. In a behavioral analysis, we find that the average price people pay is above the optimal, that there is little evidence people learn over the sequence of problems, but that there are likely significant individual differences in the way people make decisions. In a model-based analysis, we propose a set of nine possible decision strategies, based on how purchasing probabilities change according to time and the price of the ticket. Using Bayesian latent-mixture methods, we infer the strategies used by the participants, finding that some use purely time-based strategies, while others also attend to the price of the tickets. We conclude by noting the limitations in the strategies as accounts of people's decision making, highlighting the need to consider sequential effects and other context effects on purchasing behavior.

Keywords Optimal stopping · Cognitive modeling · Changing environments · Strategic decision making

JEL codes D9 · D90 · C91

✉ Michael D. Lee
mdlee@uci.edu

[1] Department of Cognitive Sciences, University of California Irvine, Irvine, CA 92697-5100, USA

1 Introduction

Buying an airline ticket is a familiar optimal stopping problem. The goal is to minimize the cost of the ticket, but this is made difficult by changes in the price over time. Part of the change in ticket prices is unpredictable fluctuation, but part is a predictable change in the price distribution, which notoriously increases rapidly as the day of travel approaches. Managing this uncertainty is the key to good decisions, since if a cheap ticket is not purchased it is not possible to go back in time, but once a ticket is purchased future prices are generally not available.

One way people could solve this problem is in terms of *time*. Classic "secretary" optimal stopping problems that present only rank information are solved this way (Ferguson, 1989). The optimal process to follow is to wait until $1/e \approx 37\%$ of the options have been viewed, and choose the next one that is best, if there is one. How people make decisions for this sort of optimal stopping problem has been widely studied in psychology, including seminal work by Amnon Rapoport and collaborators (Seale & Rapoport, 1997, 2000; Bearden et al., 2006).

Another way people could solve the flight ticket purchasing problem is in terms of *value*. Full-information secretary problems, which present the values of options and not just their ranks, are solved this way. The optimal process is to accept or reject options in terms of threshold values. The optimal thresholds depend on the distributions that the option values come from, the number of options, and the loss function being optimized (Gilbert and Mosteller, 1966). How people make decisions for this sort of optimal stopping problem has also been studied in psychology (e.g., Baumann et al., 2020; Goldstein et al., 2020; Lee, 2006), but rarely for changing distributions. One early exception is Shapira and Venezia (1981) who studied people's behavior when the means of the distributions from which values were drawn increased or decreased linearly over the sequence. They concluded that people were sensitive to the nonstationary environment, but not necessarily in optimal ways. In particular, they argued people tended to choose too late in the descending environment and too early in the ascending environment. A more recent study of optimal stopping in changing environments is provided by Lee and Courey (2021), who considered people's decision making in mate choice problems. The Gaussian distributions of values for potential mates changed in terms of both mean and variance as they aged. Based on information from dating sites, the means of the distributions changed nonmonotonically: an initial increase was followed by a sustained decrease. Lee and Courey (2021) also found that people were sensitive to the nonstationary environment, and argued that there was evidence people use relatively simple cognitive strategies to set thresholds that guided their choices

The problem of buying an airline ticket has elements of both of these previous studies.[1] Like the mate choice problem, there is evidence in both the economics

[1] There are experiments in the literature studying human decision making on optimal stopping problems that also use an airline ticket cover story (Baumann et al., 2020, 2023). In these experiments, though, the distribution from which ticket prices are drawn does not change over the course of a problem. Participants are given instruction and training so that they understand this feature of the experiment.

literature (e.g., Bilotkach et al., 2015, Table 1) and industry and media analysis (e.g., McCartney, 2014) that the changes in price are nonmonotonic. An initial premium for booking well ahead gives way to cheaper tickets a few months before the flight, and then the mean price increases quickly as the day of travel approaches. While Lee and Courey (2021) did consider a nonlinear environment structure, they required participants to choose only the best possible mate. This zero–one loss, corresponding to "only the best will do", seems inappropriate for buying airline tickets. Instead, the expected value utility used by Shapira and Venezia (1981) seems appropriate. Thus, studying how people solve the optimal stopping problem of buying airline tickets requires a novel combination of a non-monotonically changing environment with the goal of minimizing expected cost.

The goal of the current work is to study people's decision making in the flight ticket optimal stopping problem. The remainder of this article is as follows. In the next section, we describe an experiment that measures how people buy airline tickets, and characterize how optimal purchasing decisions are made in this environment. We then consider basic empirical regularities in people's behavior, and develop a set of nine cognitive strategies, based either on time or value, as possible accounts of people's decision making. We apply the strategies to the behavioral data, and find evidence for at least some use of most of the strategies. A majority of people appear to use threshold-based strategies based on the price of tickets, but some also make purchases primarily according to the time before travel. We discuss the limitations of the strategies as complete accounts of people's decision making, and suggest future research directions to overcome these limitations.

2 Experiment

2.1 Participants

A total of 46 participants recruited from the Prolific platform completed the experiment. There were 29 females, 16 males, and one non-binary or third gender participant. The average age of the participants was 37 years, ranging from a minimum of 20 years to a maximum of 61 years.

2.2 Problems

The environment we used was based on industry analysis[2] and is shown in Fig. 1. We chose 12 time points at which ticket prices were available in each problem: 12 months, 9 months, 6 months, 3 months, 2 months, 1 month, 2 weeks, 1 week, 5 days, 3 days, and 1 day before the flight. These time points were deliberately chosen to be unevenly spaced in time, in an attempt to reflect a pattern representative of the way people pay attention to changes in price as the day of travel approaches.

[2] See https://www.cheapair.com/blog/the-best-time-to-buy-flights/.

