

Baumann, Christiane ; Schlegelmilch, René ; von Helversen, Bettina

Adaptive Behavior in Optimal Sequential Search

Journal Article as: peer-reviewed accepted version (Postprint)

DOI of this document* (secondary publication): <https://doi.org/10.26092/elib/3162>

Publication date of this document: 01/08/2024

* for better findability or for reliable citation

Recommended Citation (primary publication/Version of Record) incl. DOI:

Baumann, Christiane ; Schlegelmilch, René ; von Helversen, Bettina. 2023. Adaptive behavior in optimal sequential search. In: Journal of Experimental Psychology, Vol. 152 (Issue: 3), pp. 657-692. © American Psychological Association. DOI: 10.1037/xge0001287.

Please note that the version of this document may differ from the final published version (Version of Record/primary publication) in terms of copy-editing, pagination, publication date and DOI. Please cite the version that you actually used. Before citing, you are also advised to check the publisher's website for any subsequent corrections or retractions (see also <https://retractionwatch.com/>).

©American Psychological Association, [2023]. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. The final article is available, upon publication, at: <https://doi.org/10.1037/xge0001287>

This document is made available with all rights reserved.

Take down policy

If you believe that this document or any material on this site infringes copyright, please contact publizieren@suub.uni-bremen.de with full details and we will remove access to the material.

Adaptive behavior in optimal sequential search

Christiane Baumann^{1,2}, René Schlegelmilch³, and Bettina von Helversen³

¹University of Zurich

²Harvard University

³University of Bremen

© 2022, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/xge0001287

Author Note

Christiane Baumann was supported by Swiss National Science Foundation Grant P0ZHP1_168889 and Bettina von Helversen was supported by the Swiss National Science Foundation Grant 157432. Correspondence should be addressed to Christiane Baumann at cbaumann@hks.harvard.edu.

Data and code are available on the Open Science Framework:
<https://osf.io/9rct5>(Baumann et al., 2022, April 17).

We would like to thank Vassilios Kaxiras for helping with data collection.

Abstract

Sequential decision making—making a decision where available options are encountered successively—is a hallmark of everyday life. Such decisions require deciding to accept or reject an alternative without knowing potential future options. Prior work focused on understanding choice behavior by developing decision models that capture human choices in such tasks. We investigated people’s adaptive behavior in changing environments in light of their cognitive strategies. We present two studies in which we modified (1) outcome variance and (2) the time horizon and provide empirical evidence that people adapt to both context manipulations. Furthermore, we apply a recently developed threshold model of optimal stopping to our data to disentangle different cognitive processes involved in optimal stopping behavior. The results from Study 1 show that participants adaptively scaled the values of the sampling distribution to its variance, suggesting that the value of an option is perceived in relative rather than absolute terms. The results from Study 2 suggest that increasing the time horizon decreases the initial acceptance level, but less strongly than would be optimal. Furthermore, for longer sequences, participants more weakly adjusted this acceptance threshold over time than for shorter sequences. Further correlations between individual estimates in each condition indicate that individual differences between the participants’ thresholds remain fairly stable between the conditions, pointing toward an additive effect of our manipulations.

To accept a safe option or to continue searching for a better alternative is a hallmark of everyday life, whether house hunting, selling or buying stocks, or finding a partner. In such so-called optimal stopping problems, options are presented sequentially and people must decide whether to accept an offer or to explore more alternatives (see Ferguson, 1989, for an early review). An important characteristic of such tasks is that there is no going back to an earlier option after it is initially declined. Therefore, the difficulty lies in the trade-off between accepting a possibly suboptimal option prematurely or rejecting it, hoping for a better one in the future.

Previous research has emphasized either studying choice behavior in comparison to normative models (see, e.g., Brickman, 1972; Corbin et al., 1975; Kahan et al., 1967; Rapoport & Tversky, 1970) or introducing alternative choice models to characterize optimal stopping behavior (Baumann et al., 2020; Bearden et al., 2006; Cox & Oaxaca, 1989; Goldstein et al., 2020; Guan et al., 2015; Hey, 1982; Lee, 2006; Lee et al., 2004; Sang et al., 2020; Seale & Rapoport, 2000; Zwick et al., 2003). However, an often neglected but important challenge when solving optimal stopping problems is to adapt one's behavior to changes in the environment, such as dealing with outcomes of different magnitude or limited time horizons for searching. As an example, assume it is Black Friday (sales for 3 days) and Barbara is bargain hunting for a specific video game. They explore offers from different discounters that range between \$10 and \$20, but the game sells fast and offers disappear quickly. At the same time, a luxury mall extends the Black Friday days to 2 weeks, and Barbara intends to purchase an expensive bag that is now offered for a price ranging between \$700 and \$900. However, price offers are changing from day to day and Barbara has to choose the right time to purchase the bag. How does Barbara adjust their choice strategy to different time and price scales? Here we consider an optimal stopping setting that closely reflects the above examples, where the decision maker is informed of the value of each draw and must decide after each draw whether to choose or reject it. Moreover, the sampling distribution is known and the payoff corresponds to the chosen value. In this variant, the optimal strategy dictates comparing an alternative to a position-dependent threshold that corresponds to

the expected reward of the remaining options (see Section 5b in Gilbert & Mosteller, 1966, and Appendix Text 1 & 2). If the current alternative exceeds the threshold, it should be accepted. Because of the number of remaining options in continuing sequences, the optimal decision threshold is nonlinearly increasing (becoming more relaxed with the decreasing number of remaining options). Thus, regarding variance and time horizon, two predictions immediately follow. First, since the optimal model involves a relative comparison it predicts that choice behavior should be equal regardless of whether the value variance is high or low. Second, when the total number of options changes in different time horizons the optimal model predicts that people should become pickier in longer time horizons, but also that the acceptance probability at the last possible choice (i.e., when only one future option remains) should be 50%.

However, it appears to be widely acknowledged that humans do not behave optimally in optimal stopping tasks (e.g., Baumann et al., 2020; Bhatia et al., 2021; Lee & Courey, 2021; Lee et al., 2004; Zwick et al., 2003), such that typical accounts often focus on describing deviations from optimality due to biases and noise (Baumann et al., 2020; Bhatia et al., 2021; Goldstein et al., 2020; Guan et al., 2015). Recent debates, however, have criticized the idea of capturing parametric deviations from optimal (bias) for neglecting the underlying psychological processes that drive behavior in the first place (for a general discussion, see Rahnev & Denison, 2018). In particular, a biased optimal model (BOM) suggests that people conduct the effortful calculation of optimal probabilities and then deviate from them behaviorally. Such a calculation implies that (a) people have an exact understanding of the underlying sampling distribution, and based on this distribution (b) use backward induction to calculate the outcome probabilities in each time step. However, individuals have, in general, difficulties with probabilistic concepts (e.g., Alberini et al., 2004; Goldstein & Rothschild, 2014) and backward induction has been shown to be extremely challenging to compute for just a few time steps (Huys et al., 2015). In turn, decision models that involve such computational steps assume processes that appear to be intractable and implausible for a human decision maker (e.g., Bossaerts & Murawski, 2017; Cooper, 1990; Gershman

et al., 2015; Gigerenzer & Gaissmaier, 2011; Simon, 1955; Todd et al., 2012). Moreover, several studies investigating human search in related search tasks have shown that corresponding models were not able to accommodate behavioral characteristics, leading to inferior model fits compared to heuristic choice rules (e.g., Baumann et al., 2020; Hey, 1982; Kahan et al., 1967; Rydzewska et al., 2018; Sang et al., 2020; Song et al., 2019). In turn, these researchers argued that humans adopt heuristic decision rules rather than calculating optimal outcome probabilities. Heuristic decision rules, such as simple threshold rules, emphasize the plausibility of the assumptions on the underlying decision processes. Especially in the domain of sequential search paradigms, which entail complex optimal solutions, such rules appear to be reasonably good and fairly robust and thus may constitute a better explanation (and predictor) of actual search behavior.

In particular, prior research has shown that the BOM's assumption of nonlinearly increasing thresholds is in disagreement with people's decision thresholds. A model proposing linearly increasing thresholds captures choices more accurately across changing environments (Baumann et al., 2020; Lee & Courey, 2021) and also in related sequential search domains (Sang et al., 2020; Song et al., 2019). Importantly, when studying predictions of optimal and linear threshold models (LTMs) in regard to adaptations to changing environments (e.g., left- vs. right-skewed value distributions), behavioral evidence has been shown to contradict the predictions of the optimal model but to correspond to the predictions of the LTM. Specifically, the LTM correctly predicted that people tend to search less in scarce environments (simulated using a left-skewed sampling distribution) than in environments with many good options (simulated using a right-skewed sampling; see Baumann et al., 2020), which is the opposite of the prediction of the optimal model and thus is also not predicted by BOMs.

In contrast to a model that defines human choices as deviations from optimality, the LTM (Baumann et al., 2020) allows one to directly measure the impact of task features on the cognitive parameter level. These parameters provide insights into the underlying psychological process when a person is deciding among sequentially presented options. In particular, it assumes that humans compare an observed offer to a decision threshold

at each position, and this threshold is linearly adjusted across positions. The decision maker's choice is thus determined by an aspiration level before search (initial threshold), a linear increment thereof when offers are rejected (adaptation rate), and choice sensitivity (or determinism). The initial threshold reflects the option value above which an option is accepted directly in the first decision. The adaptation rate reflects the incremental rate of change of the aspiration level with time or number of available options. Choice sensitivity reflects how sensitively people react to the extent to which the current option value deviates from the (adapted) aspiration level. Therefore, the LTM's high accuracy in capturing sequential choices across different contexts and its parameters' psychological interpretability make the LTM an appropriate framework for measuring the impact of changing task features on these mechanisms.

People's sensitivity to the task environment in an optimal stopping task has been a consistent finding across studies (e.g., Corbin et al., 1975; Cox & Oaxaca, 1989; Guan & Lee, 2018; Lee & Courey, 2021; Lee et al., 2004; Shapira & Venezia, 1981). For example, participants stopped earlier in sequences with descending compared to ascending values (Brickman, 1972; Shapira & Venezia, 1981) and increased their aspiration levels in rich compared to scarce environments (Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2014; Kahan et al., 1967; Rydzewska et al., 2018). Moreover, people adjusted to the length of the sequence by increasing their acceptance rate when fewer alternatives were available (e.g., Guan et al., 2020; Lee et al., 2004; Seale & Rapoport, 1997). However, the underlying processes that drive this adaptive behavior are unknown. From an LTM perspective, thus, the question is which cognitive parameters (e.g., initial threshold or adaptation rate) capture such behavioral trends.

Our first goal was to clarify this question in two studies when manipulating value variance (high vs. low; Study 1) and time horizon (number of options; Study 2).

Importantly, we next sought to increase the understanding of suboptimal behavior from a cognitive perspective. For this, we investigated how well people adjust within the limits of their cognitive strategy. This approach allowed us to then analyze maladaptation of cognitive variables (initial threshold and adjustment thereof) outside

the constraints of linearity, meaning they could be attributed to other psychological variables such as risk aversion in higher variance environments in Study 1 or increased search costs in longer sequences in Study 2.

Our second goal was to investigate the stability and reliability of the measured cognitive parameters across contexts and to examine their relationship with behavioral measures. This is a critical question, as attempts to link suboptimal search behavior with psychological variables, such as risk preferences, have been unsuccessful (e.g., Frey et al., 2017; Schunk, 2009). Uncovering underlying cognitive processes that are consistent across search contexts would contribute to a comprehensive understanding of individual differences in optimal stopping tasks.

In the following section, we introduce the environment manipulations and discuss how different theories may predict how people adapt. We then proceed to Study 1 and Study 2, for each, presenting first our analysis of how optimal decision makers would adjust, followed by the predictions of the LTM. In a subsequent modeling section, we use the LTM to describe the participants' behavior and compare the parametric characteristics in each condition.

Changes in decision environments

Variance. The introductory examples about purchasing a video game and a luxury bag illustrate that prices may vary for a desired object. How do humans translate different variances into representations that map onto decision criteria?

Research on simultaneous decision making suggests that a person's valuation of an option is sensitive to the range of other alternatives in the sample (e.g., Nieuwenhuis et al., 2005; Padoa-Schioppa, 2009; Rigoli et al., 2016; Tversky & Simonson, 1993), such that people value an alternative on the basis of its relative rank within a pool of alternatives and not by its absolute value (e.g., the normalization hypothesis; Rangel & Clithero, 2012). Following this line of research, we would expect that variance has little effect on search performance (i.e., on accepting optimal options) in an optimal stopping task, but that to perform this task, people transform their value representations into

value rankings, taking the variance of the alternatives into account. Interestingly, this adaptive mechanism corresponds to the prediction of the mathematically derived optimal solution, which scales the critical percentile values to the specific environment distribution. However, whether people perfectly scale these estimates (true variance) in an optimal stopping problem is an open question, which can be tested using the LTM, for instance, by estimating the scaling factor of decision thresholds between conditions. Alternatively, research on risky decision making suggests that variance in the outcomes changes people's risk preferences, where higher outcome variance implies higher perceived risk and thus results in more risk-averse behavior (Genest et al., 2016; Holt & Laury, 2002; Markowitz, 1959; Weber et al., 2004). From this perspective, an optimal stopping task can be viewed as a series of repeated risky decisions, where accepting the current option would reflect a "safe" choice and rejecting the current option would reflect choosing an uncertain outcome ("gamble"). If variance influences risk preference, we would expect that higher variance would lead to enhanced risk-averse behavior, translating into more relaxed aspiration levels, compared to lower variance. In this case, variance would affect search length, with higher variance leading to earlier stopping.

Time Horizon. Another influential factor in optimal stopping tasks is the time horizon (e.g., Lee et al., 2004; Rydzewska et al., 2018; Seale & Rapoport, 1997; Zwick et al., 2003). For example, when time is limited, as in the example above where sales appear for 3 days, rejecting an option implies that only a few more options are left over in the remaining time. The optimal solution predicts people will be pickier in longer time horizons at the beginning of the sequence, which leads to higher performance. Although increased search length in longer sequences has been empirically confirmed, it is less clear what cognitive processes drive this behavior. Which factors in the LTM need to be adapted in longer time horizons to maintain a good performance? In a simulation study (see Study 2, Adaptation to the time horizon: Optimal and linear models) we showed that to achieve the best performance, both the initial threshold and its adjustment across search should be decreased. However, the constraints of linear thresholds lower performance relative to optimality in extended time horizons.

Therefore, in Study 2 we sought to replicate earlier findings and extend them by using within-subject manipulations of time horizon, studying how different time horizons influence participants' choice behavior by comparing the cognitive parameters underlying adaptive search.

Study 1

In the first study we investigated how outcome variance affects search behavior in the optimal stopping task. In an online purchase task, we asked participants to repeatedly perform a sequential choice task in which they had to choose the cheapest airplane ticket. We manipulated the variance within participants by generating ticket prices from two Gaussian distributions with high and low variance [but constant mean; i.e., $\sim N(\mu = 180, \sigma = 10)$ or $\sim N(\mu = 180, \sigma = 40)$] and participants solved these tasks in a within-subject design.

Method

Participants. We recruited 205 participants (78 women; age range: 19–70 years) on Amazon Mechanical Turk to participate in the experiment. Participants gave informed consent, and the study design and method were approved by the ethics committee of the University of Zurich. Participants were excluded from analysis if they accepted the first option in a sequence in more than 90% of the trials, leaving 199 participants for the subsequent analysis. Participants received a fixed payment of \$4 and a bonus contingent on their search performance (see the Procedure section for the calculation of the performance-dependent bonus).

Procedure. The task was adapted from the study of Baumann et al. (2020) and imitates an airline-ticket-shopping scenario where people search for the cheapest price. In each trial, participants were presented with a sequence of 10 prices. For each price, they had to decide to accept it or reject it at their own pace. Participants were aware of the total number of prices in each trial and of the actual position in the sequence (Figure 1A). It was not possible to go back to an earlier option. If they reached the last (10th) price they were forced to choose it. When participants accepted a price, they received

feedback about how much they could have saved if they had chosen the best price in the sequence (Figure 1B). A bonus was paid proportional to the performance: The closer they were to the best price in the sequence, the more points they got (see Equation 1). All participants completed 80 sequences in each of two conditions, one with high variance and one with low variance. Within each condition, we sampled the ticket prices from either $\sim N(\mu = 180, \sigma = 10)$ or $\sim N(\mu = 180, \sigma = 40)$. The order of the two conditions was randomized and each participant encountered a newly generated sample. Prior to each of the two conditions, participants learned about the respective distribution to exclude learning effects in the stopping task. During learning, participants saw 50 ticket prices drawn from the target distribution, presented sequentially in randomly ordered trials (Figure 1C). After every 10th trial, participants had to estimate the overall average, and the correct answer was revealed (Figure 1D). Finally, participants created a histogram (by dragging the bars) of the true distribution (Figure 1E), before the correct distribution was superimposed over their estimate (Figure 1F; method proposed in Goldstein & Rothschild, 2014). If their estimates deviated more than 20% from the correct answers, they were instructed to repeat the learning phase, with a maximum of three repetitions. This procedure ensured that participants had good knowledge about the value distribution, thus minimizing any additional learning during the task.

Before the task started, participants were told their payoff would consist of a lump sum of \$4 plus a bonus between \$0 and \$4 contingent on their performance. Each trial offered a maximum of 25 points to earn, which translated into a bonus of 0.025 cents per trial. The maximum number of points was awarded when the cheapest ticket was chosen and 0 points for the most expensive one. Intermediate choices equated to

$$points_j = \frac{25 \cdot (p_j^{max} - p_j^{chosen})}{p_j^{max} - p_j^{min}}, \quad (1)$$

where p_j^{max} is the highest and p_j^{min} the lowest price in trial j . The final bonus was calculated by dividing the total number of points by 1,000. The average bonus earned

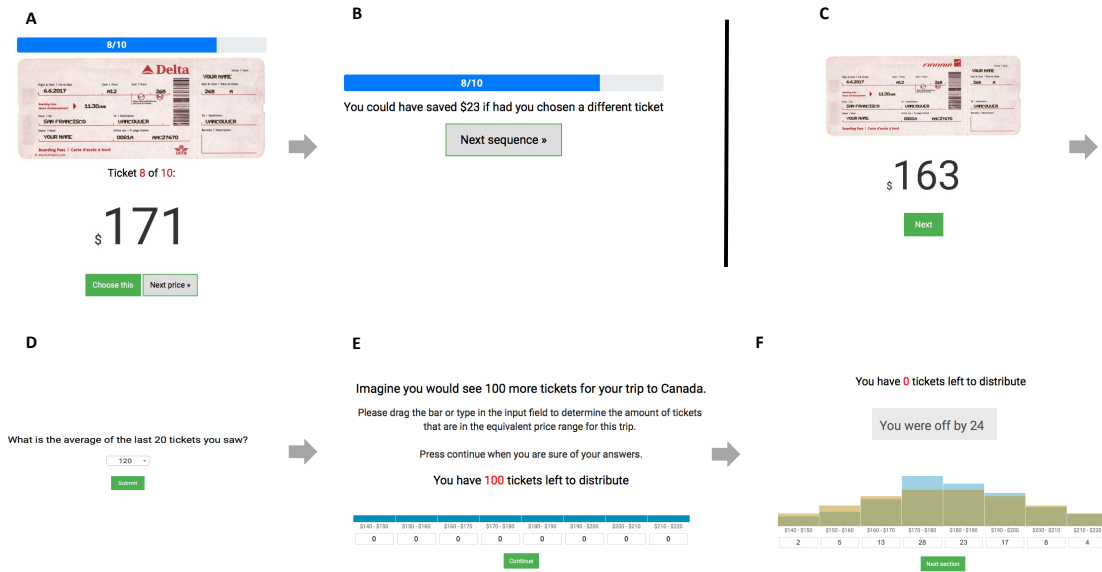


Figure 1. Screenshots of online experiment. Panels A,B: Testing Phase. Panels C–F: Learning phase. Panel A: In every trial participants encountered 10 tickets sequentially and had to decide to accept it or to continue. Panel B: Feedback was provided. Panel C: Feedback was provided by superimposing the correct distribution over their estimate. Panel D: Sequential presentation of ticket values. Panel E: After each 10 tickets, participants estimated the average of the prices just seen. Panel F: At the end of the learning phase, participants predicted how a future sample from the same predefined population might look by drawing a histogram.

was \$3.38 (min: \$2.83; max: \$3.64).

Adaptation to variance: The optimal model versus the LTM. Before we begin to analyze the behavioral data gathered from this experiment, we first discuss the adaptation to variance proposed by the optimal model and the LTM. The optimal solution to this task is to use a position-dependent optimal threshold to which an option is compared. The calculation of the optimal threshold is based on the expected reward of the prices on the remaining days when following the optimal policy. Gilbert and Mosteller (1966) derived the critical values when options are sampled from a standard normal distribution (see Gilbert & Mosteller, 1966, for a mathematical derivation of thresholds, and Appendix, Text 1). These values are then converted to the scale of the respective distribution (by multiplying the value by the distribution’s variance and adding the mean of the distribution; see Figure 2A (solid black line) for thresholds for values from $\sim \mathcal{N}(\mu = 0, \sigma = 1)$ and Figure 2B (solid black lines) for thresholds from the normal distribution with a mean of 180 and a variance of 10 and

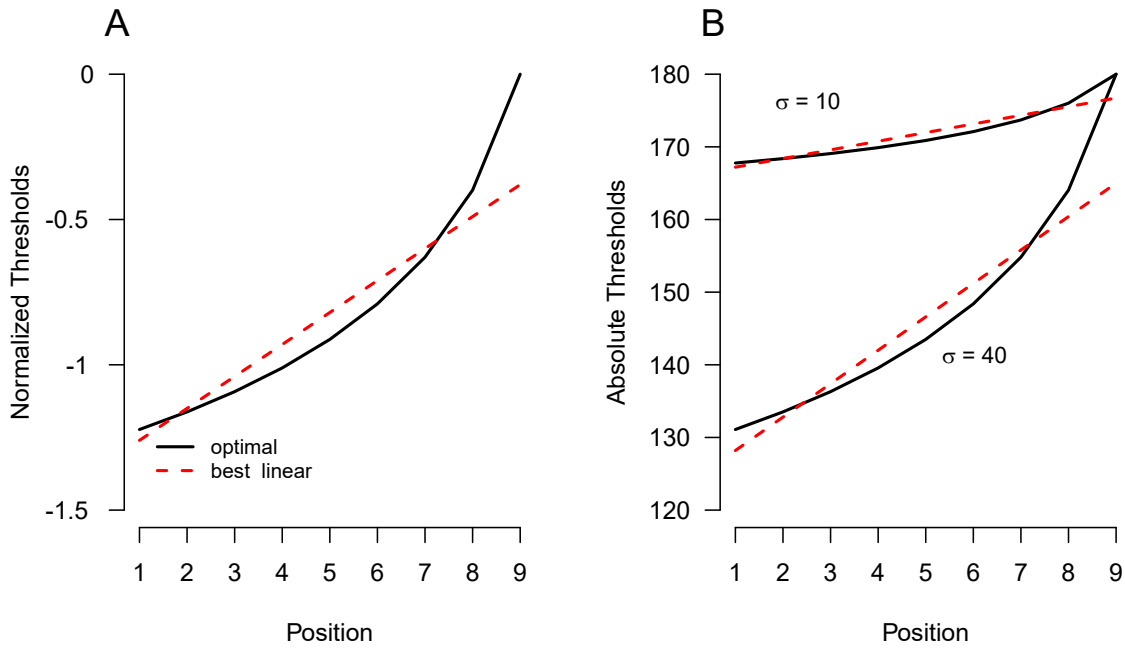


Figure 2. Optimal and best performing linear thresholds on a normalized (Panel A) and an absolute (Panel B) scale.

40). Consequently, thresholds differ on the absolute scale, with higher values in the low-variance condition (Figure 2B). However, they remain identical regarding their percentile rank within the sampling distribution. Thresholds corresponding to identical percentiles lead to the same stopping points and performances between environments that differ only in mean or variance. A simulation study¹ across the two conditions, confirmed that under the optimal model, a change in variance leads to the same performance and search length (see Figure 4, black solid lines).

To investigate the prediction of the linear strategy, we calculated those thresholds in both variance conditions (using grid search) that would lead to collecting the highest payoffs. The best performing linear thresholds in both variance conditions are shown in Figure 2B (red dashed lines), indicating an adjustment to variance, with higher threshold values in low-variance decision environments. Importantly, once these thresholds are scaled to the corresponding normalized values, they coincide (Figure 2A,

¹ For each condition ($SD = 10$ and $SD = 40$), an equivalent number of 80 trials were simulated, where values were sampled from $\sim \mathcal{N}(\mu = 180, \sigma = 10)$ and $\sim \mathcal{N}(\mu = 180, \sigma = 40)$ for a total of 240 optimal agents, which were aggregated on measures of performance and search length.

red dashed line). This result also indicates that the best performing linear decision thresholds across different variance environments match in respect to their percentile value, demonstrating that ideal adaptation to variance would lead to equal search lengths also under the cognitively less complex LTM assumptions. Vice versa, thus, if participants nonideally scale their value distributions we would generally expect differences in search length and performance between the two conditions. A simulation study for best performing linear threshold agents across the two conditions — equivalent to the above simulation study for the optimal model — confirmed that under the best performing linear threshold model, a change in variance leads to the same performance and search length (see Figure 4, red dotted lines).

Adaptation to variance: Hypotheses on psychological variables. Whereas the normalization hypothesis predicts no difference in thresholds on a relative scale between conditions, the risk hypothesis would predict a difference. Figure 3A and B depicts the adaptation of thresholds according to the normalization hypothesis, when thresholds are measured on an absolute-scale value distribution (Panel A) or on a normalized scale (i.e., standard-normal scale; Panel B). Figure 3 C represents the hypothesis that higher outcome variance leads to more risk-averse behavior, reflected in higher decision thresholds (θ_0) and potentially also higher adaptation rates (δ), when expressed on a standardized scale (see Figure 3B).

Behavioral results

Search length and performance. The data suggest that people adjust their choices according to the normalization hypothesis. Figure 4A indicates that the average accepted price differs greatly between conditions ($SD = 10$: $M = 169.4$, $SD = 1.4$; $SD = 40$: $M = 136.7$, $SD = 5.3$). However, on a normalized scale, this difference seems to disappear (Figure 4B; $SD = 10$: $M = -1.06$, $SD = 0.15$; $SD = 40$: $M = -1.08$, $SD = 0.13$). A Bayesian t test (R BayesFactor::ttestBF package, prior scale = medium; Morey et al., 2018) supports these findings, showing strong evidence for a difference in accepted prices between conditions ($M_{Diff} = 32.73$, 95% confidence interval [CI]: [32.0,




Figure 3. Threshold adaptation predictions according to the theory of normalization and the theory of risky decision making. Panels A and B: Normalization hypothesis: Thresholds differ on an absolute scale (A) but are identical on a normalized scale (B, when $\mu = 0$ and $\sigma = 1$). Panel C: Variance-risk hypothesis: High-variance environment implies higher risk, therefore thresholds in the high-variance environment are higher (more relaxed, red line) than in the low-variance environment (blue line) on a normalized scale.

33.4], Bayes factor [BF] > 300) but inconclusive evidence that normalized accepted prices differ ($M_{Diff} = 0.02$, 95% CI [0.004, 0.04], BF=0.45). Furthermore, investigating participants' performance and search length between conditions revealed substantial to strong evidence that they are equal (Figure 4C and D; search length: $M_{\sigma=10}^{SL} = 5.03$, $M_{\sigma=40}^{SL} = 4.99$, $M_{Diff}^{SL} = 0.05$, 95% CI [-0.07, 0.16], $BF_{10} = 0.11$), performance: ($M_{\sigma=10}^{Perf} = 21.24$, $M_{\sigma=40}^{Perf} = 21.36$, $M_{Diff}^{Perf} = -0.058$, 95% CI [-0.14, 0.025], $BF_{10} = 0.27$). These results speak against the risk hypothesis, which predicted differences in search length and performance. Instead, the results support the hypothesis that, from an LTM or optimal perspective, values are perceived according to their ranks within the distribution. Otherwise, we would have expected behavioral measures to differ between conditions. However, although the adaptation to variance seems in line with the optimal model (Figure 4, black circles) and the predictions of the best performing linear model (Figure 4, red triangles), the actual behavior still deviates from both predictions, as people performed suboptimally in both conditions. However, it also shows that people's deviations can at least partly be explained by the assumption on the linearity of decision thresholds. Therefore, we next investigated if the LTM or three other choice strategies describe the data, to further examine the underlying search processes that

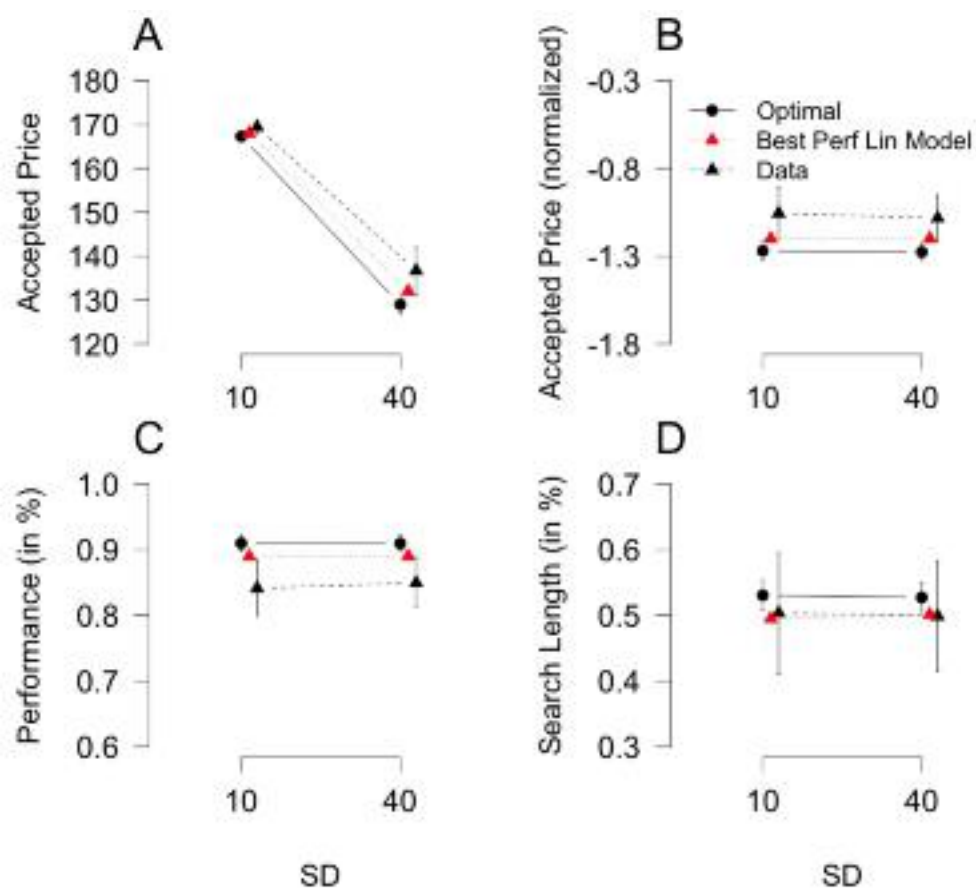


Figure 4. Search length is measured as the number of “reject” decisions in a sequence plus 1. Performance was defined as the average sum of accumulated points in each trial. In red: predictions of the best performing linear threshold model.

may lead to suboptimal behavior.

None of our models assume any learning over trials. This assumption was in line with an analysis of the accepted prices across trials. A linear mixed model on accepted prices per trial with trial number as fixed effect and by-participant random intercepts for trial number showed no significant effect of trial number in either variance condition [low variance: $F(1, 118) = 2.1, p = .15$, high variance: $F(1, 102) = 1.9, p = .16$].

Modeling Participants’ Decisions

In this section we explore participants’ choice strategies by comparing possible cognitive models that could account for their choice behavior. Furthermore, we analyse adaptive choice behavior based on the underlying choice processes captured in the model’s cognitive parameters. Besides the LTM we consider three models that have been used in the optimal stopping literature: the BOM (a variant of the bias-from-optimal model

described in Guan et al., 2015) that determines people’s thresholds in terms of how much they deviate from optimality, a fixed threshold model (FTM; Goldstein et al., 2020; Lee et al., 2004) that assumes a stable threshold across search, and an independent threshold model (ITM; Baumann et al., 2020; Guan et al., 2015) that assumes no relationship between thresholds. The ITM thus allows for any shape of thresholds across search and sets the benchmark on how any threshold model can describe optimal stopping behavior. Sampling models that base their stopping decisions on information gained from an initial sample of options are not considered here as studies have shown that such models are not appropriate for capturing human behavior in the full-information version of the optimal stopping task (see e.g. Baumann et al., 2020; Lee et al., 2004; Sang et al., 2020).

The models were implemented in a *hierarchical Bayesian statistical framework* using JAGS (Plummer, 2003). In a Bayesian framework, information regarding model parameters is represented by probability distributions. The data were used to update *prior distributions*, resulting in *posterior distributions*, which were used for inference. A hierarchical implementation allowed us to fit data on the individual-trial level while simultaneously taking into account information shared across participants via group-level distributions. Fitting involved running four independent chains, each with 2,000 samples drawn from the posterior distribution, with a burn-in period of 100 samples. The chains converged, which was monitored via the calculation of Gelman–Rubin statistics on the four chains, autocorrelation plots, and visual inspection of the chains. In the subsequent analysis, we evaluated the parameter estimates in both variance conditions and analyzed them in terms of stability and predictability of choice behavior. A description of all models’ priors can be found in the Appendix (Method section). The following section introduces the four models using the example used in our studies.

Models

We consider a decision maker who encounters a sequence of ticket prices with values denoted by x_1, \dots, x_{10} who wants to find the minimum value in the sequence. If the decision maker accepts the price x_i , the sequence terminates and they have to pay x_i ; otherwise, they continue to the next price. When the last ticket is reached, it must be accepted.

All models assume that the decision maker relies on a probabilistic threshold to make the decision to accept or reject a price—that is, price x_i on position i is compared to a position-dependent threshold t_i . This comparison yields an acceptance probability θ_i based on a sigmoid choice function with sensitivity parameter β and

$$\theta_i = \frac{1}{1 + \exp\{\beta(x_i - t_i)\}}. \quad (2)$$

The parameter β (sensitivity) governs how deterministic the decision threshold is. Small values of β reflect stochasticity in decisions, whereas the policy approaches determinism with larger values. In other words, large values of β indicate that a decision maker applies the decision threshold very consistently, and low values indicate sporadic violations (e.g., either strategically to explore remaining options once in a while, or due to less precise cognitive processes during the evaluation; Bendor et al., 2001).

The LTM (Baumann et al., 2020) postulates that the decision threshold changes over time in a linear fashion. Formally, it is defined by the first threshold θ_0 and the incremental value δ , which is added to θ_i whenever a price is rejected:

$$\theta_{i+1} = \theta_0 + \delta \cdot i. \quad (3)$$

An increasing θ_i over time (i.e., positive values of δ , in the current example) means that the decision maker becomes more likely to accept more expensive prices compared to the previous sequence positions. In sum, this model entails three adjustable parameters,

the first threshold θ_0 , the increment of the threshold δ , and the choice sensitivity β .

The corresponding psychological interpretations are listed in Table 1.

Table 1

Linear Threshold Model Parameters and Their Psychological Interpretation

Parameter	Cognitive function
θ_0	Aspiration level before search
δ	Adaptation rate during search
β	Sensitivity to differences between aspiration level and actual price

The FTM is a simplified version of the LTM where the decision maker compares an offer to a single fixed threshold θ_F (Goldstein et al., 2020; Lee & Courey, 2021). This model is implemented as a nested model of the LTM by allowing the slope parameter δ to have the value 0.