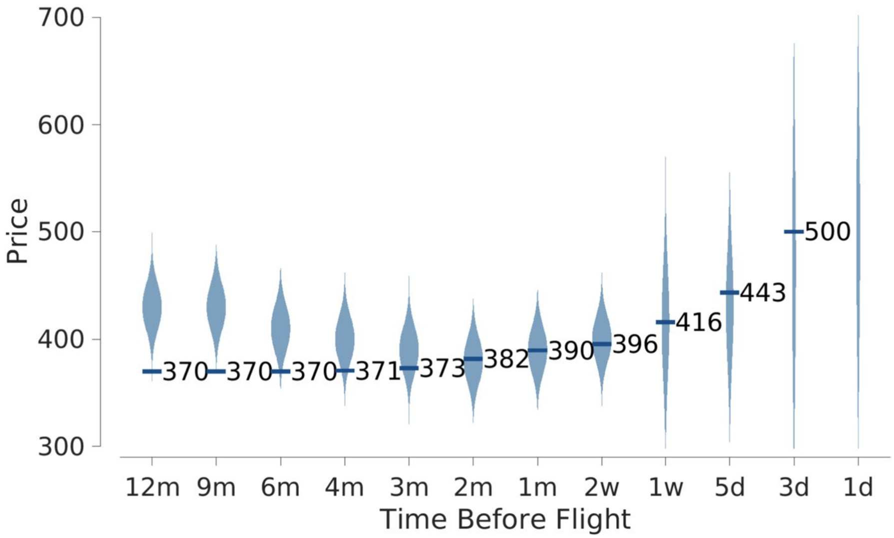

Fig. 1 The violin plots show the distribution of possible flight price options from 12 months to 1 day before travel. The numbered vertical lines show the optimal thresholds for minimizing the expected price

Figure 1 shows with violin plots the truncated Gaussian distributions from which specific ticket prices were sampled at each of these time points. The modes[3] of the distributions change nonlinearly, with a slight decrease from 6 months to 2 months followed by a sharper increase toward the day of travel. The variance is stable until 1 week before travel, at which time it increases significantly.

The numbered horizontal lines in Fig. 1 show the optimal thresholds that minimize expected price for the environment. They are initially flat around \$370, but then increase about 2 months before travel. These thresholds were calculated using the methods developed by Gilbert and Mosteller (1966, Section 5). The basic approach is to reason backwards, starting with the second last option in the sequence, which is the last time a decision can be made. A purchase should be made 3 days before travel if the price on that day is less than the expected price of a ticket 1 day before travel. In other words, a purchase should be made if the expectation is that a more expensive ticket will be purchased by not purchasing. This defines the threshold maximum price 3 days before travel. Assuming adherence to this optimal threshold further defines an expected price paid from 3 days before travel onward: it depends on the probability a price 3 days before travel will be below the threshold, the expected price below that threshold according to the price distribution for 3 days before travel, and the price distribution for 1 day before travel in the case that the price 3 days before travel is too high and is rejected. Given this expectation from 3 days before travel onward, the same argument can be applied to determine the

[3] Because the Gaussian distributions are truncated between \$300 and \$700, the parameter that is usually the mean is technically the mode of the distribution.

threshold for 5 days before travel. A purchase should be made 5 days before travel if the price on that day is less than the expected price of a ticket that would follow by not purchasing. This reasoning can continue to be applied recurrently to find optimal thresholds for all of the options in the sequence, as shown in Fig. 1.

We developed a set of 50 optimal stopping problems based on this environment. The set of problems was chosen to be representative with respect to the optimal decision process. To do this, we started by determining the distribution of how often each option is chosen by the optimal decision process, which was calculated by applying the process to a very large number of problems generated from the environment shown in Fig. 1. We then searched for a set of 50 problems for which the optimal decision process chooses options that follow the same distribution as closely as possible. The agreement achieved is very good, and is shown by a figure in the supplementary material. Making the 50 problems representative in this way helps justify comparing people's decisions to optimal decisions, because it means that the optimal thresholds shown in Fig. 1 apply similarly to the specific problems that people encountered.

2.3 Procedure

Each participant completed the 50 representative optimal stopping problems in a random order. Instructions were provided at the beginning of the experiment explaining that the goal was to minimize the price paid for an airline ticket, that there were 12 opportunities to purchase, and that the distribution from which prices were drawn changed over the possible purchase times. It was emphasized that once a ticket was purchased there was no opportunity to change the purchase, but that if no ticket was purchased by the time it was 1 day before travel, the final ticket price available on that day had to be purchased.

At each time point during a problem, participants were given a ticket price (in whole dollars) that they could purchase or reject. In addition they were shown the distributions of all 12 time points, with the distribution for the current time point highlighted. That is, they saw all of the 12 blue distributions in Fig. 1, but with the current time point highlighted in yellow.

Once a ticket was purchased, the next problem started. No feedback was provided about the ticket prices that would have been available, nor was any feedback provided about the accuracy of the purchase decision or the average price paid for tickets. The median time for participants to complete the experiment was 18 min, with an interquartile range of 13–27 min.

2.4 Payment

Participants were paid a flat hourly rate of $8 that was not contingent on performance. This is standard practice in many fields of experimental psychology, but less common in experimental economics (Camerer & Hogarth, 1999; Hertwig & Ortmann, 2001). Some evidence that it is a reasonable approach is provided by Campbell and Lee (2006), who found that financial rewards determined by a

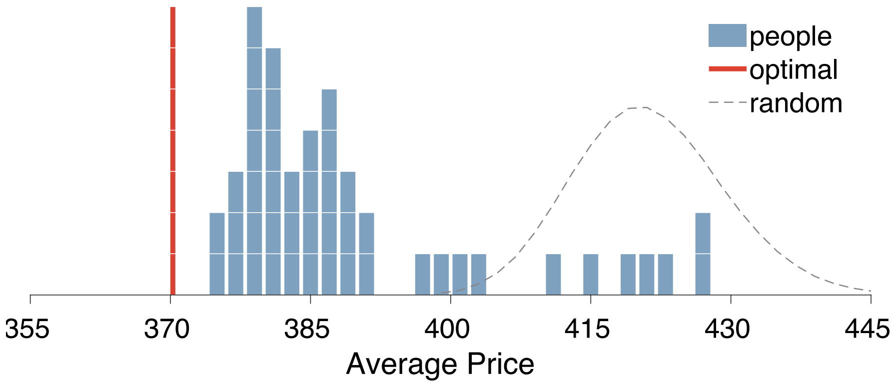

Fig. 2 The blue histogram shows the distribution of the average ticket price paid by the participants. The red line is the average paid by following the optimal decision process. The dotted line shows the distribution of average ticket price paid if a random ticket among the options for each problem is chosen

performance-contingent quota scheme (e.g., Bonner et al., 2000) did not change people's performance on optimal stopping problems.