The BOM is based on the idea (Guan et al., 2015) that people use thresholds that deviate systematically from the optimal thresholds. The optimal thresholds θ_i^* for each position i are derived by determining the expected reward of the remaining options (for the mathematical derivation, see Gilbert & Mosteller, 1966, Section 5b, Table 13 and the Appendix Text 1 & 2). The model entails a systematic bias parameter γ that reflects the divergence of the human threshold from the optimal one. Additionally, the thresholds depend on a parameter α that determines how much their bias increases or decreases as the sequence progresses.

$$\theta_i = \theta_i^* + \gamma + \alpha \cdot i. \quad (4)$$

When γ and α are set to 0, the thresholds represent the optimal thresholds that lead to best performance. This model is therefore defined by three free parameters, γ , α , and the choice sensitivity β . Possible threshold shapes for different parameter combinations

are shown in the Appendix Figure A1.

Finally, the ITM estimates $N - 1$ independent threshold parameters $\theta_1, \dots, \theta_{N-1}$ freely, one for each position in the sequence. The thresholds can take any value across positions and therefore provide an upper limit for how well any threshold model can describe an individual's decision given the assumption of a probabilistic threshold.

Modeling results

We first compared the estimated thresholds from the BOM and the LTM with the estimates from the ITM, which serves as an upper limit in accuracy when describing people's thresholds. Figure 5A shows that the LTM's estimated thresholds closely match the estimated thresholds of the ITM in both variance conditions. Because of the thresholds' nonlinear increase across positions, the BOM is less capable of capturing people's decision thresholds. A quantitative analysis of the performance of the models (using the deviance information criterion, DIC; see Spiegelhalter et al., 2002) both variance conditions suggests a preference for the LTM compared to the BOM and the ITM (see Appendix Table 15). To test if the FTM would better represent participants' strategies, we investigated if the group-level posterior distribution (95% CI) of the slope parameter δ includes 0. For both variance conditions, this was not the case ($\sigma = 10$: 0.7, 95% CI [0.6, 0.8]; $\sigma = 40$: 2.4, 95% CI [2.2, 2.8]; see Table 2), suggesting that thresholds are adjusted across search. Inspection of the distribution across the estimated individual δ s for both conditions provided further support for this finding (see Appendix Figure 3).

To further investigate the accuracy of the LTM in describing choices, we studied the agreement of the acceptance probabilities between participants and the model. Figure 5B (low variance condition) and C (high variance condition) shows participants' acceptance probabilities for prices split into quantiles on each position (black lines) together with the posterior predictive means of the LTM, suggesting a fairly accurate agreement. We conducted the same analysis for the BOM (see Appendix Figure 2), which shows a weaker agreement with the data, in particular due to the model's

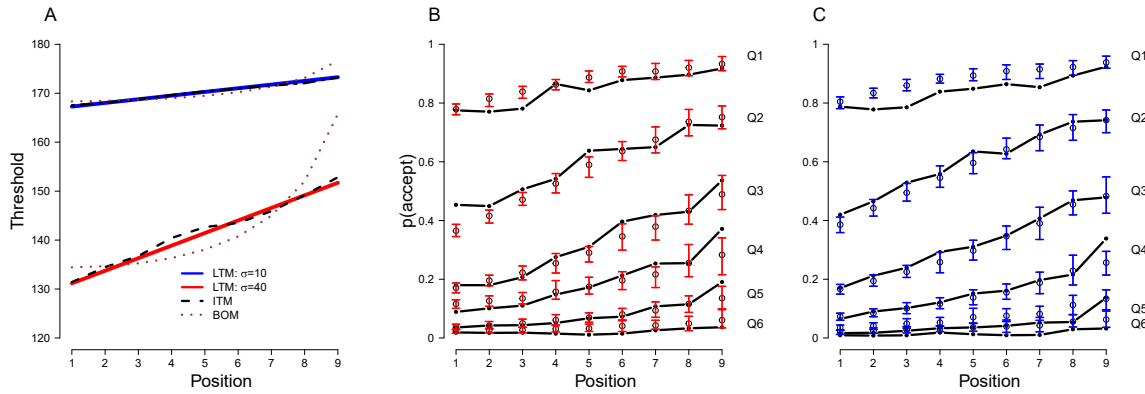


Figure 5. Panel A: Estimated thresholds in both conditions for three models (see legend). Panels B (high variance) and C (low variance): Probability of accepting across positions, agreement between data (black lines) and the LTM (posterior predictive means and 95% confidence intervals). Q1: Prices in the first quantile; Q2: prices between the first and second quantile, etc. Only the first six quantiles are shown (of 10); BOM = biased optimal model; ITM = independent threshold model; LTM = linear threshold model.

assumption of a nonlinear threshold increase.

Before analyzing the best fitting parameter values of the LTM, it is important to check whether the fitting procedure gives meaningful parameter values when fitting fake data where the “true” parameter values are known (Nilsson et al., 2011). Such a procedure is known as parameter recovery, and it is a crucial part of a model-based analysis.

Scatterplots and correlations are shown in Appendix Figure 4, revealing a fairly good agreement between the LTM’s generating and fitted parameter values ($\rho > .85$ for all parameter correlations).

Cognitive parameter estimates in the low- and high-variance conditions.

The behavioral evidence of equal search length and performance across conditions indicated that participants adjusted their decision thresholds according to the normalization hypothesis. The LTM allows one to directly estimate this threshold and thereby to substantiate this claim in both variance conditions. The corresponding mean threshold parameters are plotted in Figure 6 (group-level means and individual estimates). As can be seen, the absolute thresholds differed substantially between the two conditions (Panel A) but not when the model was applied to the normalized values (Panel B). The corresponding group-level parameters of θ_0 and δ are listed in Table 2 and reveal no difference between the two variance conditions (see third column).

The normalized values further allow a direct interpretation according to z values, indicating the initial acceptance percentiles and the rate of subsequent adaptation (e.g., “I accept prices that belong to the 0.1 quantile of the distribution, and in each search step I increase the quantile by 0.01”). Interestingly, as the final threshold values (normalized) were still below 0 (50%), this also means that almost all participants were risk seeking in the ninth trial. In the following section, we investigate the robustness of the cognitive parameters between conditions to understand if these processes provide a more stable measurement than search length or performance.

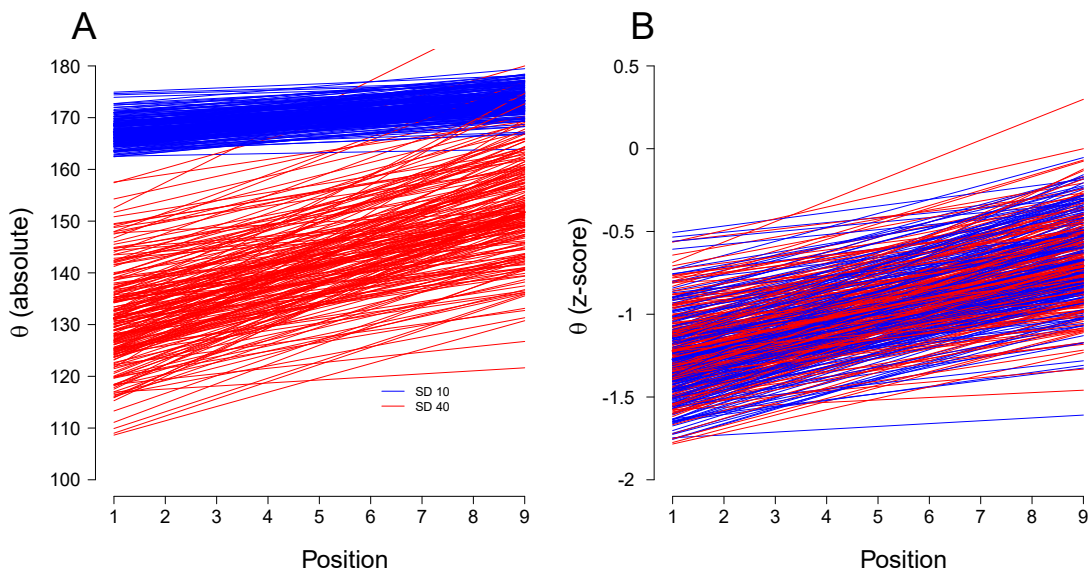


Figure 6. Individual-level threshold parameters for both variance conditions (blue: low variance, red: high variance). The two darker lines show the thresholds based on the group-level parameters. Panel A: Thresholds estimated using absolute prices. Panel B: Thresholds estimated using normalized prices ($\mu = 0$ and $\sigma = 1$).

Stability of cognitive parameters between conditions. To interpret the robustness of parameter estimates, we first examined how behavioral measures, such as search length and performance, correlated between the variance conditions. That is, the previous modeling analyses indicated that performance and search length depend on multiple variables (initial threshold, adaptation, and sensitivity). This means these behavioral measures might be more variable (correlate less strongly) than their underlying cognitive parameters. Thus, if the parameter estimates were meaningful, then, ideally, they should correlate between conditions more strongly than—or at least as much as—the behavioral measures, which provides a reference point for evaluation.

Table 2

Posterior group-level mean parameters of the linear threshold model (LTM) and 95% confidence intervals (in square brackets) for the low- and high-variance conditions. Top three rows: Absolute parameter values; bottom three rows: normalized ($\mu=0$ and $\sigma=1$) parameter values.

LTM parameter	Group-mean parameter estimates		
	$\sigma = 10$	$\sigma = 40$	Difference ($\sigma_{10} - \sigma_{40}$)
θ_0 first threshold	167.6 [167.2, 168]	131.8 [130.3, 133.2]	35.8 [34.36, 37.32]
δ slope	0.74 [0.68, 0.8]	2.67 [2.45, 2.9]	-1.9 [-2.17, -1.7]
β choice sensitivity	0.40 [0.37, 0.41]	0.1 [0.09, 0.11]	0.3 [0.27, 0.32]
θ_0 normalized	-1.17 [-1.21, -1.14]	-1.20 [-1.24, -1.17]	-0.04 [-0.08, 0.02]
δ normalized	0.07 [0.068, 0.08]	0.06 [0.061, 0.07]	0.007 [-0.001, 0.15]
β normalized	3.61 [3.43, 3.78]	3.89 [3.69, 4.01]	-0.21 [-0.26, 0.026]

Search length between conditions correlated with $r = .51$ (95% CI [.4, .6], BF > 300) and performance with $r = .50$ (95% CI .38, .60, BF > 300; R

BayesFactor::correlationBF package, prior scale = medium; Morey et al., 2018).

Indeed, the initial acceptance level θ_0 yields a stronger correlation of $\rho = .62$, 95% CI [.59, .65] (see Appendix Table 1 for an overview of all correlations) and the adaptation rate of $\rho = .53$, 95% CI [.41, .61]. These associations are higher than or equal to,

respectively, the correlations of search length and performance, indicating that the cognitive parameters of the LTM represent a more robust measure between variance conditions. Furthermore, the LTM parameters did not strongly correlate with each other, except for θ and δ , indicating mostly independent contributions to behavior.

When the LTM parameters more robustly describe behavior than the behavioral measures themselves, then the consequent question is how the parameters relate to performance measures and search length. Since both measures did not differ between the variance conditions we readdress these questions in the next study in which, to foreshadow, search length was affected by time horizon, which allowed us to delineate participant characteristics from the effects of additional manipulations. Note here that

the initial threshold was the primary predictor of search length and performance in both variance conditions, illustrated in the scatterplot provided in the Appendix Figure 5.

Discussion Study 1

The goal of Study 1 was to understand how people adjust to environments changing in option values' variance when sequentially searching for the best option. Whereas economic theories (Rangel & Clithero, 2012; Stewart, 2009) suggest that people transform absolute values (and thus variance) into a normalized scale on a cognitive level (normalization hypothesis), some theories in risky decision making (Genest et al., 2016; Holt & Laury, 2002; Weber et al., 2004) postulate that higher variance in the price distribution causes more risk-averse behavior. To test this, we manipulated the variance of ticket prices in a sequential stopping task. The comparison of search length and performance between the two variance conditions did not reveal any effect. This finding suggests that outcome variance in a sequential search task does not affect risk aversion, because otherwise we would have expected shorter search in high-variance environments. To specify people's choice strategies we implemented different models from the optimal stopping literature in a computational framework and demonstrated that an LTM best approximated the behavioral data (regarding acceptance proportions in value quantiles), supporting the appropriateness of our account. In line with the behavioral finding, the LTM (Baumann et al., 2020) captured the lack of difference in search performance by scaling the thresholds relative to the variance in the corresponding price distributions. This implies that participants accepted prices in both conditions within the same percentile on each position, supporting the assumptions of the normalization hypothesis.

The model parameters consistently correlated between the low- and the high-variance conditions, ranging between $\rho = .53$ and $\rho = .62$, which is higher than correlations between behavioral measures such as search length and performance. In particular the initial threshold, which yields a correlation of $\rho = .62$, seems to be a robust underlying parameter when adjusting between contexts. Studies examining test–retest correlations

of behavioral measures in related search paradigms (e.g., the balloon analogue risk task; Frey et al., 2017; White et al., 2008) have reported lower correlations ($\rho = .52$) in behavioral measures, suggesting that the underlying search parameters depicted in the LTM may uncover a more stable process of search across contexts.

Study 2

In the second study, we focused on the influence of the time horizon on adaptive choice behavior in optimal stopping tasks. The sequences contained 5, 10, or 20 options. Otherwise, the task was equivalent to that in Study 1, where participants were conducting an online ticket-shopping task with the goal of finding the cheapest price.

Method

Participants. We recruited 70 participants (20 women; age range: 25–60 years) on Amazon Mechanical Turk to participate in the three-condition, within-subject experiment. We selected a sample size equivalent to that in a similar study (Baumann et al., 2020). Participants gave informed consent, and the University of Zurich Committee on the Use of Human Subjects approved the experiments. Participants were excluded from analysis if they accepted the first option in a trial in more than 95% of the trials, leaving 66 participants in the subsequent analysis.² They received a fixed payment of \$2 and a performance-dependent bonus ranging between \$0 and \$6 (see Procedure for a description of the calculation).

Procedure. The second experiment employed the same task as in Study 1. The sequence length was varied, offering 5, 10, or 20 values. All prices were generated from the same distribution throughout the experiment [$\sim \mathcal{N}(\mu = 180, \sigma = 20)$]. All participants worked on a total of 180 trials, split into two parts of 90 trials. Both parts consisted of three blocks with 30 trials, where trials within a block were of length $N = 5, 10, \text{ or } 20$. The order of the blocks was randomized. The maximum number of 33

² We examined the individual-level data (plotting accept probabilities for each individual in each condition) and are confident that random behavior (e.g., randomly choosing one of 10 positions, $p = .1$, or randomly choosing to accept with a probability of 50%) was mostly absent (i.e., almost all participants showed an increasing acceptance trend with ongoing positions).

points was awarded when the cheapest ticket was chosen and 0 points for the most expensive one. Intermediate choices equated to

$$points_i = \frac{33 \cdot (p_{max} - p_{chosen})}{p_{max} - p_{min}}, \quad (5)$$

where p_{max} represents the highest price and p_{min} the lowest price in sequence i .

Participants received a base payment of \$2 and earned between \$0 and \$4 additionally depending on their performance. The final bonus was calculated by dividing the total collected points by the maximum number of points. The average bonus earned was \$4.87 (min: \$2.98, max: \$5.30)

Adaptation to the time horizon: Optimal and linear models

To better understand the theoretical predictions regarding the effects of time horizon we first simulated the optimal model and the LTM, beginning with the former. In the optimal model, the decision thresholds change across position and correspond to the expected reward of the remaining options. Thus, initial thresholds are lower for longer time horizons (in the case of finding the minimum price offer, see Appendix Figure 6, black lines). Optimal thresholds therefore depend solely on the remaining options in the sequence, regardless of the absolute sequence length. This also predicts that the choice rates on the last decision (e.g., on Position 4 with Length 5, or on Position 9 with Length 10) will be equal under the optimal model in all length conditions.

For detailed predictions regarding acceptance rates, search length, and performance of an optimal agent, we simulated for each condition ($N = 5, 10,$ and 20) an equivalent number of 60 trials for values sampled from $\sim N(\mu = 180, \sigma = 20)$ for a total of 100 agents. The average accepted ticket price decreases for longer time horizons ($N = 5: M = 161.5; N = 10: M = 154.7; N = 20: M = 148$; see Figure 8A) whereas performance, which is measured as the average accumulated points in each trial divided by the maximum points, increases ($N = 5: M = 0.88; N = 10: M = 0.9; N = 20: M = 0.93$; Figure 8B). In absolute terms, an optimal agent searches more in longer time horizons

($N = 5$: $M = 2.9$; $N = 10$: $M = 5.3$; $N = 20$: $M = 9.9$; Figure 8C), but in relative terms, search decreases (search length/total sequence length, $N = 5$: $M = 0.58$; $N = 10$: $M = 0.53$; $N = 20$: $M = 0.49$; Figure 8D).

If the participants' strategies correspond to the architecture of the LTM, this would also have implications for how parameters should change to maintain near-optimal performance. Thus, to understand how people *might* adjust to the time horizon, we conducted a series of model simulations, illustrated in Figure 7.³ First, Figure 7A shows the predictions of performance (Panel A1) and search length (Panel A2) in the three time horizon conditions (x axes; $N = 5$ vs. $N = 10$ vs. $N = 20$) depending on different aspiration levels (colored circles; θ_0), without changing them over time (i.e., with δ set to 0). Figure 7A1 shows that, within each condition, the maximum performance predicted by the LTM is a curvilinear function of the initial aspiration level. Intuitively speaking, when searching for the minimum, too high aspiration levels lead to accepting too many suboptimal options and too low aspiration levels lead to rejecting too many optimal options.

Second, the connected dots represent identical threshold values across conditions (indicated by their color in Figure 7). As can be seen, the threshold that yields optimal performance with a time horizon of 5 (orange circles) predicts only suboptimal performance with a time horizon of 10 (the same holds for a comparison between 10 and 20). This means, according to the LTM, if participants seek to optimize their behavior, relative to a time horizon of 5, they have to decrease their aspiration levels (less tolerant thresholds). Specifically, with the assumption that there is no adaptation across sequence positions ($\delta = 0$), the LTM predicts that the optimal initial thresholds for the conditions of $N = 5$, 10, and 20 are -0.4, -0.8, and -1.2 (on a normalized scale), respectively. Correspondingly, for these optimal initial thresholds, the LTM predicts that the relative search length (relative to the time horizon; Figure 7A2) decreases with

³ For each sequence length ($N = 5$, 10, and 20), an equivalent number of 60 trials was simulated, where values were sampled from $\sim N(\mu = 180, \sigma = 20)$ for a total of 100 participants, which were aggregated on measures of performance and search length. We approximated the range of the parameters to the range of values obtained in Study 1 and used any possible combination of initial threshold (θ_0) and rate of increase (δ) to simulate choices.

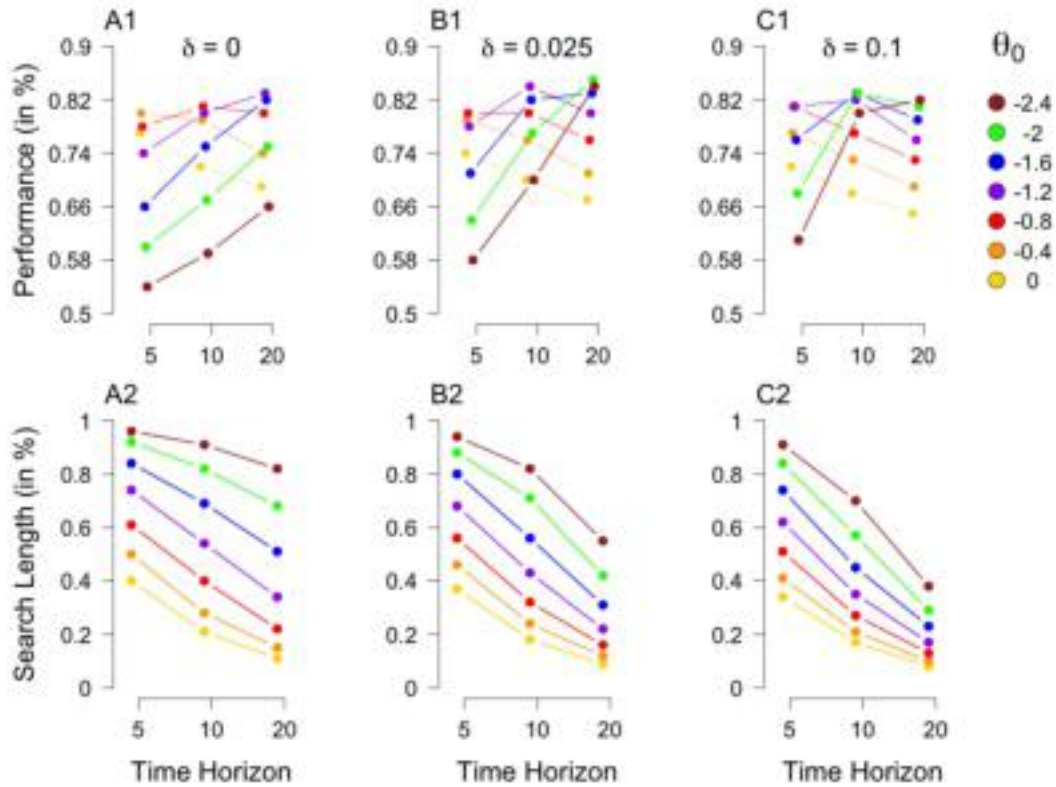


Figure 7. Linear Threshold Model simulation study. Panels A1 and A2: Performance (average number of points/maximum number of points) and search length (i/N ; i = accepted position, N = time horizon) for a range of θ_0 values (see legend) and $\delta = 0$. Optimal thresholds differ between time horizons: $N = 5$: $\theta_0 = -0.4$ (orange); $N = 10$: $\theta_0 = -0.8$ (red); $N = 20$: $\theta_0 = -0.8$ (purple). Panels B1 and B2: $\delta = 0.025$; θ_0 values according to legend. Panels C1 and C2: $\delta = 0.1$; θ_0 values according to legend. Points connected with a line represent the same θ_0 values. In Panels A2, B2, and C2, search length is calculated in relative terms (i.e., i/N ; i = accepted position, N = time horizon).

increasing time horizon.

Up to now, our simulation has assumed that the thresholds remain fixed across the sequence. In the next step, we allowed the initial threshold (θ_0) to change across positions via adaptation rate δ . The simulation results indicate that δ also contributes to optimal performance in the LTM, suggesting differences between the time horizon conditions, as illustrated in Figure 7B (B1 for predictions on performance and B2 for predictions on search length) and C (C1 and C2, as described for B1 and B2). First, with $\delta = 0$ in Figure 7A1, introducing a small adaptation rate ($\delta = 0.025$ in Figure 7B1) leads the previously optimal threshold (e.g., orange circles in Panel A1 for $N = 5$) to become suboptimal. Further, whereas a small adaptation of δ ($\delta = 0.025$) leads to a higher performance when $N = 20$ (Panel B1), a large adaptation (in Panel C1) seems to impede performance in this condition (when $\delta = 0.1$).

This model dynamic can be plausibly verbalized in a psychological way. That is, if participants adapt to the task, they might approach the different time horizons with different initial aspiration levels in the first place (longer time horizon = lower aspiration level θ_0 ; see Panel A1). In the second place, they would regulate their adjustment of the aspiration level over time according to the time horizon, with lower increases in long time horizons, because too strongly adapting the threshold over long time horizons leads to prematurely accepting suboptimal offers. Of course, the LTM does not explain *how* participants would know what the optimal adaptations are, but here we seek to accurately capture these trends in order to pin down the contributing factors, and later we discuss why they might or might not reach optimal levels from an LTM perspective. To compare performances between the optimal and a best linear threshold strategy, we calculated the best performing linear thresholds for each time horizon using grid search to find the parameter values (θ_0 and δ) that result in the highest reward.⁴ As proposed by the simulation study above, optimal adaptation of linear thresholds to time horizons involves decreasing the first threshold and lowering the adjustments across search (see Figure 9B, black dashed lines). Figure 8B (red line) shows that using the best performing linear thresholds reduces performance only slightly relative to the optimal model, by 1.1% when $N = 5$, by 3% when $N = 10$, and by 6.5% when $N = 20$. This result reveals that an ideal adaptation of the linear thresholds to time horizon would lead to almost optimal performance and emphasizes the LTM's potential to efficiently solve the optimal stopping task. Of course, misadjustments might also lead to much worse performance, which we discuss below. Figure 8D (red lines) displays the relative search length by a decision maker who uses the best performing linear thresholds and reveals that relative search length decreases as well with increasing time horizons ($N = 5$: 58%, 95% CI [0.57, 0.59], $M = 2.9$, $SD = 0.2$; $N = 10$: 53%, 95% CI [0.52, 0.54], $M = 5.3$, $SD = 0.36$; $N = 20$: 48%, 95% CI [0.47, 0.49], $M = 9.6$, $SD = 0.6$). Although this result suggests that the linear and the optimal threshold strategy lead to the same aggregated measure of search length, they differ in terms of their prediction on the

⁴ We used the same procedure as in the simulation study that describes the relationship between parameters and behavioral measures (Figure 7).

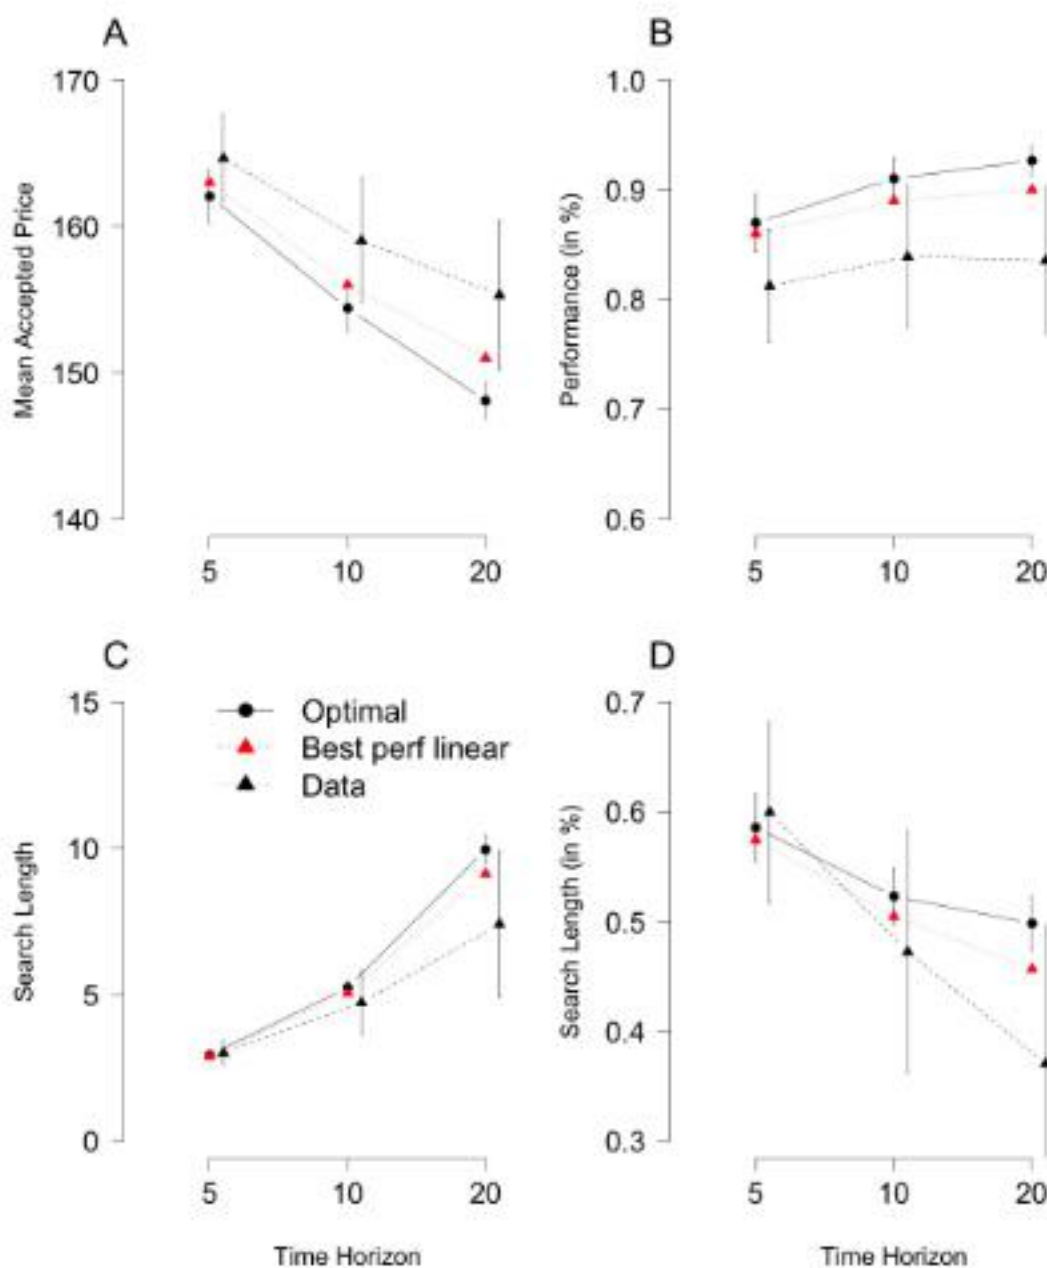


Figure 8. Summaries of (Panel A) mean accepted price, (Panel B) performance (as percentage), (Panel C) search length (absolute), and (Panel D) search length (as percentage) for participants and the optimal model. In red: predictions of the best performing linear thresholds.

distribution of accepted position across trials (e.g., 50% acceptance at Position 9 of 10 in the optimal model but 35% in the LTM; see Appendix Figure 7 for $N = 20$).

Behavioral results

Search length, performance, and choice probabilities. The results reveal strong evidence that the mean accepted ticket prices decreased with increasing time horizons (Figure 8A, $N = 5$: 164.7; $N = 10$: 159.0; $N = 20$: 155.3; Bayesian t tests between $N = 5$ and $N = 10$: $\mu_{Diff}^{price} = 5.4$, 95% CI=[4.1, 6.5], $BF_{10} > 300$ and between $N = 10$ and $N = 20$: $\mu_{Diff}^{price} = 3.5$, 95% CI=[2.6, 4.4], $BF_{10} > 300$), replicating participants' search behavior found in previous studies (Guan et al., 2015; Lee et al., 2004). In other words, participants became more selective in sequences involving a higher number of offers.

Participants' performance (accumulated points per trial/total number; see Figure 8B) increased from $N = 5$ to $N = 10$ ($N = 5$: 81.2%, 95% CI [0.8, 0.3]; $N = 10$: 83.9%, 95% CI [0.82, 0.86]), supported by a t test producing a BF of 12.5 in favor of a difference in performance. However, when $N = 20$, performance did not increase further, reaching on average 83.6% (95% CI [0.82, 0.85], BF = 0.15), indicating that there is no difference between performance when $N = 10$ and $N = 20$. Thus whereas an optimal agent would still improve performance in an $N = 20$ environment (see Figure 8B, black solid line), participants did not.

Considering absolute search length, we found that participants searched longer in extended time horizons, illustrated in Figure 8C ($N = 5$: $M = 3$, 95% CI [2.9, 3.1]; $N = 10$, $M = 4.7$, 95% CI [4.5, 5.0]; $N = 20$, $M = 7.4$, 95% CI=[6.8, 8.0]). However, whereas we found no difference in search length when $N = 5$ between participants and the optimal model (BF = 0.17), participants searched significantly less when $N = 10$ and $N = 20$ (BF > 300 for both comparisons between search length), which was most pronounced when $N = 20$, as also reflected in the participants' relative search lengths (Figure 8D): Whereas in a short time horizon ($N = 5$) participants searched on average 60% of the options (95% CI [0.58, 0.62]), they decreased their relative search length to 47.7% (95% CI [0.45, 0.5]) when $N = 10$ and to 37% (95% CI [33.9, 40.1]) when $N = 20$. To what extent do participants behave optimally under the constraints of linearity? A Bayesian t test indicates that compared to participants' relative search length, there is no evidence of a difference in search length when $N = 5$ ($\mu_{Diff}^{SL} = -0.001$, 95% CI [-0.02,

0.02], $BF_{10} = 0.17$) but strong evidence of a difference when $N = 10$ ($\mu_{Diff}^{SL} = 0.08$, 95% CI [0.06, 0.1], $BF_{10} > 300$) and $N = 20$ ($\mu_{Diff}^{SL} = 0.11$, 95% CI [0.09, 0.14], $BF_{10} = 0.17$; see Figure 8D). These results show that in longer sequences, participants' search length was notably reduced compared to a policy that assumes best performing linear thresholds, by 8% when $N = 10$ and by 11% when $N = 20$. Furthermore, participants achieved up to 92% of the performance when $N = 5$, 91% when $N = 10$, and 89% when $N = 20$, which indicates of reduction of optimal adaptation in longer sequences (Figure 8B). Yet, some participants arrived at levels of performance competitive with the best performing linear threshold rule, whereas others achieved only 47%. First, these findings suggest that suboptimal behavior cannot be explained by strategic constraints (optimal vs. best performing LTM) alone; otherwise, participants might have reached the LTM's best performance. Second, there are clear individual differences in how people adapted to time horizons. Although some participants may have behaved optimally relative to the best performing linear thresholds, others showed suboptimal adaptation to extended time horizons. The causes of individual differences in adaptive behavior may be found in the maladaptation of initial threshold and adaptation rate. In particular, participants' decrease in search length (and thus lower performance) in longer sequences implies that they either placed their initial threshold too liberally or adapted it too strongly over time, which we further address via modeling below.

Modeling results

We applied the same four models as in Study 1 to participants' choices: the LTM, the BOM, the FTM, and the ITM. We then analyzed how underlying parameters are adjusted between time horizons and then investigated the reasons for participants' performance decrease in longer sequences. Finally, we investigated the parameters related to behavioral measures across changing time horizons.

Model comparison. Figure 9A shows the hierarchical group-mean thresholds for the LTM (dark solid lines), the BOM (orange dashed lines), and the ITM (dotted blue lines) for all three time horizon conditions. Apparently, the LTM and the ITM

thresholds are almost congruent when $N = 5$ and $N = 10$. However, when $N = 20$, the overlap of the thresholds is reduced, mainly due to a spike in the probability of accepting higher ticket prices on Position 19, which cannot be captured by a linear threshold. Although this spike seems to be in line with the BOM, it is much less pronounced than predicted, indicating that the BOM is not well suited to capturing participants' choices. We note that in this case, a more flexible modeling approach proposed by Lee (2006) that describes behavior by allowing a combination of fixed and linear adjustments across search could capture such a spike in the acceptance rate at the end of the extended sequence, which should be addressed in further studies. A quantitative analysis of the performance of the models (using the DIC; see Spiegelhalter et al., 2002)) for the three length conditions suggests a clear preference for the LTM compared to the BOM and the ITM (see Appendix Table 16).

Qualitative evidence of the agreement between the LTM and the data is provided in Appendix Figures 8–10; The LTM fairly adequately captures acceptance probabilities for each quantile at every position, in all three time horizon conditions. Agreements between the data and the BOM are shown in Appendix Figure 11, showing that the BOM's assumption of nonlinearly increasing thresholds is implausible and is most pronounced when $N = 5$ and $N = 10$.

We used a parameter recovery analysis to ensure that the LTM provides meaningful parameters in extended or shortened time horizons. The correlation of the simulated and recovered parameters is rather high in all time horizon conditions, ranging between .85 and .97. We refer to the Appendix Figure 12 for the correlation plots and the description of the procedure.