3 Empirical results

Figure 2 provides a simple analysis of people's performance. It shows the distribution of the average price paid over all 50 problems by each participant. Also shown is the price paid by following the optimal decision process, and the distribution of average price paid by buying at random one of the 12 possible prices for each problem. The results suggest that people make suboptimal but thoughtful decisions. People are generally paying about $10–$20 per ticket more than optimal, but their average prices are too good to be explained by random choices. This conclusion holds for a different conception of thoughtless behavior in which one of the first three available tickets is chosen at random. The distribution of average price for this behavior is very similar to the random one shown in Fig. 2.

Figure 3 shows the change in the average ticket price purchased by individual participants, and the average over all participants, as their sequence of 50 experimental problems progressed. There is no evidence of learning for either individuals or overall. This is consistent with many (e.g., Lee, 2006, Figure 7; Guan et al., 2014, Figure 1; Guan & Lee, 2018, Figure B2; Lee & Courey, 2021, Figure 3), but not all (e.g., Goldstein et al., 2020), previous findings testing for learning in optimal stopping experiments.

Figure 4 provides an analysis of participants' decision making. For each time point, it shows the distribution of ticket prices purchased versus rejected, as well as the optimal threshold for each time point. It is clear that most tickets are purchased among the first six or seven offered, up to about 1 month before travel. For many time points, such as 6 months or 2 months before travel, the distribution of

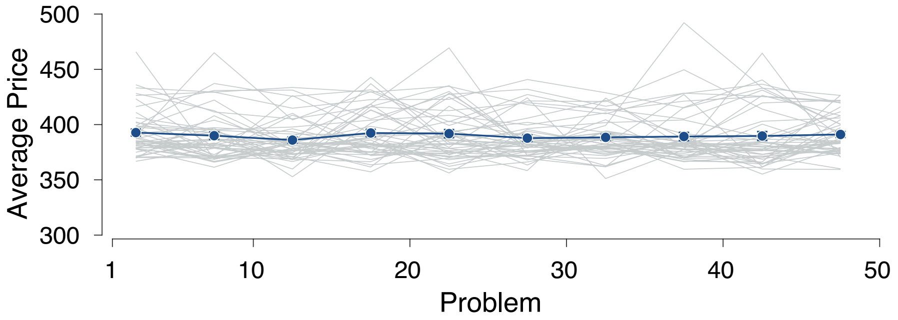

Fig. 3 The lack of learning across problems. The gray lines correspond to the average price paid by each individual participant as they completed the 50 experimental problems. The blue line shows the average price paid across all participants

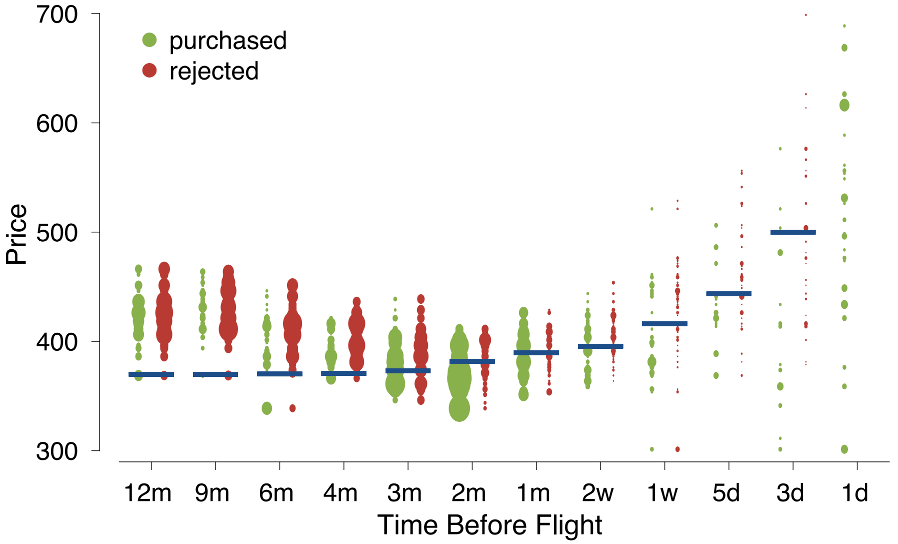

Fig. 4 The green distributions show for each time point how often an airline ticket at a given price was purchased by a participant. The red distributions show ticket prices that were rejected. The horizontal lines show the optimal thresholds to minimize the expected price paid

purchased prices is a little cheaper than those rejected. But this effect is not strong, and there is significant overlap between the distribution for most time points. This suggests that people's decision making is based on factors other than comparing the current price to a threshold. It also makes clear that people are not following the optimal thresholds. For most time points there are many purchased tickets with prices significantly more expensive than the optimal threshold.

The behavior shown in Fig. 4 is unlike the behavior in other optimal stopping tasks for which the same analysis is available (Guan & Lee, 2018, Figure 3; Lee & Courey, 2021, Figure 2; Guan et al., 2014, Figures 2, 3). These earlier analyses showed a clearer separation between selected and rejected values and provided evidence of the use of thresholds. The current data appear more complicated in this regard, perhaps because of the different expectation-based utility function rather than zero-one loss.

Reconciling participants' reasonably effective purchasing behavior in Fig. 2 with the overlap in purchased versus reject ticket prices in Fig. 4 can potentially be achieved by allowing for individual differences in the decision strategies used. It is logically possible that the overlapping distributions in Fig. 4 arise from decision making based on thresholds, but with different participants using different ones. It seems more likely that the individual differences also involve qualitatively different strategies. Accordingly, in the next section we pursue a model-based analysis using a set of different decision strategies people could use.[4]

4 Models of human decisions

In this section we develop nine possible strategies based on either time or value. We formalize all of the strategies as generative probabilistic models to make them compatible with fully Bayesian inference. A visual schematic summarizing each strategy is shown in Fig. 5. Time strategies are based on probabilities of purchasing a ticket at each time point and are depicted as bar graphs that represent the probabilities. Value strategies are based on threshold prices for tickets at each time point and are depicted by lines that represent the thresholds.

4.1 Time strategies

4.1.1 Constant probability strategy

The constant probability strategy assumes there is the same probability of a ticket being chosen at each point in time. The strategy can be formalized by:

$$y_{ijk} \sim \text{Bernoulli}(\theta_i) \tag{1}$$

$$\theta_i \sim \text{uniform}(0, 1), \tag{2}$$

where $y_{ijk} = 1$ if the ith person on their jth problem chooses the kth option, and $y_{ijk} = 0$ otherwise. The probability θ_i is the constant probability the ith person uses for all problems and options.

[4] In so doing, we either do not take much heed of or, perhaps, fully embrace Pinker's (1997, p. 282) admonition: "In psychology, invoking 'strategies' to explain funny data is the last refuge of the clueless".

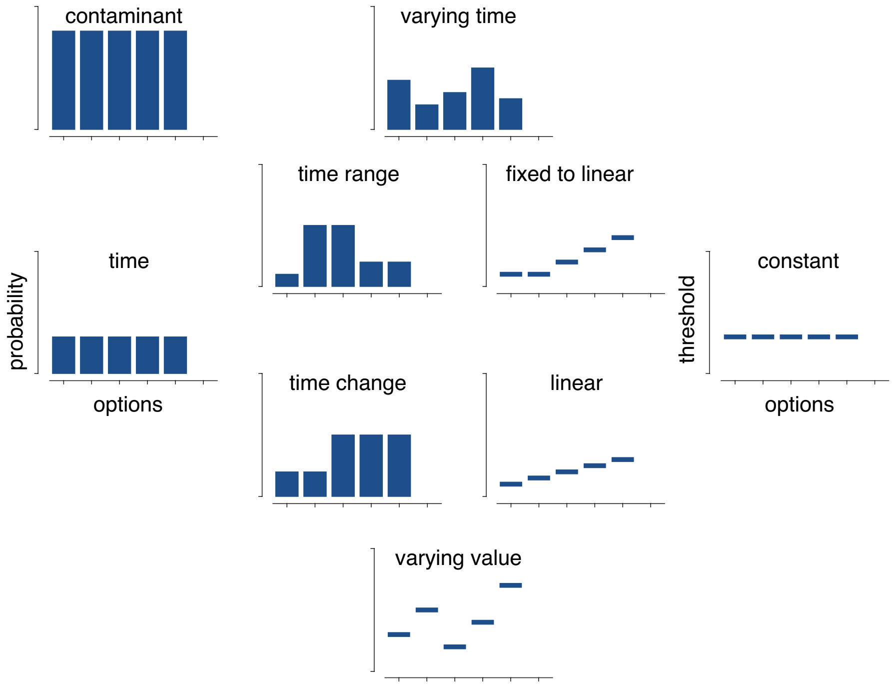

Fig. 5 Schematic representations of the nine cognitive strategies. Histograms represent the probability of purchasing a ticket at different time points for time-based strategies. Horizontal lines represent thresholds for purchasing tickets with lower values at different time points for value-based strategies

4.1.2 Contaminant strategy

The contaminant strategy is a special case of the time strategy, in which there is a high probability of choosing the ticket at any time period. This makes it likely that a ticket will be chosen among the first few options. This is consistent with unthinking contaminant behavior, in which a participant's goal is to complete the task as quickly as possible. Formally, the contaminant model changes Eq. 2 in the time strategy to:

$$\theta_i \sim \text{beta}\left(4, 1\right). \tag{3}$$

4.1.3 Time change strategy

The time change strategy assumes that there is a point at which the probability of purchasing a ticket increases. Formally, the strategy is defined by:

$$y_{ijk} \sim \text{Bernoulli}\left(\theta_{ik}\right) \tag{4}$$

$$\theta_{ik} = \begin{cases} \alpha_{i1} & \text{if } k < \lambda_i \\ \alpha_{i2} & \text{if } k \geq \lambda_i \end{cases} \tag{5}$$

$$\alpha_{i1}, \alpha_{i2} \sim \text{uniform}\big(0, 1\big) \,:\, \alpha_{i2} > \alpha_{i1} \tag{6}$$

$$\lambda_i \sim \text{categorical}\Big(0, \frac{1}{n-2}, \ldots, \frac{1}{n-2}, 0\Big). \tag{7}$$

Because the probability of buying at ticket now depends the time of the option, the observed choice behavior y_{ijk} now depends on a buying probability θ_{ik} that varies for the ith person and kth option, but is still assumed to be the same for all problems. The probability before the increase is α_{i1} for the ith person and α_{i2} after. The order constraint $\alpha_{i2} > \alpha_{i1}$ in Eq. 6 ensures that the probability increases. The increase happens at option λ_i. The categorical prior in Eq. 7 prevents the first or last option being the change point, since that would mean only a single purchase probability was used for all options, contrary to the meaning of the strategy.

4.1.4 Time range strategy

The time range strategy assumes that there is a bounded time period with increased probability of purchasing. Formally, the strategy is defined similarly to the time change strategy, by:

$$\theta_{ik} = \begin{cases} \alpha_{i1} & \text{if } k < \lambda_{i1} \\ \alpha_{i2} & \text{if } \lambda_{i1} \leq k \leq \lambda_{i2} \\ \alpha_{i3} & \text{if } k > \lambda_{i2} \end{cases} \tag{8}$$

$$\alpha_{i1}, \alpha_{i2}, \alpha_{i3} \sim \text{uniform}\big(0, 1\big) \,:\, \alpha_{i1} < \alpha_{i3} < \alpha_{i2} \tag{9}$$

$$\lambda_{i1}, \lambda_{i2} \sim \text{categorical}\Big(0, \frac{1}{n-2}, \ldots, \frac{1}{n-2}, 0\Big) \,:\, \lambda_{i1} < \lambda_{i2} \tag{10}$$

The time points at the beginning and end of the range are λ_{i1} and λ_{i2} for the ith person. The order constraints in Eq. 9 of $\alpha_{i1} < \alpha_{i3} < \alpha_{i2}$ ensures that that the probability of buying increases from α_{i1} to α_{i2} during this range, then decreases to α_{i3} when the range is finished. As before, the categorical priors in Eq. 10 prevent the first or last option being change points, and the additional order constraint $\lambda_{i1} < \lambda_{i2}$ means that λ_{i1} is the first change point that leads to the increase.

4.1.5 Varying time strategy

The most flexible time strategy allows each option to have a different probability of buying a ticket. Formally, the strategy is defined by:

$$\theta_{ik} \sim \text{uniform}\left(0, 1\right), \tag{11}$$

so that the probabilities are independent of each other.