Best fit parameters. To address the question of which cognitive factors contributed to the suboptimal adjustments to time horizons, we compared the LTM's estimated parameters between the conditions. The estimated by-position thresholds for each time horizon condition are plotted in Figure 9B (solid lines), along with the thresholds obtained with the individual-level parameters. Table 3 summarizes the group-mean parameter values of θ_0 , δ , and β in each time horizon condition and shows that both θ_0

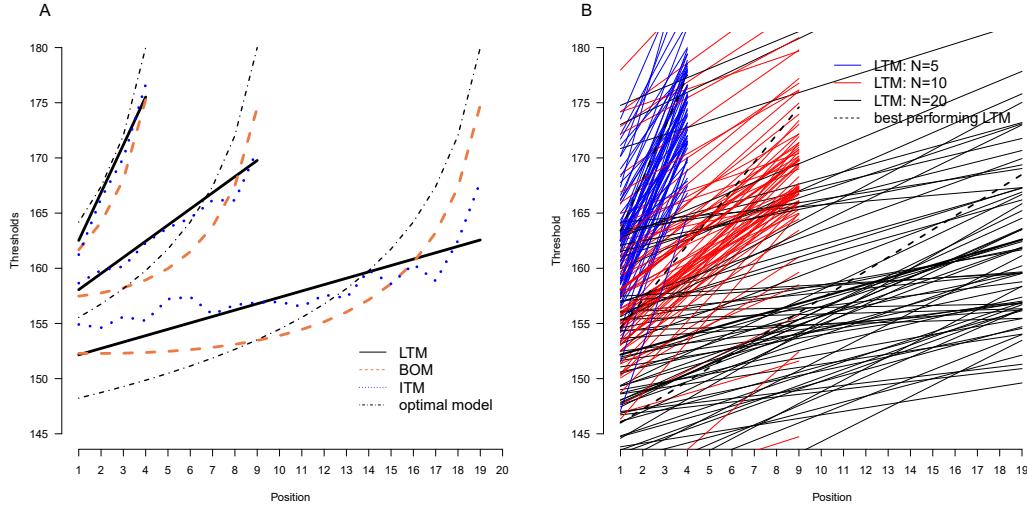


Figure 9. Panel A: Hierarchical group-mean thresholds estimated from the linear threshold model (LTM; dark solid lines), the independent threshold model (ITM; dotted blue lines), the biased optimal model (BOM; orange dashed lines), and the optimal model (black dot-dashed line). Panel B: Just the LTM, repeating the group means from Panel A for this model and in addition showing the individual-level threshold. Black dashed lines: Best performing linear threshold parameters.

Table 3

Posterior group-level mean parameters for each sequence length and difference between the posterior distributions ($N5 - N10$, $N10 - N20$)

LTM parameter	Group-mean parameter estimates				
	$N = 5$	$N = 10$	$N = 20$	$N5 - N10$	$N10 - N20$
θ_0	161.1 [159.4, 162.7]	156.2 [154.2, 158.2]	151 [148.6, 153.3]	4.9 [3.5, 6.2]	5.2 [2.1, 8.3]
δ	4.4 [3.8, 4.9]	1.48 [1.24, 1.71]	0.58 [0.46, 0.71]	2.9 [2.4, 3.3]	0.9 [0.6, 1.1]
β choice sensitivity	0.21 [0.18, 0.23]	0.19 [0.16, 0.21]	0.17 [0.15, 0.19]	0.01 [-0.02, 0.04]	0.019 [-0.01, 0.05]

Note: Values in square brackets are the 95% confidence intervals. LTM = linear threshold model.

and δ are reliably reduced with increasing time horizon. Columns 4 and 5 (Table 3) report the means of the difference of the population parameter's posterior distributions and indicate strong evidence of an adaptation of the first aspiration level (θ_0) and its adjustment across time (δ) between time horizons (difference is credibly different from zero). These adjustments are thus in line with the type of adjustment required to become more optimal, as predicted in the simulation study (Figure 7). However, the main driver of the too short search length was the initial aspiration level, which becomes evident when correlating the parameters with the behavioral measures (see below).

To understand the reasons for participants' decrease in search length and performance in extended time horizons, we compared the underlying parameters with the parameter

of the best performing LTM (Fig. 9B, dashed lines). The best performing linear threshold rule predicts an adaptation of the initial thresholds proportional to time horizon (doubling the sequence leads to a decrease in the initial threshold by a constant amount) and the adjustment across position (doubling the sequence leads to halving the adjustment rate). Participants followed the same trend: The initial threshold was decreased by a constant amount when the sequence size was doubled, from $N = 5$ to $N = 10$ by 4.9 (95% CI [3.5, 6.2]) and from $N = 10$ to $N = 20$ by 5.2 (95% CI [2.1, 8.3]; see column 4 in Table 3). The adjustment rate of the threshold, δ , is reduced with increasing sequence length, from 4.4 (95% CI [3.8, 4.9]) when $N = 5$ to 1.48 (95% CI [1.24, 1.71]) when $N = 10$ and to 0.58 (95% CI [0.46, 0.71]) when $N = 20$. Therefore, although participants appear to have followed the same regularities in adapting the initial threshold θ_0 and its adjustment rate δ to time, they decreased θ_0 too little (thus accepting too many options in the beginning of the sequence) but reduced δ too strongly (thus departing too little from the initial threshold across search) with longer time horizons. This result shows that both the maladaptation of the initial threshold and the adjustment of the threshold contributed to the increased deviation from optimality in extended sequences.

Stability of LTM parameters. As for Study 1, we next calculated the correlation (R BayesFactor::correlationBF package, prior scale = medium; Morey et al., 2018) of the individuals' performance measures between conditions as a reference point for evaluating model parameter correlations. Search length between conditions revealed a correlation of $\rho = .75$ between $N = 5$ and $N = 10$ (95% CI [.6, .8], BF > 300) and $\rho = .77$ (95% CI [.71, .83], BF > 300) between $N = 10$ and $N = 20$. The correlation of performance between conditions is lower between $N = 5$ and $N = 10$, $\rho = .59$ (95% CI [.42, .7], BF > 300), and between $N = 10$ and $N = 20$, $\rho = .79$ (95% CI [.7, .82], BF > 300).

Regarding the LTM parameters (see Appendix Table 2 for all correlations), θ_0 (initial decision threshold) showed the strongest correlations ($\rho = .8$, 95% CI [.7, .87] between $N = 5$ and $N = 10$ and $\rho = .73$, 95% CI [.61, .83] between $N = 10$ and $N = 20$),

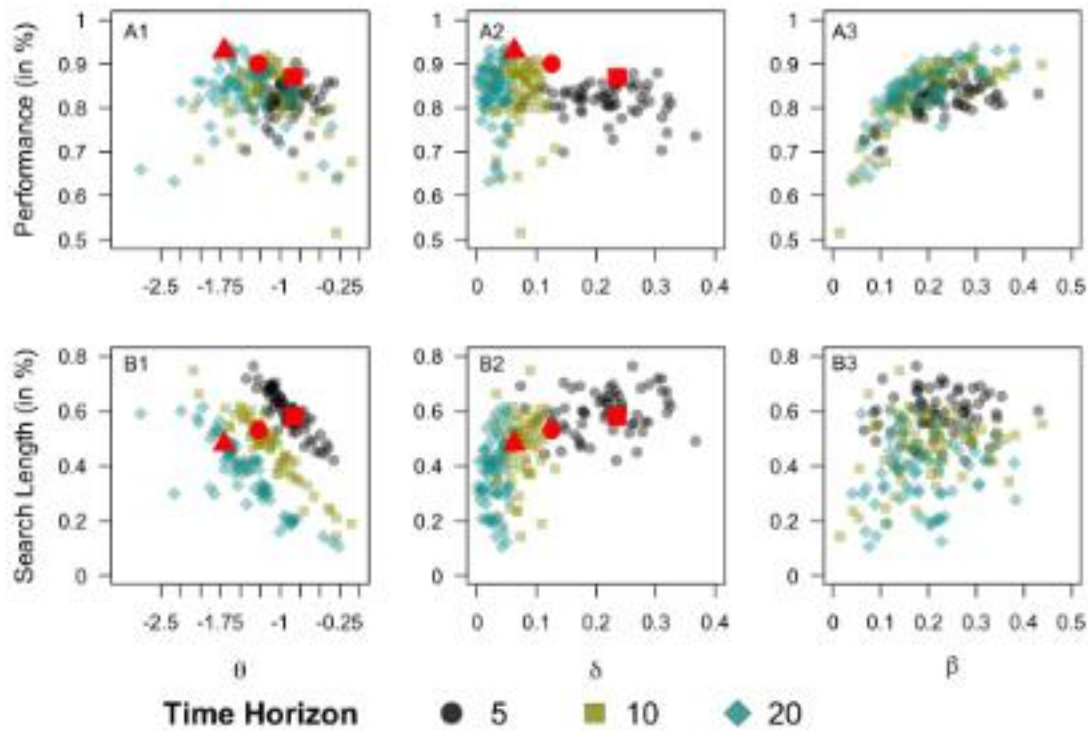


Figure 10. Study 2: Panels A1–A3: Scatterplots of search length (as percentages, search length/total sequence length) versus parameters (θ_0 , δ , and β). Panels B1–B3: Performance (as percentages, points/maximum points) versus parameters. Black: $N = 5$, pea green: $N = 10$, blue green: $N = 20$. In red: Best performing linear threshold parameters and their performance and search length.

yielding a stronger stability between conditions than behavioral measures. The correlation coefficient of δ (adaptation rate) was smaller ($\rho = .51$, 95% CI [.32,.67] between $N = 5$ and $N = 10$ and $\rho = .27$, 95% CI [.04, .47] between $N = 10$ and $N = 20$), which could be an indication that particularly in long sequences ($N = 20$) this parameter is estimated with higher uncertainty in light of relatively weak adaptation rates. In any case, especially the high correlation of θ_0 suggests that the initial aspiration level is an individually stable predictor of search length, and that the manipulation of time horizon had an additive effect on the participants' aspiration level, which is surprisingly equal in magnitude for all participants, which we address next.

Cognitive parameters capturing behavioral patterns in changing time

horizons. Since the LTM parameters are fairly stable and the initial threshold yields an even higher correlation than the behavioral measures themselves, we address here the question of how they relate to each other. The scatterplots in Figure 10 display the relation between parameters and behavioral measures for each condition, as indicated

by the symbol colors. The scatterplot displaying the relationship between performance and parameters (Figure 10 A1) suggests an inverse U-shaped relationship between θ_0 and performance, consistent with the prediction from the simulation study of the LTM (Figure 7). A Bayesian regression analysis (R `BayesFactor::lmBF` package, prior scale = medium; Morey et al., 2018) reveals that in short sequences ($N = 5$), θ_0 is not a significant predictor of performance, whereas in longer sequences ($N = 10$ and $N = 20$), the best model is composed of θ_0^2 , accounting for about 45% and 44% (respectively) of the variance in performance (see Appendix Tables 9 and 10 when $N = 20$ and Tables 13 and 14 when $N = 20$).

This result indicates that with increasing time horizon, the initial threshold becomes predictive of performance, whereas setting it too high or too low leads to poorer performance. Comparing participants' initial threshold with the best performing linear thresholds (Figure 10 A1, red shapes) shows that about 34% of the participants set their initial threshold higher than the best performing linear threshold when $N = 5$ and thus were less picky. This proportion increases with extended time horizons, to 67% when $N = 10$ and 77% when $N = 20$, revealing that the reduction in performance in longer sequences originates from participants' too relaxed initial thresholds.

The correlation between the initial threshold and search length reveals a negative linear relationship in all time horizons (Figure 10 B1) and also indicates the systematic differences between conditions. A Bayesian linear regression analysis (R `BayesFactor::lmBF` package, prior scale = medium; Morey et al., 2018) supports these assumptions: For each time horizon condition, that is, $N = 5$, $N = 10$, and $N = 20$, θ_0 is an important predictor of search length and accounts for 76%, 81%, and 66% (respectively) of the variation in search length (see Appendix Tables 3 and 4 when $N = 5$, Tables 7 and 8 when $N = 10$, and Tables 11 and 12 when $N = 20$). These results suggest that the initial aspiration level is the main predictor for participants' search length in different time scales, and that time horizon has an almost constant additive effect regardless of a participant's aspiration level.

The threshold adjustment across search, captured in the parameter δ , does not appear

to be related to performance or search length in all conditions (Figure 10 A2 and B2), as found in Study 1. A regression analysis confirmed this finding, showing that δ is not an important zero-order predictor of performance. Additionally, we observed that when $N = 5$, about 60% of the participants used a lower δ value than predicted by the best performing linear thresholds. With increasing time horizon, however, δ was not optimally adjusted, as 97% of the participants used lower adjustments across search when $N = 10$ and even 100% of the participants when $N = 20$. This result illustrates that increased length leads to both a too relaxed initial threshold and a more pronounced adherence to this initial threshold, thus contributing to a lower performance in longer optimal stopping tasks.

The parameter β (choice sensitivity) is a further strong predictor of the variation in performance, where higher values of β result in higher payoffs. However, the manipulation of time horizon did not affect the level of choice sensitivity, as evident in Figure 10 A3. For additional comparisons between high- and low-performing individuals in relation to the choice sensitivity in each length condition, see Appendix Figure 13. According to Figure 10 B3, there is no relationship between search length and the parameter β .

Discussion Study 2

The main purpose of the second study was to investigate the systematic influence of time horizon on adaptive choice behavior in optimal stopping tasks (searching for the cheapest ticket) and its cognitive description using the LTM (Baumann et al., 2020). Participants' search (in relative terms, search length/total length) decreased in longer sequences as predicted by the optimal model. However, the reduction in search when $N = 20$ is more pronounced. This trend is reflected in participants' performance, which increased from $N = 5$ to $N = 10$ but did not further improve when $N = 20$, which is in contrast to the optimal model's prediction of steadily increasing performance with extended time horizons.

We used a Bayesian hierarchical approach to compare different choice strategies and

found that the LTM accurately describes choice behavior across time horizons. This analysis allowed us to further study the underlying processes that govern adaptive behavior. The effect of time horizon was captured by changes in initial decision thresholds (aspiration level) and the adaptation rate of the thresholds with ongoing sequence positions (i.e., with growing number of rejections). More specifically, with longer time horizons, the initial aspiration level decreased, reflecting more frequent initial rejections for longer time horizons. Similarly, the adaptation rate of the thresholds across sequence positions was weaker for longer time horizons, reflecting a slower change in the probability of accepting ticket prices.

However, a comparison with a best performing LTM (using grid search to find the best performing linear thresholds) revealed that not all participants adapted their thresholds perfectly and thus some had suboptimal performances. Especially in longer sequences, as for $N = 10$ and $N = 20$, some participants tended to set their initial aspiration level too high and adjusted it too little, which resulted in a suboptimal early termination. Moreover, people showed a higher resistance to adjusting their initial aspiration level in extended time horizons, which further contributed to a lower performance. Taken together, the results indicate that participants' performance reduction relative to the optimal model is in line with the predictions of the LTM. However, participants failed to perfectly adjust the initial aspiration level and its adjustment across search to a longer time horizon, leading to a more pronounced decrease in performance. Several reasons could account for people's maladaptation to search in extended sequence lengths, such as increased search costs or higher uncertainty, which we explore in the General Discussion.

We found fairly high correlations between the parameter values, indicating that these processes reflect stable processes in individual search behavior. Importantly, as found in Study 1, the initial threshold reveals a higher correlation across time horizons than behavioral measures and it is highly predictive for search length and performance. This result highlights the importance of considering the underlying processes to better understand sequential search. Indeed, research on other search domains has frequently

emphasized that individual search might underlie a general mechanism that affects search across different domains and tasks (Hills et al., 2008; Hills et al., 2015; Mata & von Helversen, 2015; Pirolli, 2007). However, attempts to link behavior across different exploration–exploitation tasks have been unsuccessful (e.g., von Helversen et al., 2018). Stable cognitive process parameters in the sequential search tasks may thus help reveal the common mechanism of search across different exploration–exploitation tasks.

General Discussion

Our goal was to learn about the aspects of adaptive behavior in light of outcome variance (Study 1) and various time horizons (Study 2) in an optimal stopping task. Research has shown that people use simple strategies to solve the task that may limit how well they can adjust to these contexts. Here we studied adaptive behavior by taking into account the constraints imposed by people’s cognitive choice strategies. The knowledge we gained then allowed us to identify underlying processes responsible for maladaptations to decision environments that go beyond such constraints.

Our analysis shows that an LTM (Baumann et al., 2020; Lee & Courey, 2021; Song et al., 2019) accurately captures human choices across different contexts. The model assumes that the decision maker’s choice strategy can be described by three parameters, the initial aspiration level, its adaptation over time, and choice sensitivity (response determinism). The model makes it possible to directly measure behavioral adaptation on the parameter level and thus is quite suitable for our studies on context effects.

The optimal solution for the variant of the optimal stopping task studied here is based on thresholds that are determined by their relative rank within the distribution and scaled to the specific value distribution. Such an adaptation leads to identical behavioral measures between contexts that differ only in variance. A simpler linear threshold strategy is capable of adjusting to variance, as proposed by the optimal model, thus leading to no loss in performance between conditions. We found that participants’ performance was unchanged between variances, reflected in linear thresholds that are identical between variance conditions when standardized. This

finding is in line with the normalization hypothesis (Rangel & Clithero, 2012), which assumes that the value of an option is perceived according to its rank within the sample. In Study 2 we investigated adaptive behavior in an environment of changing time horizons. To achieve the highest performance in extended time horizons, the linear threshold strategy proposes decreasing both the initial threshold and its adjustment. Participants followed this pattern but failed to adjust both cognitive parameters sufficiently in longer time horizons.

Normalization and its implications

Study 1 shows that outcome variance has an effect on people's decision thresholds, whereas low variance leads to increased thresholds. However, the decision thresholds between the low- and high-variance conditions are practically identical on a normalized scale ($\mu = 0, \sigma = 1$), showing that people accept prices within the same percentile on each position, regardless of the variance. This result suggests that the price's value is assessed on its percentile rank within the sample rather than on its absolute number. Our finding is thus in line with the normalization hypothesis (Rangel & Clithero, 2012), which assumes that the value of an option is computed under a normalized code and corresponds to its relative position in the distribution of options. It is also consistent with the line of research suggesting that the value of a single option is calculated by comparing it to a sample of attribute values drawn from memory and its rank within the sample (decision by sampling; Stewart et al., 2006).