4.2 Value strategies

4.2.1 Fixed strategy

The core psychological assumption of the fixed strategy is that a person has a threshold ticket value that they apply to every available ticket, and purchase a ticket if its price is below this threshold. Based on these thresholds alone, the strategy would make deterministic predictions. It would predict the ticket chosen was the first one in each problem below the threshold, or the final ticket if none of the previous ticket prices were below the threshold. To make the strategy probabilistic, we augment the core assumption to allow for variability in the application of the threshold, by introducing a sigmoid function that converts a difference between the current price and the threshold into a probability of purchasing.

The sigmoid mapping function can be conceived as arising from moment-to-moment fluctuations in the threshold price the person uses. Sigmoid functions are widely used in psychological modeling to make deterministic decision rules probabilistic (Kuss et al., 2005), and have been successfully used for modeling optimal stopping behavior (Baumann et al., 2023). Other modeling assumptions that make deterministic rules probabilistic have been used in previous modeling of optimal stopping behavior. In particular, trembling hand or error-of-execution processes, in which a deterministic choice is not followed with some small probability, have been used for optimal stopping problems with zero–one loss (e.g. Guan & Lee, 2018; Lee & Courey, 2021). The graded nature of the sigmoid assumption seems better suited to problems like the current one with expected loss.

Formally, we use a logistic choice rule to model the probability the ith person on the jth problem will purchase the kth ticket:

$$\theta_{ijk} = 1/\left(1 + \exp\left(-\beta_i\left(\tau_i - v_{jk}\right)\right)\right). \tag{12}$$

The probability of purchasing θ_{ijk} now depends on the problem because, unlike time-based strategies, value-based strategies are influenced by the ticket price that is available, and this changes for each problem. The fixed threshold used by the person has prior

$$\tau_i \sim \text{uniform}\left(300, 700\right). \tag{13}$$

The lower and upper bounds of \$300 and \$700 are based on the price distributions in Fig. 1, and define an informative prior (Lee & Vanpaemel, 2018). The assumption is that nobody uses a fixed threshold outside of this range for this particular environment.

The parameter β_i in Eq. 12 corresponds to how deterministically the ith person makes decisions. For β_i values near zero, behavior is near deterministic, with all ticket prices below the threshold very likely to be purchased, and all tickets with prices above the threshold very unlikely to be purchased. As β_i increases behavior becomes progressively more probabilistic. In the limit, as β_i becomes very large, θ_{ijk} is close to $\frac{1}{2}$ no matter what the price of the ticket is. An informative prior is also assumed for this parameter, in the form of a truncated positive Gaussian distribution:

$$\beta_i \sim \text{Gaussian}_+\left(0, 1/0.1^2\right). \tag{14}$$

This choice of prior is also tuned to the range of ticket values in the environment. It is designed to allow near-deterministic step functions and near-flat constant mapping functions, but place most of the prior probability on intermediate probabilistic possibilities. Lee (2018) provides a worked example of constructing an informative prior for logistic choice rules to meet these modeling goals, and the analysis on which Eq. 14 is based is available in the supplementary materials.

4.2.2 Linear strategy

The linear strategy generalizes the fixed strategy by allowing the threshold to increase at a constant rate as time progresses. Formally, the strategy is defined by replacing Eq. 13 with

$$\tau_{ik} = \alpha_i + k\gamma_i, \tag{15}$$

so that the threshold price for the first time point is α_i for the ith person, and the threshold increases by γ_i at each subsequent time point. The starting point and increase are given informative priors

$$\alpha_i \sim \text{uniform}\left(300, 700\right) \tag{16}$$

$$\gamma_i \sim \text{Gaussian}_+\left(0, 1/30^2\right). \tag{17}$$

4.2.3 Fixed-to-linear strategy

The fixed-then-linear strategy combines and generalizes the constant and linear strategies by assuming thresholds are constant up to some time point, and then increase linearly. Formally, the strategy defines thresholds as

$$\tau_{ik} = \begin{cases} \alpha_i & \text{if } k < \lambda_i \\ \alpha_i + \left(k - \tau_i\right)\gamma_i & \text{if } k \geq \lambda_i. \end{cases} \tag{18}$$

The change from a constant to an increasing threshold happens at time point λ_i, which is given the same prior used in the time change and time range strategies:

$$\lambda_i \sim \text{categorical}\left(0, \frac{1}{n-2}, \dots, \frac{1}{n-2}, 0\right). \tag{19}$$

4.2.4 Varying value strategy

The final strategy is the varying value strategy, which allows any threshold at any time point. Formally, thresholds are simply defined as

$$\tau_{ik} \sim \text{uniform}\big(300, 700\big). \tag{20}$$

This is clearly the most flexible strategy. By setting thresholds appropriate, the strategy can always describe a participant's behavior on any specific strategy. It is not, however, a saturated model, with the ability to describe any pattern of choices in the experiment, because the same thresholds are assumed to apply to every problem.

Having presented all nine strategies in terms of the intuitive time-based or value-based distinction, we note that this distinction is a conceptual approximation. All of the time strategies do not depend on the price of the ticket being considered. In this sense, they are pure time-based strategies. The value-based strategies, however, with the exception of the constant strategy, are really mixed strategies. The constant strategy uses one threshold for all points in time. The other three, however, use thresholds that change over time, and so really involve both time and value considerations. This is especially true for the final varying value strategy, which is extremely flexible. The unconstrained thresholds in this strategy could be conceived as being set based on considerations of both time and value.

5 Modeling results

We used a latent-mixture approach in JAGS (Plummer, 2003) to infer posterior probabilities that each participant uses each of the nine strategies. Given the nature of the strategies, with both continuous and discrete parameters, order constraints on some sets of parameters, and the use of informative priors, the heuristic of counting parameters to measure model complexity is clearly inadequate (Villarreal et al., 2023). The latent-mixture approach has the advantage of controlling for model flexibility and complexity according to prior predictive distributions, as in model comparison based on Bayes factors.

Figure 6 shows the results of the latent-mixture analysis, detailing the posterior probability for each participant using each strategy. There is clear evidence of significant individual differences in strategy use across the participants. All of the strategies except the time strategy were most likely used by at least one participant.

A majority of the participants are inferred to be likely to be using the value-based linear or fixed-then-linear strategies, with about twice as many using fixed-then-linear. A minority of participants are inferred to be using a time-based strategy, with about four or five participants most likely using each of the time change, time range, and varying time strategies.

For most participants, there is some uncertainty about which strategy they are using, but most of this uncertainty suggests the alternatives are closely-related strategies. For example, the participants most likely using the linear strategy could

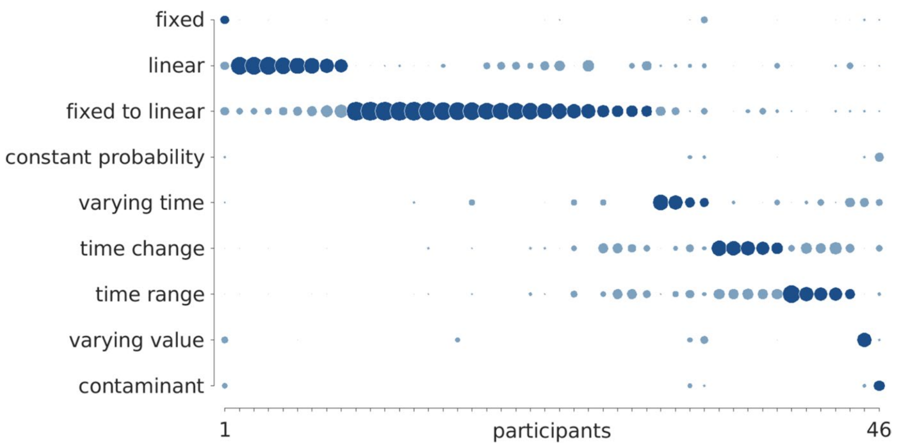

Fig. 6 Inferences about strategy use for all 46 participants. Circles show posterior probabilities for each strategy for each participant. The most likely strategy for each participant is highlighted in darker blue

instead be using the fixed-then-linear strategy. Only for six or seven participants is there significant uncertainty ranging across time-based and value-based strategies.

An advantage of including very simple and very flexible strategies in the set of those considered is that they provide “bookends” (Lee et al., 2019a). One participant is inferred to use the simplest contaminant strategy and one participant is inferred to use the most flexible varying value strategy. The remaining participants are inferred to use strategies with intermediate levels of complexity. This means inferences about their strategy use are not based on goodness-of-fit or complexity alone, but by striking an effective balance between these two considerations (Pitt et al., 2002).

5.1 An example of time-based strategy use

Figure 7 shows an example of a participant inferred to use a time-based strategy. The participant was inferred to use the time-change strategy, increasing their probability or purchasing a ticket from near 0% to about 50% two months before travel. These inferences match the observed behavior of the participant. They purchased a ticket only once before 2 months, but made regular purchases between 2 months and 2 weeks before travel.

Figure 7 quantifies the agreement between the modeling inferences and observed behavior in a number of ways. At the level of individual purchase or reject decisions, the posterior predictive agreement is 0.84. That is, the inferences made by the model about time-change strategy use and the purchase probabilities before and after the point of change, on average give a probability of 0.86 to the decision the participant made. Another way of measuring agreement is to focus on the purchased ticket for

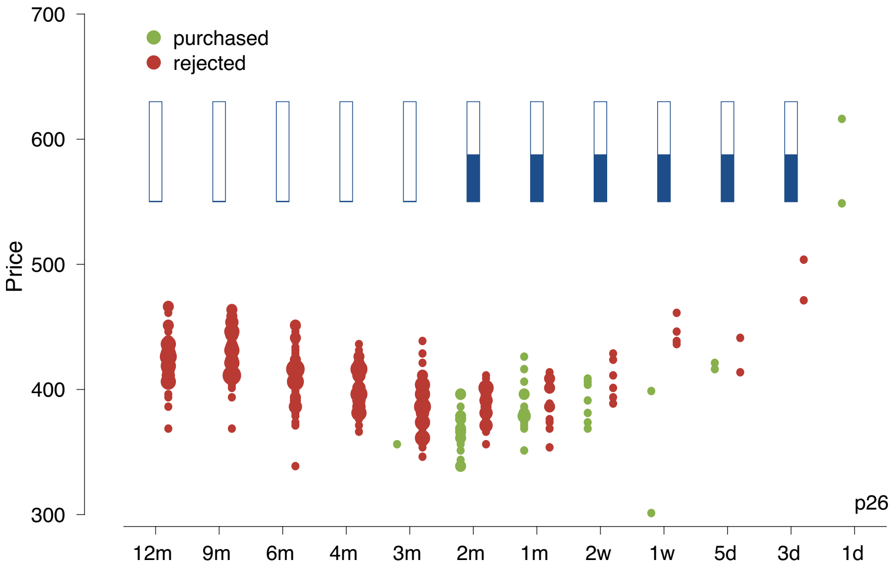

Fig. 7 An example of a participant inferred to use the time-change strategy. The green distributions show for each time point how often an airline ticket at a given price was purchased by the participant. The red distributions show ticket prices that were rejected. The blue bars show the inferred probability of purchasing a ticket at each time point up to 1 week before travel

each problem as a whole, rather than the individual decisions that lead to that choice. This requires the model to capture exactly the sequence of reject decisions a participant made, finishing with the final purchase decision. Under this stricter measure, the posterior predictions of the model on average give a 0.28 probability to the ticket finally purchased by the participant.

5.2 An example of value-based strategy use

Figure 8 shows an example of a participant inferred to use a value-based strategy. The participant was inferred to use the linear strategy, increasing the price they were willing to pay from a little over \$300 to about \$450. These inferences generally match the observed behavior of the participant. The thresholds at 6 months and 1 week make perfect divisions, but only for one purchased ticket. Similarly, the thresholds at 12 months, 9 months, 5 days, 2 days, and 3 days also work perfectly because they are smaller than all of the rejected prices. More substantially, at the time points 1 month and 2 months before travel, the thresholds mostly divide purchased prices from rejected ones. There is more overlap 3 months before travel that prevents the threshold dividing the decisions cleanly. This account of the participant's decision making has a posterior predictive agreement probability of 0.87 for individual decisions and 0.43 for the sequence of decisions terminating in a purchase.