Consequently, the finding that the individual acceptance level is formulated on a percentile level (e.g., "I accept all prices in the first position that belong to the 12% lowest prices within the sampling distribution") may lead to distinct choice behavior predictions for other distributional environments, for example, when options are sampled from a skewed distribution. Indeed, previous studies could show that people's qualitative threshold adjustments are in line with the predictions of the normalization hypothesis (Baumann et al., 2020; Guan & Lee, 2018). That is, the same percentile rank ("I accept only prices that belong to the 12% lowest ones in the sample") maps to

a higher value in environments with only a few good options (left skewed) compared to plentiful environments (right skewed, assuming the same mean). Additionally, Lee et al. (Guan & Lee, 2018) have shown that thresholds that people use when searching for the maximum [when values $\sim \text{beta}(4, 2)$] can be modeled in the same way as the thresholds that people use when searching for the minimum [when values $\sim \text{beta}(2, 4)$], which further supports our findings. Future work should attempt to replicate our analysis in settings where options are sampled from different distributions and thus investigate if thresholds formulated on a percentile level represents a stable cognitive mechanism of the valuation process in sequential search tasks.

Our analysis is based on a variant of the optimal stopping task where the sampling distribution is known and stable across the sequences. Such a scenario is required when assuming that values defined on a percentile rank are successfully mapped to their absolute values. However, in unknown or unstable environments, the process implied by the normalization hypothesis could lead to strong biases. For example, in an environment with an unknown value distribution, one has to learn it during search. Calculating the rank of the current alternative within an initial, small (and plausibly unrepresentative) sample would lead to a high distortion in the acceptance threshold. For example, a decision maker who has a threshold to accept a price that belongs to the 10% best prices would accept too early in an initial unrepresentative sample of increased prices or too late if it contains many low prices. Indeed, there is evidence that people's search behavior in unfamiliar or unstable environments is highly sensitive to the history of samples during search (e.g., Brickman, 1972; Corbin et al., 1975; Lee & Courey, 2021; Shapira & Venezia, 1981). The models presented in this paper do not account for learning and thus are mute about how thresholds are formed. To understand the impact of unknown or unstable value distributions on participants' decision thresholds, one would need to extend cognitive models of optimal stopping search by incorporating a learning component, which we believe would be a useful direction for future research.

Adaptivity to time horizons

Study 2 shows that time horizon affects decision thresholds, whereby extended time results in lower first thresholds and lower adjustment rates across positions. In other words, participants are less selective at the beginning of the sequence in shorter time horizons and also adapt their acceptance rate more strongly across positions. This strategy is intuitively correct since shorter sequences imply lower chances along the way of a good option appearing. The qualitative adaptation of participants' initial thresholds and its adjustment over time is identical to the predictions of the LTM simulation about optimal parameter adaptation across time. However, with increasing time horizons, participants failed to adapt the parameters optimally, which led to a reduced search length and lower performance.

Given that the LTM parameters are more stable than the behavioral measures themselves, we addressed the question of how they relate to each other. Our result indicates that the initial aspiration level, set prior to search, is the main factor predicting participants' search length. Indicated by the negative linear correlation coefficient across the tasks in both studies, higher initial aspiration levels led to less search and vice versa. Importantly, the rate of the threshold's adjustment during search had no additional impact on search length, indicating that participants' modification of their aspiration level during search had little impact on how long they searched.

Furthermore, this result suggests that a too high aspiration level prior to search may serve as an explanation for the general finding of undersearch (Hey, 1982; Rapoport & Tversky, 1970; Schotter & Braunstein, 1981) and thus deserves special consideration.

The relationship between the initial threshold and performance is nonmonotonic, indicating that too extreme values (too high or too low) in the initial threshold led to worse performance, revealing that not only too little, but also too much search can be suboptimal. However, despite participants performing surprisingly well throughout all the tasks, we found that a longer sequence length (20 alternatives, Study 2) led to more pronounced undersearch for a larger number of participants (elongated longer tail in search length distribution) and thus led to worse performance. Yet some participants

achieved almost optimal performance across the tasks (Study 2) and thus adapted their decision thresholds perfectly to the time horizons. Taken together, these results show that participants differed not only in their aspiration levels but also in how well they adapted them across time horizons.

Our investigation further suggests that participants' initial thresholds and adjustment rates were adapted systematically to time horizons in a nonlinear way, indicating that they might have been following scaling regularities across time scales. Similarly, a study using a simplified explore–exploit task (Sang et al., 2020) has shown that participants' linearly decreasing thresholds declined in proportion to the length of the game. Indeed, the best performing linear thresholds predicted a proportional adaptation of the initial threshold and its adjustment rate and in this sense, participants' adaptation was optimal. Nevertheless, participants adapted too little in extended time horizons, leading to stronger deviation from optimality. Additional cognitive mechanisms that were not considered in our analysis, such as risk preference or search costs, could be reasons for the suboptimal adaptation to extended time scales, which we further discuss in subsequent sections.

Psychological reasons for maladaptation in extended time horizons

Risk preferences. Some studies have suggested that people use biased decision thresholds and thus search too little because they tend to be risk averse (Bhatia et al., 2021; Schotter & Braunstein, 1981; Schunk, 2009). Unlike the risk-neutral best performing LTM, human decision makers are often found to be risk averse (e.g., Pedroni et al., 2017), preferring the safe over the risky option. Given this assumption, suboptimal adaptation of higher initial thresholds to longer time horizons may reflect risk aversion that may become more visible with longer time horizons. However, there are conflicting findings regarding the impact of risk aversion on optimal stopping behavior. One study reported that there was no relationship between search length in an optimal stopping task and risk preferences elicited in a series of lottery tasks (Schunk, 2009). On the other side, a recent study about optimal sequential search with recall

found evidence that in this paradigm, risk aversion is one of the main factors guiding people's behavior (Bhatia et al., 2021). Moreover, Sonsino et al. (2002) suggested that the probability of choosing a given option decreases with the relative complexity of that alternative. In that sense, participants might perceive a higher time horizon as more risky or more complex and thus devalue the option of search more heavily, leading to behavior that deviates from optimality. However, it should be noted here that participants' choice behavior in Study 1 was not affected by an increase in variance, a manipulation that is often used as a measure of risk in the risky decision making literature (e.g., Weber et al., 2004); higher variance resulted in increased risk aversion.

Search costs. Alternatively, Seale and Rapoport (2000) hypothesized that stopping too early in the optimal stopping task could result from the existence of endogenous search costs, that is, psychological costs arising from considering and evaluating options (see also Bhatia et al., 2021). Although our experiment used a somewhat different design, participants may have been motivated to end the experiment as quickly as possible and thus considered the time spent observing and reviewing the prices as a kind of search cost. This idea is further supported by the literature on the accuracy–effort trade-off reported by Payne et al. (1993) and could also explain why search was more strongly reduced with longer time horizons. In our study, search costs may have been integrated into the evaluation of the decision thresholds, whereby extended sequences would accumulate higher search costs (more steps to go through all the alternatives) and thus exert a greater impact on decision thresholds.

Misrepresentation of underlying sampling distribution. A simple explanation for the weak adjustment of the threshold across positions (relative to the adjustment of the best performing linear thresholds) in extended time horizons might be that it results from initial thresholds being higher. Alternatively, the origins of maladaptation of the thresholds within a sequence might be that participants misrepresented the underlying outcome distribution. In a sequential search paradigm such as the optimal stopping task, the outcome distribution changes dynamically and thus has to be recalculated at each time step. Since people seem to lack the ability to infer many

characteristics of aggregated outcome distributions correctly (Benartzi & Thaler, 1999; Klos et al., 2005), simpler strategies are used solve the task (see also Speekenbrink & Konstantinidis, 2015; Wilson et al., 2014, for a similar discussion of bandit tasks). Indeed, results from the balloon analogue risk task (a nonstationary environment) revealed that the best fitting model to two data sets assumes a stationary representation (Bishara et al., 2009; Wallsten et al., 2005), showing that people simplified the balloons' exploding behavior. Future work is needed to incorporate such mental representations and updating mechanism into optimal stopping models to explain "how" participants arrive at the lower thresholds.

The role of choice sensitivity. Participants' response sensitivity is highly predictive of their performance, with higher values (and thus more determinism) leading to better performance. Indeed, high-performing participants' choice curves are more consistent compared to those of low-performing participants, showing less deviance from the intended choice (see Appendix Figure 13). However, we found no relationship between participants' choice sensitivity and their search length. This result could shed light on the complex relationship between risk preferences and search behavior. On one side, studies relating risk preferences elicited in lottery tasks with search length in optimal stopping tasks have not found any correlation (Schunk, 2009). On the other side, studies on related sequential risk-taking tasks, such as the the angling risk tasks or the balloon analogue risk task, reported sensitivity to the outcome evaluation as being predictive for individual differences of real-world risk takers (Pleskac, 2008; Wallsten et al., 2005). Exploring the relationship between choice sensitivity in search tasks and risk preferences would be a further useful direction for future research.

LTM versus psychologically motivated optimal models

As mentioned in the Introduction, we think that the main advantages of the LTM are (a) its plausible cognitive process assumptions, (b) that it captures linearly increasing thresholds as observed across different contexts, and (c) that it offers psychological interpretations of its parameters. The model fits in both of our studies support the

assumptions of the LTM and thus provide further evidence of the LTM's robustness in describing choice behavior. The BOM, an optimal model with systematic biases, was less capable of accommodating people's choices. Interestingly, a recent study introduced a model that assumes decision thresholds are based on the optimal thresholds but explains deviations thereof by psychological factors such as risk preference and effort search costs (Bhatia et al., 2021). Such models represent an important source for quantifying and understanding deviations from optimality especially in simpler search task with tractable optimal solutions. However, there are so far no direct comparisons between the LTM and the psychologically motivated optimal model and therefore it is a matter of debate whether one model is better than the other. An additional important point when deciding on the models' validity is to investigate the interpretability of both models' parameter values. Indeed, previous studies have identified individual characteristics that could account for the variation in optimal stopping behavior, such as intelligence (Burns et al., 2006), working memory capacity (people with higher capacity tend to explore more; Hills & Pachur, 2012), age (older people may adapt less; Mata et al., 2013; Rydzewska et al., 2018), or depressive symptoms (higher depression leads to more search; von Helversen et al., 2011). Therefore, future studies should design studies that question if individual differences in the models' parameters are related to more traditional measures such as problem solving, decision making biases or IQ.

Limits of Generality

In this section we address the question of our results' generality (Simons et al., 2017). Our research focused on an optimal stopping task in which people knew the underlying value distribution. For this reason, we included a learning phase in which participants were exposed to samples from the respective distribution. This framework ensured we could measure adaptive behavior in response to context in optimal stopping tasks by minimizing additional processes such as learning or adaptation to changing environments. In such problems we find that human behavior is well described by an LTM, which has been replicated in other studies despite variations in the testing

context (Baumann et al., 2020; Lee & Courey, 2021; Sang et al., 2020; Song et al., 2019), and even in environments in which the value distribution was not explicitly learned (Baumann et al., 2020, Study 3). Therefore we expect that our findings are reproducible in the full-information version of the optimal stopping task with stable environments. However, unknown or unstable environments require participants to learn the value distribution during search. Such environments require behavioral models with an account of learning, which is not integrated in our models.

The task in our studies used an airline-ticket-purchasing task as a cover story and thus one could argue that participants' intuition may have been shaped by the expectation that prices increase when approaching the end of search. On the other hand, searching for online tickets represents a natural decision setting that has no clear implications on the trend of ticket prices. Moreover, the extensive training before the actual decision phase and additional test trials ensured that participants were aware that the price distribution was not changing across the sequence. Furthermore, we did not find any learning across trials, indicating that people's belief about the stationarity of the ticket distribution was stable throughout the experiment. Therefore, we think that our results are likely to generalize to optimal stopping tasks that use distinct cover stories.

Context of Our Research

Overall, our work belongs in a stream of research that uses cognitive modeling to identify the choice strategies and environmental factors that influence search behavior in optimal stopping tasks (Baumann et al., 2020; Bhatia et al., 2021; Guan et al., 2020; Guan & Lee, 2018; Lee, 2006) and sequential decision making tasks more broadly (e.g., Reutskaja et al., 2011; Solway & Botvinick, 2015; Song et al., 2019; Wallsten et al., 2005). Furthermore, our work aims to uncover underlying processes that capture stable individual differences in sequential decisions that go beyond behavioral measures and thus contributes to the research area investigating a general mechanism of search across different domains (Hills et al., 2008; Hills et al., 2015; Mata & von Helversen, 2015; Pirolli, 2007).

In addition, the work is part of a current research project with the goal of clarifying the relationship between decisions in sequential search tasks and risky decisions. Sequential decisions can be seen as risky choices that require one to decide between a safe (accept) and an uncertain (reject) alternative. However, attempts to link risk preferences measured in lottery tasks to search behavior in sequential decisions have been unsuccessful (e.g., Frey et al., 2017; Pedroni et al., 2017). Our studies have repeatedly found that in optimal stopping tasks risk preference appears to depend on position: That is, people make risk-averse choices in the beginning of the sequence but turn to risk-seeking behavior on the last positions. In a related project, we are currently investigating how the characteristics of optimal stopping tasks—for example, the sequential presentation of options—affect behavior beyond stable risk preferences to gain a more comprehensive understanding of the relationship between risk preferences and choices in sequential decisions.

Conclusion

We inquired into whether and how context task variables affect optimal choice behavior in optimal stopping tasks by taking into account the constraints imposed by participants' cognitive strategies. Results show that people adapt their decision thresholds to variance and time horizon as described by the choice strategy, but more pronounced deviations from optimality suggest that additional variables may affect human search behavior. By identifying and quantifying the cognitive processes involved in participants' adaptation strategies, we show that the adaptation to variance leads to identical decision thresholds on a percentile scale, suggesting that the value of an option is determined by the rank within the pool of alternatives. Extended time horizons affect both people's initial threshold level and its adjustment across search, both contributing to less search and weaker performance. This work contributes to the understanding of the adaptive processes that underlie sequential decisions and points to the importance of taking into account the fit between the cognitive strategy and the environmental condition when interpreting human behavior.

References

- Alberini, A., Cropper, M., Krupnick, A., & Simon, N. B. (2004). Does the value of a statistical life vary with age and health status? evidence from the us and canada. *Journal of Environmental Economics and Management*, *48*(1), 769–792.
- Baumann, C., Schlegelmilch, R., & von Helversen, B. (2022, April 17). Adaptive behavior in optimal sequential search. Retrieved from *osf.io/9rct5*.
- Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *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.
- Benartzi, S., & Thaler, R. H. (1999). Risk aversion or myopia? choices in repeated gambles and retirement investments. *Management Science*, *45*(3), 364–381.
- Bendor, J., Mookherjee, D., & Ray, D. (2001). Aspiration-based reinforcement learning in repeated interaction games: An overview. *International Game Theory Review*, *3*(2-3), 159–174.
- Bhatia, S., He, L., Zhao, W. J., & Analytis, P. P. (2021). Cognitive models of optimal sequential search with recall. *Cognition*, *210*, Article 104595.
- Bishara, A. J., Pleskac, T. J., Fridberg, D. J., Yechiam, E., Lucas, J., Busemeyer, J. R., Finn, P. R., & Stout, J. C. (2009). Similar processes despite divergent behavior in two commonly used measures of risky decision making. *Journal of Behavioral Decision Making*, *22*(4), 435–454.
- Bossaerts, P., & Murawski, C. (2017). Computational complexity and human decision-making. *Trends in Cognitive Sciences*, *21*(12), 917–929.
- Brickman, P. (1972). Optional stopping on ascending and descending series. *Organizational Behavior and Human Performance*, *7*(1), 53–62.

- Burns, N. R., Lee, M., & Vickers, D. (2006). Are individual differences in performance on perceptual and cognitive optimization problems determined by general intelligence? *The Journal of Problem Solving*, *1*(1), Article 3.
- Cooper, G. F. (1990). The computational complexity of probabilistic inference using bayesian belief networks. *Artificial Intelligence*, *42*(2-3), 393–405.
- Corbin, R. M., Olson, C. L., & Abbondanza, M. (1975). Context effects in optional stopping decisions. *Organizational Behavior and Human Performance*, *14*(2), 207–216.
- Cox, J. C., & Oaxaca, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, *2*(3), 301–329.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4*(3), 282–289.
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, *3*(10), Article e1701381.
- Genest, W., Stauffer, W. R., & Schultz, W. (2016). Utility functions predict variance and skewness risk preferences in monkeys. *Proceedings of the National Academy of Sciences of the United States*, *113*(30), 8402–8407.
- Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482.
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, *61*(313), 35–73.
- 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), Article 1.

- Guan, M., Stokes, R., Vandekerckhove, J., & Lee, M. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision Making, 15*, 823–850.
- Guan, M., & Lee, M. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision, 5*(4), 339–361.
- Guan, M., Lee, M., & 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 meeting of the cognitive science society* (pp. 553–558).
- Guan, M., Lee, M., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. In D. Noelle & R. Dale (Eds.). *Cognitive Science Society*.
- Hey, J. D. (1982). Search for rules for search. *Journal of Economic Behavior & Organization, 3*(1), 65–81.
- Hills, T. T., & Pachur, T. (2012). Dynamic search and working memory in social recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38*(1), 218–228.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science, 19*(8), 802–808.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & Group, C. S. R. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences, 19*(1), 46–54.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review, 92*(5), 1644–1655.
- Huys, Q. J., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., Dayan, P., & Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences of the United States, 112*(10), 3098–3103.

- Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception and Psychophysics*, *2*(8), 374–376.
- Klos, A., Weber, E. U., & Weber, M. (2005). Investment decisions and time horizon: Risk perception and risk behavior in repeated gambles. *Management Science*, *51*(12), 1777–1790.
- Lee, M. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26.
- Lee, M., & 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., 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).
- Markowitz, H. (1959). *Portfolio selection: Efficient diversification of investments*. John Wiley.
- Mata, R., & von Helversen, B. (2015). Search and the aging mind: The promise and limits of the cognitive control hypothesis of age differences in search. *Topics in Cognitive Science*, *7*(3), 416–427.
- Mata, R., Wilke, A., & Czienskowski, U. (2013). Foraging across the life span: Is there a reduction in exploration with aging? *Frontiers in Neuroscience*, *7*, Article 53.
- Morey, R., Rouder, J., Jamil, T., Urbanek, S., Forner, K., & Ly, A. (2018). *BayesFactor: Computation of Bayes Factors for common designs* [R package version 0.9.12-4.2]. <https://cran.r-project.org/web/packages/BayesFactor/index.html>
- Nieuwenhuis, S., Heslenfeld, D. J., von Geusau, N. J. A., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *Neuroimage*, *25*(4), 1302–1309.
- Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, *55*(1), 84–93.

- Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *Journal of Neuroscience*, *29*(44), 14004–14014.
- Payne, J. W., Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, *1*(11), 803–809.
- Pirolli, P. (2007). *Information foraging theory: Adaptive interaction with information*. Oxford University Press.
- Pleskac, T. J. (2008). Decision making and learning while taking sequential risks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(1), 167.
- Plummer, M. (2003). Jags: A program for analysis of Bayesian graphical models using gibbs sampling. *Proceedings of the 3rd international workshop on distributed statistical computing*, *124*(125.10), 1–10.
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, *41*.
- Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: Theory and evidence. *Current Opinion in Neurobiology*, *22*(6), 970–981.
- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Decision Processes*, *5*(2), 105–120.
- Reutskaja, E., Nagel, R., Camerer, C. F., & Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, *101*(2), 900–926.
- Rigoli, F., Friston, K. J., & Dolan, R. J. (2016). Neural processes mediating contextual influences on human choice behaviour. *Nature Communications*, *7*(1), 1–11.
- Rydzewska, K., von Helversen, B., Kossowska, M., Magnuski, M., & Sedek, G. (2018). Age-related within-task adaptations in sequential decision making: Considering cognitive and motivational factors. *Psychology and aging*, *33*(2), 297.

- Sang, K., Todd, P. M., Goldstone, R. L., & Hills, T. T. (2020). Simple threshold rules solve explore/exploit trade-offs in a resource accumulation search task. *Cognitive Science*, *44*(2), Article e12817.
- Schotter, A., & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, *19*(1), 1–25.
- Schunk, D. (2009). Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence. *Journal of Economic Dynamics and Control*, *33*(9), 1719–1738.
- 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*(4), 391–411.
- Shapira, Z., & Venezia, I. (1981). Optional stopping on nonstationary series. *Organizational Behavior and Human Performance*, *27*(1), 32–49.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.
- Simons, D. J., Shoda, Y., & Lindsay, D. S. (2017). Constraints on generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, *12*(6), 1123–1128.
- Solway, A., & Botvinick, M. M. (2015). Evidence integration in model-based tree search. *Proceedings of the National Academy of Sciences of the United States*, *112*(37), 11708–11713.
- Song, M., Bnaya, Z., & Ma, W. J. (2019). Sources of suboptimality in a minimalistic explore–exploit task. *Nature Human Behaviour*, *3*(4), 361–368.
- Sonsino, D., Benzion, U., & Mador, G. (2002). The complexity effects on choice with uncertainty—experimental evidence. *The Economic Journal*, *112*(482), 936–965.

- Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2), 351–367.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the royal statistical society: Series b (statistical methodology)*, 64(4), 583–639.
- Stewart, N. (2009). EPS prize lecture: Decision by sampling: The role of the decision environment in risky choice. *Quarterly Journal of Experimental Psychology*, 62(6), 1041–1062.
- Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive psychology*, 53(1), 1–26.
- Todd, P. M., Gigerenzer, G., & the ABC Research Group. (2012). *Ecological rationality: Intelligence in the world*. Oxford University Press.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, 39(10), 1179–1189.
- von Helversen, B., Mata, R., Samanez-Larkin, G. R., & Wilke, A. (2018). Foraging, exploration, or search? On the (lack of) convergent validity between three behavioral paradigms. *Evolutionary Behavioral Sciences*, 12(3), 152–162.
- 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. *Journal of Abnormal Psychology*, 120(4), 962–968.
- Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological Review*, 112(4), 862–880.
- Weber, E., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111, 430–445.
- White, T. L., Lejuez, C. W., & de Wit, H. (2008). Test-retest characteristics of the balloon analogue risk task (bart). *Experimental and Clinical Psychopharmacology*, 16(6), 565–570.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma.

Journal of Experimental Psychology: General, *143*(6), 2074–2081.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, a. V. (2003). Consumer

Sequential Search: Not Enough or Too Much? *Marketing Science*, *22*(4), 503–519.

Appendix: Adaptive Behavior in Optimal Sequential Search

Text

Calculation of optimal thresholds

We describe the calculation of optimal thresholds applied to our scenario, where payoff is proportional to the chosen value and the goal is to find the cheapest ticket price. We first derive the optimal solution mathematically based on the paper of Gilbert and Mosteller (Gilbert & Mosteller, 1966, Section 5b) and further provide a more intuitive explanation.

Let us assume a sequence of n ticket prices that are drawn from a standard normal distribution with density

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}} \quad (6)$$

and the goal is to find the lowest ticket price in this sequence.

The *optimal strategy* for this task is as follows: If $n = 1$, the decision maker is forced to accept the ticket. The threshold on the last position (T_1) is set to ∞ and ticket prices below this threshold are accepted:

$$T_1 = \infty \quad (7)$$

Therefore the expected price of the last ticket (P_1) is the mean (μ) of the distribution.

If $n = 2$, the decision maker decides to keep the first option or to reject it and to go on to the second one. If they go on, their expected ticket price is μ . They keep the current one, x , if $x < \mu$, reject it if $x > \mu$, and are indifferent if $x = \mu$. Thus, the expected price of the last ticket (P_1) is also the threshold for the second to last option:

$$T_2 = P_1 \quad (8)$$

Then for $n = 2$, the decision maker's expected price (P_2) is

$$P_2 = \int_{-\infty}^{P_1} f(x) \cdot x dx + P_1 \cdot \int_{P_1}^{\infty} f(x) dx \quad (9)$$

The remaining terms of the sequence can be computed in a recursive manner. For each n , the decision maker accepts the ticket if its price is lower than the expected price of the remaining $n - 1$ tickets ($x < P_{n-1}$) but rejects it if the ticket's price is higher than the remaining expected prices ($x > P_{n-1}$). Therefore the threshold on the n th position (T_n) is

$$T_n = P_{n-1} \quad (10)$$

Accordingly the expected price (P_n) is

$$P_n = \int_{-\infty}^{P_{n-1}} f(x) \cdot x dx + P_{n-1} \cdot \int_{P_{n-1}}^{\infty} f(x) dx \quad (11)$$

Intuitive explanation

The optimal thresholds T_n for maximizing the payoff are calculated by working backward from the last ticket price: The threshold of the final item (T_1) is ∞ , because the rules of the task stipulate that the final item must be accepted if no earlier item has been chosen. The thresholds for the previous items are determined by working backward from the final item, using conditional expectations. First, we calculate the expected value of the final item (P_1). For the last item, this is the expectation of the

overall probability distribution from which the options are sampled. Therefore, to maximize expected reward on the second to last position, one's policy should be to accept a particular option if it is better (in our case smaller) than the expected reward if one continues under the optimal policy. The second to last item should be accepted if its value is smaller than the expected value of the final item. This means that the threshold of the second to last item (T_2) is the expected value of the last item (P_1). The expected value of the second to last item (P_2) is the expected value of the part of the probability distribution that is better (in our case smaller) than the threshold (T_2) for the second to last item. The probability of this expected value is the area under the probability distribution that is better than this threshold. The overall expected reward at the second to last position [P_2 ; and therefore the threshold for the third to last item (T_3)] is calculated as follows: We multiply the expected value for the second to last item with its probability plus the expected value of the last item multiplied with its probability (which is equal to 1 minus the probability of the second to last item). The remaining thresholds are calculated in the same way.

Method

- *Linear Threshold Model (LTM):*

$$\beta_j \sim N(\mu_g^\beta, \sigma_g^\beta)$$

$$\delta_j \sim N(\mu_g^\delta, \sigma_g^\delta)$$

$$t_{0j} \sim N(\mu_g^{t_0}, \sigma_g^{t_0})$$

$$\mu_g^\beta \sim N(0, 100), \sigma_g^\beta \sim U(0.1, 10)$$

$$\mu_g^\delta \sim N(0, 100), \sigma_g^\delta \sim U(0.1, 10)$$

$$\mu_g^{t_0} \sim U(100, 200), \sigma_g^{t_0} \sim U(0.1, 10)$$

- *Biased Optimal Model (BOM):*

$$\beta_j \sim N(\mu_g^\beta, \sigma_g^\beta)$$

$$\gamma_j \sim N(\mu_g^\gamma, \sigma_g^\gamma)$$

$$\alpha \sim N(\mu_g^\alpha, \sigma_g^\alpha)$$

$$\mu_g^\beta \sim N(0, 100), \sigma_g^\beta \sim U(0.1, 10)$$

$$\mu_g^\gamma \sim N(0, 100), \sigma_g^\gamma \sim U(0.1, 10)$$

$$\mu_g^\alpha \sim N(0, 100), \sigma_g^\alpha \sim U(0.1, 10)$$

- *Independent Threshold Model (ITM):*

$$\beta_j \sim N(\mu_g^\beta, \sigma_g^\beta)$$

$$t(p)_j \sim N(\mu_g^{t(p)}, \sigma_g^{t(p)})$$

$$\mu_g^\beta \sim N(0, 100), \sigma_g^\beta \sim U(0.1, 10)$$

$$\mu_g^{t(p)} \sim U(0, 220), \sigma_g^\beta \sim U(0.1, 20)$$




Figure A1. Biased optimal model: Thresholds for different parameter values of γ and α .




Figure A2. Biased optimal model (BOM): Probability of acceptance across positions, data versus BOM's posterior predictions. Empirical data appear as solid lines; black circles and bars represent the posterior predictive means and 95% confidence intervals. The different lines represent prices ranging from the first quantile (Q1) to the sixth (Q6). Q1: Prices in the first quantile; Q2: prices between the first and second quantile, etc.




Figure A3. Distribution of posterior individual-level mean slope parameters for both variance conditions.

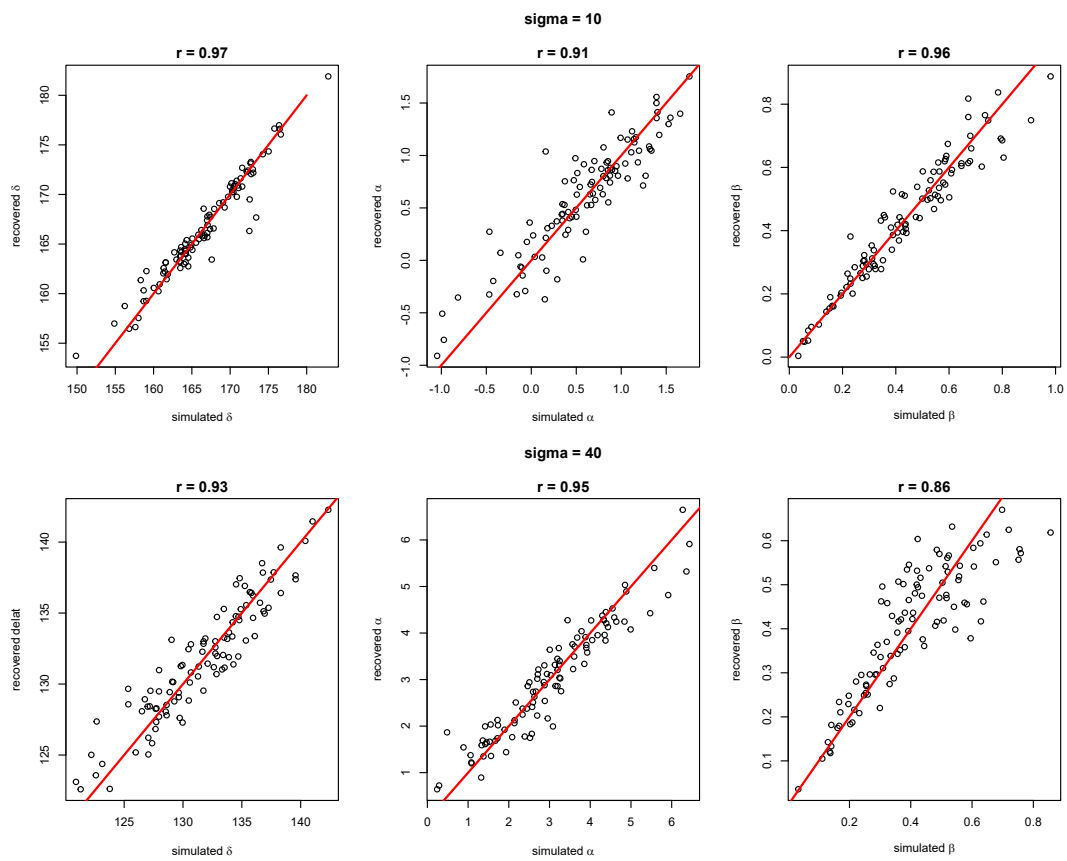


Figure A4. Study 1: Parameter recovery for the linear threshold model's parameters for both variance conditions ($\sigma = 10$ and $\sigma = 40$). Parameter values in the range of estimated values from model fitting. Each point corresponds to one of 100 trials with 100 sequences.

Table A1
Study 1: Correlations between parameters

	θ_{10}	θ_{40}	δ_{10}	δ_{40}	β_{10}
θ_{10}	1				
θ_{40}	0.62 [0.59, 0.65]	1			
δ_{10}	-0.35 [-0.47, -0.23]	-0.26 [-0.38, -0.13]	1		
δ_{40}	-0.06 [-0.12, 0.07]	-0.34 [-0.45, -0.21]	0.53 [0.41, 0.61]	1	
β_{10}	0.08 [-0.05, 0.21]	0.01 [-0.13, 0.15]	0.08 [-0.05, 0.21]	-0.13 [-0.26, 0.00]	1
β_{40}	0.23 [0.11, 0.36]	0.19 [0.06, 0.32]	0.03 [-0.10, 0.17]	0.19 [0.05, 0.32]	0.56 [0.45, 0.64]

Correlations between parameters with 95% confidence intervals (square brackets; bold numbers: Bayes factor > 150). Parameter's subscript corresponds to the variance condition (10 = SD10, 40 = SD40). There is very strong evidence that the cognitive parameters are moderately to strongly correlated between the two conditions. θ_0 , reflecting participants' aspiration level before search, reveals the strongest correlation ($\rho = .62$), pointing to a stable cognitive process across the two tasks. Note that the first threshold (θ_0) and the adjustment rate across search (δ) are negatively correlated to some extent, showing a typical parameter dependency as in all regression-like (intercept plus slope) models. That is, if the initial threshold is very high in the first place, then the ceiling is almost reached and there is no point in further increases, and starting with a very low threshold, naturally, allows for the strongest increase. Importantly, all other correlations are close to 0, suggesting an independent contribution to predictions.




Figure A5. Panels A1–A3: Scatterplots of performance (as percentages, points/maximum points) and parameters (θ_0 , δ , and β). Panels B1–B3: search length (as percentages, search length/total sequence length) and parameters. Black: low variance ($SD = 10$); green: high variance ($SD = 40$). Red cross: best linear threshold parameter.




Figure A6. Optimal thresholds for each time horizon ($N=5, 10, 20$). Red dashed line: Identical thresholds as the remaining number of options is the same. Optimal threshold depends only on the remaining options, therefore the initial decision threshold when $N=5$ is equal to the threshold on Position 6 when $N=10$ and to the threshold on position 16 when $N=20$.

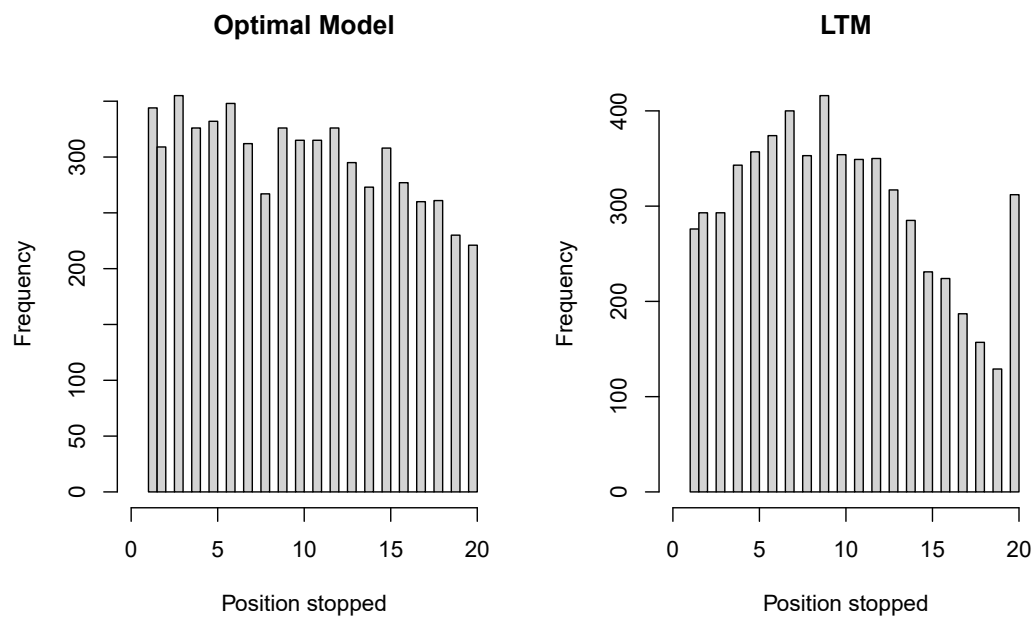


Figure A7. $N = 20$: Distribution of the stopping positions for 66 decision makers for 60 trials using the optimal model versus the linear threshold model (LTM).




Figure A8. $N = 5$: Probability of acceptance across positions: data versus the linear threshold model's posterior predictions. Empirical data appear in blue; black lines and bars represent the posterior predictive means and the 95% confidence intervals. The different lines represent prices ranging from the first quantile (Q1) to the sixth (Q6). Q1: Prices in the first quantile; Q2: prices between the first and second quantile, etc.




Figure A9. $N = 10$: Probability of acceptance across positions; data versus the linear threshold model's posterior predictions. Empirical data appear in blue; black lines and bars represent the posterior predictive means and the 95% confidence intervals. The different lines represent prices ranging from the first quantile (Q1) to the fifth (Q5). Q1: Prices in the first quantile; Q2: prices between the first and second quantile, etc.