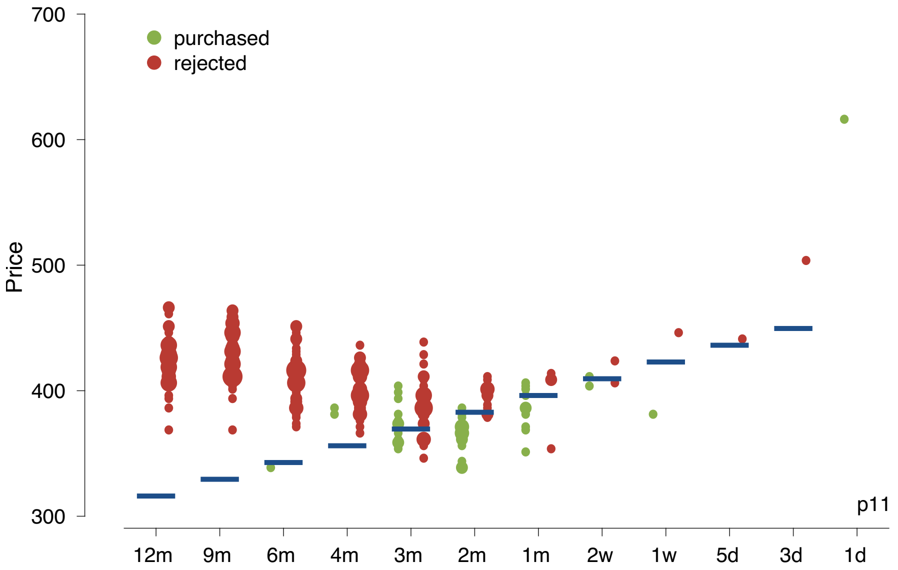

Fig. 8 An example of a participant inferred to use the linear strategy. The green distributions show for each time point how often an airline ticket at a given price was purchased by the participant. The red distributions show ticket prices that were rejected. The blue bars show the inferred price thresholds at each time point

5.3 Limitations of the strategies

The results for the specific participants presented in Figs. 7 and 8 are representative of the results for the majority of participants.[5] Across all participants, the expected posterior predictive agreement for an individual decision, based on the most likely strategy for each participant, is 0.78. This level of agreement ranges over participants from 0.50 to 0.94. The expected agreement for the sequence of decisions leading to a final purchasing decision is 0.28 on average, ranging from 0.09 to 0.93. A reasonable summary of these overall results is that the strategies we considered typically explain around half of participants' fine-grained decision making and a quarter of their observed final choices.

Failures of models in terms of goodness-of-fit are most important when they arise because the proposed cognitive model cannot capture empirical regularities in people's behavior. There is at least one such regularity in our data. It involves the early purchase of relatively expensive tickets, but in a pattern that is not simple contaminant behavior. An example of a participant who behaves this way is shown in Fig. 9. They often purchase tickets 6 months and 4 months before travel that are not cheap relative to the distribution of prices at those time points. They do not, however, show

[5] The term "representative" means here what it often means in experimental psychology: at the more impressive end of what is observed, but not cherry picked nor hiding gross failures. Modeling results for all participants are available in the supplementary material.

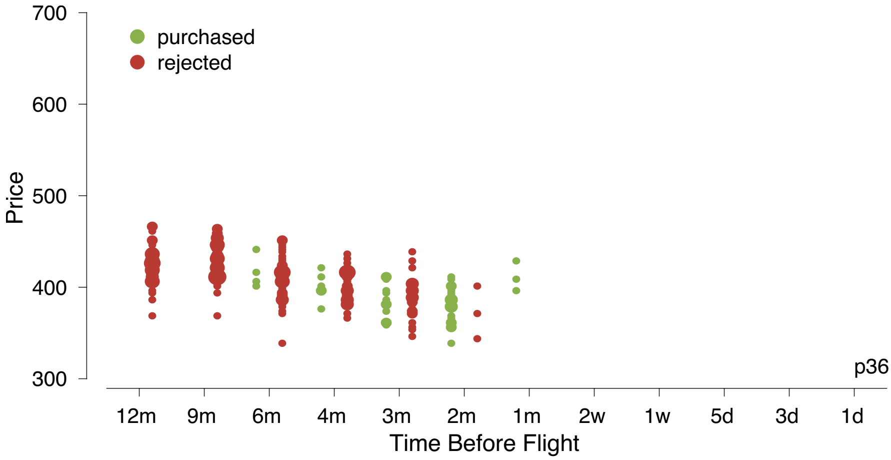

Fig. 9 An example of a participant who behaves in a way that is difficult to model with the current strategies. The green distributions show for each time point how often an airline ticket at a given price was purchased by the participant. The red distributions show ticket prices that were rejected

the same behavior 12 months or 9 months before travel. None of our nine time-based or value-based strategies can provide a natural account of this behavior.

6 Discussion

Our results suggest that there are significant individual differences in the strategies people use to solve the optimal stopping task of purchasing a plane ticket. Some participants make purchases consistent with time-based strategies like "I will buy a ticket between three months and one month before travel." Others make purchases consistent with value-based strategies like "I will buy a ticket as soon as the price goes under \$300." But our results also suggest these sorts of strategies do not provide a complete explanation. The same participant, for different problems, sometimes purchases but sometimes rejects essentially the same price offered at the same time. There is no way to explain this behavior without extending the cognitive account beyond the bounds of just time-based or value-based strategies.

One possibility is that the inconsistency is caused by learning, adapting, or shifting strategies (Lee et al., 2019; Lee & Courey, 2021). The analysis in Fig. 3, however, makes this unlikely. It shows that there is no evidence of significant learning at the level of individual participants or overall. This does not prove that no form of adaptation is responsible, but it would have to change behavior significantly without affecting performance, which seems unlikely. Of course, our experimental design provided the relevant environmental distributions directly to participants, which is a best-case scenario. A more complete cognitive model would seek to account for how people learn and maintain approximations to the price distributions. But, for our task, the analysis in Fig. 3 suggests participants did not systematically change strategies over the course of the experiment.

Another possibility is that there are sequential effects. These were studied empirically by Shapira and Venezia (1981) and considered later in a model-based analysis by Guan and Lee (2018). It seems possible that people could become especially excited by a large drop in price and buy the current ticket on that basis. It also seems possible that people regret missing a good price, and this prompts them to buy the next ticket. Incorporating one or both of these two effects, however, would add significant flexibility to any decision strategy. If both were implemented in an unconstrained way, they seem like to be able to describe almost any pattern of observed behavior. Accordingly, it would be important to make theoretical assumptions to provide constraints on sequential effects in order to explore how the context of previous prices within a problem, or purchasing decisions for previous problems, impacts decision making.