Figure A10. $N = 20$: Probability of acceptance across positions; data versus the linear threshold model's posterior predictions. Empirical data appear in blue; black lines and bars represent the posterior predictive means and the 95% confidence intervals. The different lines represent prices ranging from the first quantile (Q1) to the fourth (Q4). Q1: Prices in the first quantile; Q2: prices between the first and second quantile, etc.

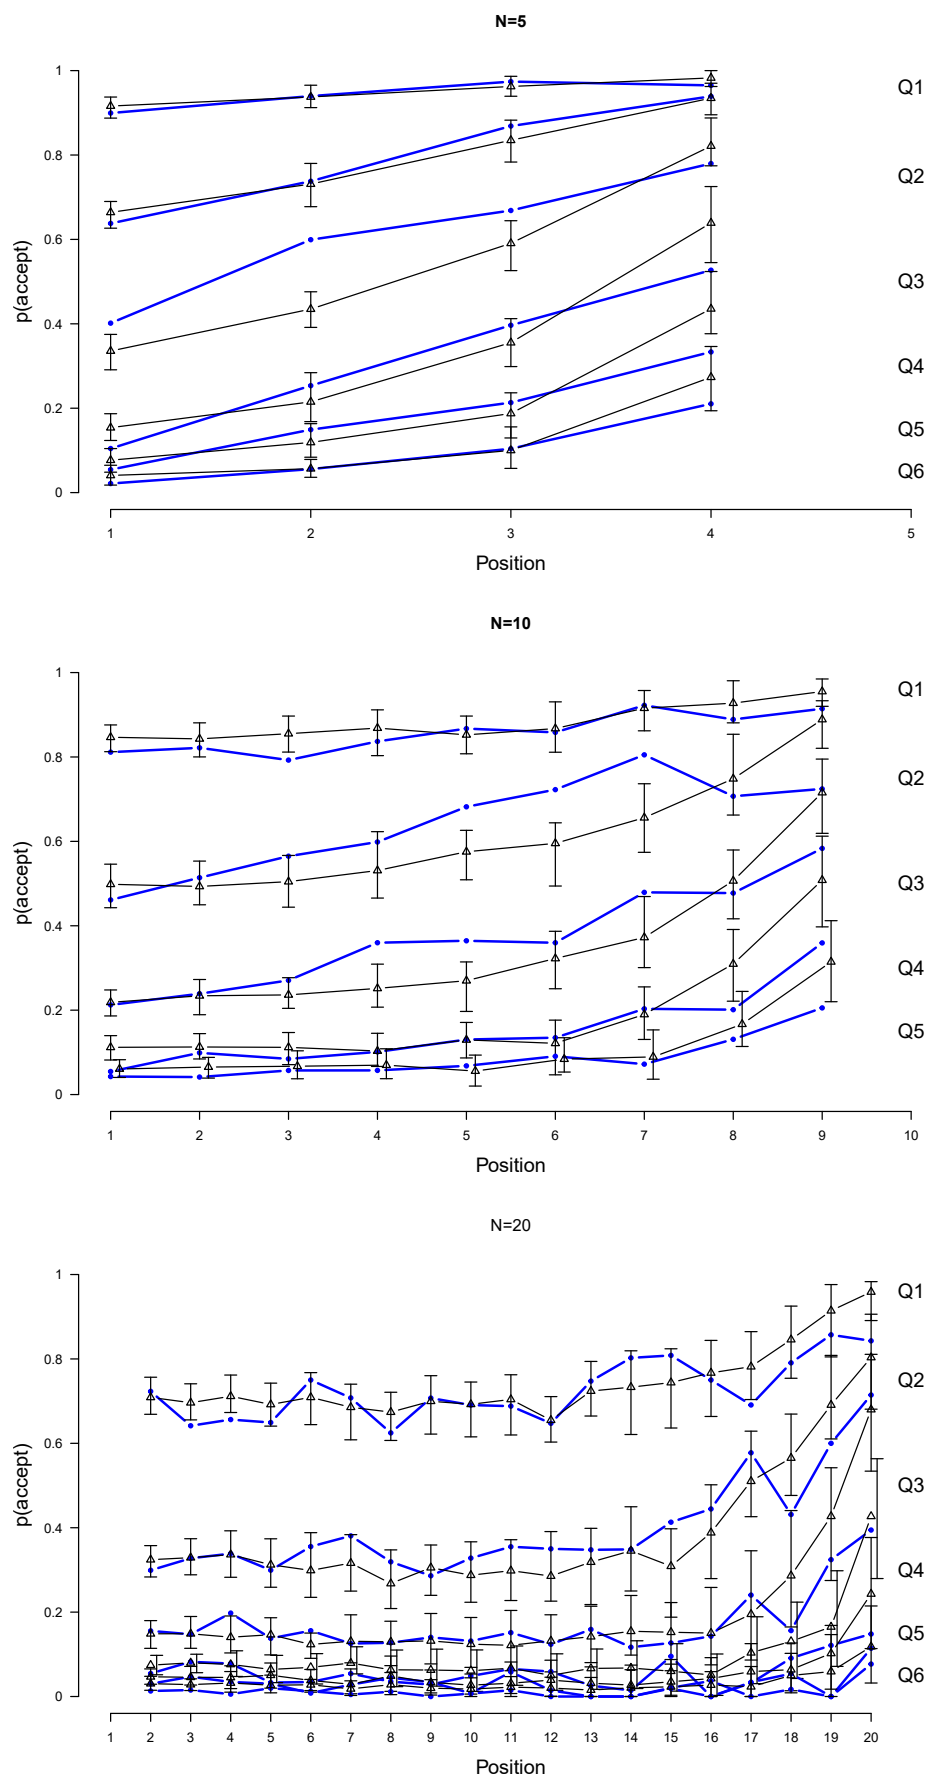


Figure A11. From top to bottom: $N = 5$, $N = 10$, and $N = 20$: Probability of acceptance across positions; data versus the biased optimal model's posterior predictions. Empirical data appear in blue; black lines and bars represent the posterior predictive means and the 95% confidence intervals. The different lines represent prices ranging from the first quantile (Q1) to the fourth (Q4). Q1: Prices in first quantile; Q2: prices between the first and second quantile, etc.




Figure A12. Parameter recovery for the linear threshold model when $N = 5$, $N = 10$, and $N = 20$. Parameter values in the range of estimated values from model fitting. Each point corresponds to one of 2,000 trials with 60 sequences.

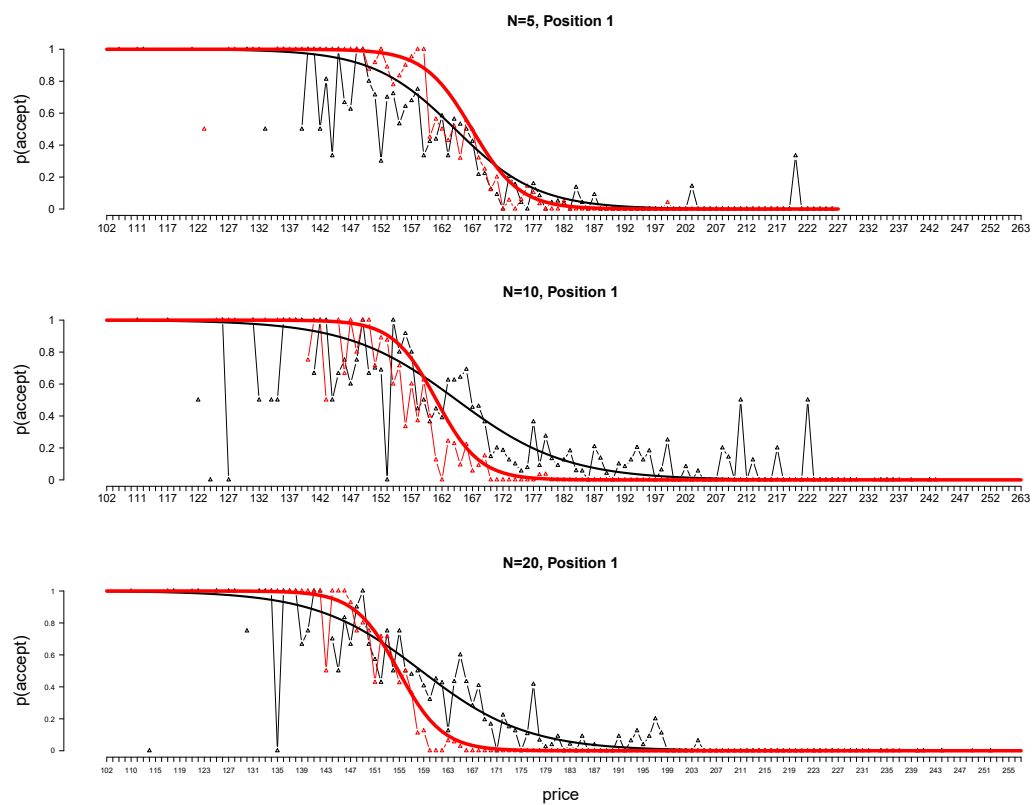


Figure A13. Choice curves split into participants that belong to the 20% highest (red lines) and the 20% lowest performing group (black lines), for each search length and on the first position. Well-performing participants exhibit higher determinism in choices than low-performing participants. Solid lines are logistic curves fitted to the data (red line: fitted to data of participants in the high-performing group; black line: fitted to data of participants in the low-performing group.)

Table A2
Study 2: Correlations between LTM parameters

θ_0^5	θ_0^{10}	θ_0^{20}	δ^5	δ^{10}	δ^{20}	β^5	β^{10}	
θ_0^5	1							
θ_0^{10}	0.80 [0.69, 0.87]	1						
θ_0^{20}	0.61 [0.45, 0.74]	0.73 [0.61, 0.83]	1					
δ^5	-0.25 [-0.4, -0.2]	-0.12 [-0.3, -0.1]	-0.16 [-0.4, 0.06]	1				
δ^{10}	-0.2 [-0.4, 0.02]	-0.14 [-0.34, 0.07]	-0.19 [-0.4, 0.02]	0.51 [0.32, 0.67]	1			
δ^{20}	-0.2 [-0.41, 0.01]	-0.11 [-0.32, 0.11]	-0.15 [-0.36, 0.07]	0.31 [0.1, 0.5]	0.27 [0.04, 0.47]	1		
β^5	0.00 [-0.2, 0.2]	-0.15 [0.36, 0.07]	-0.08 [-0.3, 0.14]	-0.02 [-0.4, -0.02]	-0.12 [-0.3, 0.09]	0.00 [-0.23, 0.22]	1	
β^{10}	0.00 [-0.2, 0.2]	-0.05 [-0.27, 0.17]	-0.05 [-0.28, 0.17]	-0.13 [-0.35, -0.1]	0.00 [-0.22, 0.2]	-0.05 [-0.27, 0.18]	0.62 [0.45, 0.74]	1
β^{20}	0.00 [-0.2, 0.2]	-0.05 [-0.2, 0.1]	0.07 [-0.15, 0.3]	-0.06 [-0.28, 0.16]	-0.1 [-0.3, 0.12]	-0.05 [-0.28, 0.16]	0.66 [0.51, 0.77]	0.74 [0.61, 0.83]

Correlations between linear threshold model parameters when $N = 5, 10,$ and $20,$ with 95% confidence intervals (bold numbers: Bayes factor > 150)

Table A3

Study 2: $N=5$: Model comparison for the analysis of search length using Bayes Factor

Model	B_{m0}	B_{mf}
$\theta + \delta + \beta$	1.9×10^{25}	1
$\theta + \beta$	1.6×10^{25}	0.8
$\theta + \delta$	3×10^{20}	1.5×10^{-5}
$\theta + \beta$	1.02×10^{20}	5×10^{-6}
θ	0.64	3.3×10^{-26}
β	0.26	1.3×10^{-26}

$N = 5$: Ranking of the different models for the analysis of the search length using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A4

Study 2: $N = 5$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	2.97[2.94, 3.01]		
θ	-0.08 [-0.09, -0.07]	-0.87 [-0.92, -0.79]	.76
δ	-0.03 [-0.06, -0.00]	0.16 [-0.07, 0.4]	.03
β	1.36 [0.88, 0.82]	0.17 [-0.06, 0.4]	.03

Analysing search length: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A5

Study 2: $N=5$: Model comparison for the analysis of performance using Bayes Factor

Model	B_{m0}	B_{mf}
β	1.9×10^4	1
$\theta + \theta^2$	6.4×10^3	0.32
$\theta + \beta$	4.6×10^3	0.23
$\delta + \beta$	4.1×10^3	0.21
$\theta + \theta^2 + \beta$	2.3×10^3	0.12
$\theta + \beta + \delta$	1.6×10^3	0.08

$N = 5$: Ranking of the different models for the analysis of the performance using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A6

Study 2: $N = 5$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	27[26.7, 27.3]		
β	0.7 [-0.42, 0.98]	0.53 [0.35, 0.67]	.28

Analysing performance: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A7

Study 2: $N=10$: Model comparison for the analysis of search length using Bayes Factor

Model	B_{m0}	B_{mf}
$\theta + \beta$	4.2×10^{28}	1
$\theta + \beta + \delta$	5.2×10^{27}	0.13
θ	1.4×10^{24}	3.3×10^{-5}
$\theta + \delta$	1.4×10^{23}	3.3×10^{-6}
β	1.46	3.5×10^{-29}
β	1.01	2.4×10^{-29}

$N = 10$: Ranking of the different models for the analysis of the search length using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A8

Study 2: $N = 10$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	4.7[4.6, 4.81]		
θ	-0.14 [-0.16, -0.13]	-0.90 [-0.93, -0.85]	.81
β	3.06 [2.0, 4.1]	0.22 [0.00, 0.43]	.05

Analysing search length: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A9

Study 2: $N=10$: Model comparison for the analysis of performance using Bayes Factor

Model	B_{m0}	B_{mf}
$\theta + \theta^2 + \beta$	4.0×10^{16}	1
$\theta + \theta^2 + \delta + \beta$	5.4×10^{15}	0.13
$\theta^2 + \beta$	1.1×10^{15}	0.03
$\theta^2 + \beta + \delta$	1.7×10^{14}	0.004
$\theta + \beta$	8.4×10^{12}	2.0×10^{-4}
$\theta + \delta + \beta$	1.2×10^{12}	2.6×10^{-5}

$N = 10$: Ranking of the different models for the analysis of the performance using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A10

Study 2: $N = 10$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	27.5 [27.2, 27.8]		
θ	-0.55 [-0.88, -0.24]	-0.30 [-0.50, -0.07]	.09
θ^2	-0.54 [-0.76, -0.32]	-0.67 [-0.79, -0.53]	.45
β	1.36 [0.99, 1.7]	0.71 [0.58, 0.81]	.50

Analysing performance: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A11

Study 2: $N=20$: Model comparison for the analysis of search length using Bayes Factor

Model	B_{m0}	B_{mf}
$\theta + \beta$	9.1×10^{18}	1
$\theta + \beta + \delta$	9.5×10^{17}	0.10
θ	6.1×10^{15}	6.7×10^{-4}
$\theta + \delta$	6.0×10^{14}	6.5×10^{-5}
β	0.5	5.11×10^{-20}
β	0.3	3.2×10^{-20}

$N = 20$: Ranking of the different models for the analysis of the search length using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A12

Study 2: $N = 20$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	7.4[7.11, 7.68]		
θ	-0.25 [-0.28, -0.21]	-0.81 [-0.88, -0.72]	.66
β	8.9 [5.2, 12.7]	0.06 [-0.16, 0.29]	.004

Analysing search length: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A13

Study 2: $N=20$: Model comparison for the analysis of performance using Bayes Factor

Model	B_{m0}	B_{mf}
$\theta + \theta^2 + \beta$	1.4×10^{20}	1
$\theta + \theta^2 + \delta + \beta$	3.1×10^{19}	0.23
$\theta^2 + \beta$	3.3×10^{18}	0.02
$\theta^2 + \beta + \delta$	1.6×10^{18}	0.01
$\theta + \beta$	1.3×10^{11}	9.3×10^{10}
$\theta + \beta + \beta$	2.2×10^{10}	1.6×10^{-10}

$N = 20$: Ranking of the different models for the analysis of the performance using Bayes factors. Selected model is in bold. The six models with the highest support are presented. B_{m0} : Bayes factor for row model against the Null model. B_{mf} : compares the row model against the model with the highest Bayes Factor

Table A14

Study 2: $N = 20$: Posterior summaries of coefficients

Parameter	μ	ρ	R^2
<i>Intercept</i>	727.6[27.3, 27.8]		
θ	-0.44 [-0.68, -0.19]	-0.07 [-0.3, -0.16]	.005
θ^2	-0.63 [-0.78, -0.47]	-0.66 [-0.78, -0.51]	.44
β	1.6 [1.08, 1.62]	0.70 [0.57, 0.81]	.49

Analysing performance: Mean of the posterior distribution of the covariates from the model with the highest support. Each parameter and its correlation with search length is associated with the lower and upper limits of the 95% highest density interval.

Table A15

Study 1: Deviance Information Criterion (DIC) for the LTM, BOM, and ITM for both conditions ($SD = 10$ and $SD = 40$).

Condition	Model	DIC
$SD = 10$	LTM	21,874
	ITM	26,874
	BOM	27,636
$SD = 40$	LTM	30,170
	ITM	30,687
	BOM	30,843

Note: Lower numbers indicate a better fit. Lowest DIC in bold. LTM: Linear Threshold Model, BOM: Biased Optimal Model, ITM: Independent Threshold Model

Table A16

Study 2: Deviance Information Criterion (DIC) for the LTM, BOM, and ITM for the three length conditions ($N = 5, 10, \text{ and } 20$).

Length	Model	DIC
$N = 5$	LTM	4,735
	ITM	4,845
	BOM	4,934
$N = 10$	LTM	7,174
	ITM	7,240
	BOM	7,570
$N = 20$	LTM	9,966
	ITM	9,969
	BOM	10,489

Note: Lower numbers indicate a better fit. Lowest DIC in bold. LTM: Linear Threshold Model, BOM: Biased Optimal Model, ITM: Independent Threshold Model