To provide a better account of people's decision making, it is also likely to be worth expanding the strategies considered. For example, the participant in Fig. 9, besides sometimes choosing expensive tickets early, seems to have a hard deadline of one month for making a purchase. As discussed earlier, the varying value strategy can in principle capture this behavior. Setting the 1 month threshold to $0 would implement the deadline. But the varying value strategy makes this possible only through extreme flexibility. For novel strategies to be inferred from limited behavioral data likely requires them to be formalized more specifically, once again based on theoretical assumptions. A possible psychological basis for this behavior is that people's utilities depend on time as well as price. Our task asked people to minimize the cost of tickets. But there is a psychological cost to the uncertainty of delaying purchasing in search of cheaper prices. It is very possible that some people value the certainty that comes with early and predictable purchasing well before the day of travel. This idea could be explored in a modified task that incorporated both time and price considerations in the criterion to be optimized.

Deciding when to buy an airline ticket over a prolonged period leading up to travel is an interesting example of an optimal stopping problem. It is one people understand intuitively, which makes it a good cover story for the experimental study of optimal stopping in changing environments. It is also a real-world problem which could benefit from understanding how people decide when to purchase tickets. The pricing of tickets is based in part on consumer demand, which arises collectively from individual purchasing decisions. This study is a first attempt to model the time-based and value-based strategies that people may use. It provides some insight into people's decision making, especially by highlighting large potential individual differences in strategy use. It also makes clear that people's behavior is complicated, and the use of simple strategies is only part of the full story.

Funding This research was supported by UROP and SURP funding from the University of California Irvine to Sara Chong.

Data availability The replication and supplementary material for the study is available at http://doi.org/10.17605/OSF.IO/NVE5B.

Declarations

Conflict of interest The authors have no competing financial or non-financial interests to declare.

References

Baumann, C., Gershman, S. J., Singmann, H., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Science, 117*(23), 12750–12755. https://doi.org/10.31234/osf.io/epd2x

Baumann, C., Schlegelmilch, R., & von Helversen, B. (2023). Adaptive behavior in optimal sequential search. *Journal of Experimental Psychology: General, 152*(3), 657–672. https://doi.org/10.1037/xge0001287

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. https://doi.org/10.1287/mnsc.1060.0535

Bilotkach, V., Gaggero, A. A., & Piga, C. A. (2015). Airline pricing under different market conditions: Evidence from European low-cost carriers. *Tourism Management, 47*, 152–163. https://doi.org/10.1016/j.tourman.2014.09.015

Bonner, S. E., Hastie, R., Sprinkle, G. B., & Young, S. M. (2000). A review of the effects of financial incentives on performance in laboratory tasks: Implications for management accounting. *Journal of management accounting research, 12*(1), 19–64. https://doi.org/10.2308/jmar.2000.12.1.19

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty, 19*, 7–42. https://doi.org/10.1007/978-94-017-1406-8_2

Campbell, J., & Lee, M. D. (2006). The effect of feedback and financial reward on human performance solving 'secretary' problems. In R. Sun (Ed.), *Proceedings of the 28th Annual conference of the cognitive science society* (pp. 1068–1073). Erlbaum.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science, 4*(3), 282–296. https://doi.org/10.1214/ss/1177012493

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal, 61*(313), 35–73. https://doi.org/10.1007/978-0-387-44956-2_22

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science, 66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Guan, H., Lee, M. D., & Silva, A. (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). Cognitive Science Society.

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision, 5*(4), 339–361. https://doi.org/10.1037/dec0000081

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences, 24*(3), 383–403. https://doi.org/10.1017/s0140525x01004149

Kuss, M., Jäkel, F., & Wichmann, F. A. (2005). Bayesian inference for psychometric functions. *Journal of Vision, 5*(5), 478–492. https://doi.org/10.1167/5.5.8

Lee, M. D. (2006). A hierarchical Bayesian model of human decision making on an optimal stopping problem. *Cognitive Science, 30*(3), 555–580. https://doi.org/10.1207/s15516709cog0000_69

Lee, M. D. (2018). Bayesian methods in cognitive modeling. In J. Wixted & E.-J. Wagenmakers (Eds.), *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience. Volume 5: Methodology* chapter 2 (4th edn, pp. 37–84). Wiley.

Lee, M. D., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior, 4*, 1–17. https://doi.org/10.1007/s42113-020-00085-9

Lee, M. D., Criss, A. H., Devezer, B., Donkin, C., Etz, A., Leite, F. P., Matzke, D., Rouder, J. N., Trueblood, J. S., White, C. N., & Vandekerckhove, J. (2019a). Robust modeling in cognitive science. *Computational Brain & Behavior, 2*, 141–153. https://doi.org/10.31234/osf.io/dmfhk

Lee, M. D., Gluck, K. A., & Walsh, M. M. (2019b). Understanding the complexity of simple decisions: Modeling multiple behaviors and switching strategies. *Decision, 6*(4), 335–368. https://doi.org/10.1037/dec0000105

Lee, M. D., & Vanpaemel, W. (2018). Determining informative priors for cognitive models. *Psychonomic Bulletin & Review, 25*, 114–127. https://doi.org/10.3758/s13423-017-1238-3

McCartney, S. (2014). The best day to buy airline tickets. *The Wall Street Journal*. https://www.wsj.com/articles/the-best-day-to-buy-airline-tickets-1413999377

Pinker, S. (1997). *How the mind works*. W. W. Norton & Co.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109*(3), 472–491. https://doi.org/10.1037/0033-295x.109.3.472

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), Proceedings of the 3rd international workshop on Distributed Statistical Computing (DSC 2003), vol. 124 (pp. 123 Computational Brain & Behavior 1–10). Vienna, Austria.

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. https://doi.org/10.1006/obhd.1997.2683

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks. *Journal of Behavioral Decision Making, 13*(4), 391–411. https://doi.org/10.1002/1099-0771

Shapira, Z., & Venezia, I. (1981). Optional stopping on nonstationary series. *Organizational Behavior and Human Performance, 27*(1), 32–49. https://doi.org/10.1016/0030-5073(81)90037-4

Villarreal, M., Etz, A., & Lee, M. D. (2023). Evaluating the complexity and falsifiability of psychological models. *Psychological Review, 130*(4), 853–872. https://doi.org/10.1037/rev0000421