Setting and Adjusting Thresholds in an Optimal Stopping Task: Model Predictions and Empirical Results

Erin H. Bugbee (ebugbee@cmu.edu)
Department of Social and Decision Sciences, Carnegie Mellon University
Pittsburgh, PA 15213 USA

Cleotilde Gonzalez (coty@cmu.edu)
Department of Social and Decision Sciences, Carnegie Mellon University
Pittsburgh, PA 15213 USA

Abstract

Researchers have proposed that people set thresholds to decide when to stop searching in optimal stopping tasks with full information, where option values are known. Most models assume that individuals set internal thresholds to guide their stopping decisions. However, whether humans actually set and adjust thresholds with experience remains unexamined. This experiment investigates how people set and adjust thresholds and whether this affects search behavior and learning over time. We designed an optimal stopping task where participants either report a threshold before seeing the option's value or proceed without setting one. In addition, we varied whether the set threshold was binding for stopping decisions. Our findings, based on model predictions and empirical data, suggest that setting thresholds leads to more errors and lower accuracy. Accuracy is lowest when thresholds are non-binding. Participants often deviate from their set thresholds and perform better for doing so. These results challenge the assumption that people rely on thresholds for stopping decisions. Instead, they seem to learn from experience to improve accuracy and reduce errors, offering new insights into sequential decision making.

Keywords: optimal stopping; sequential decision making; sequential search; thresholds; instance-based learning theory

Introduction

Imagine that you have been considering buying a new winter coat. The coat is expensive at $200, so you decide to wait until it goes on sale. You will purchase it if it is marked 25% off at a price of $150. Suddenly, it is November and the price has not dropped as much as you hoped, so you decide that you are willing to purchase the coat at 15% off for $170 because you do not want to be cold for winter. If winter starts and the price has not fallen below your threshold, you plan to buy the coat for full price. At the end of November, the coat is on sale for 20% off and since this exceeds your threshold, you purchase the coat. This is an example of how a decision maker may decide when to stop searching and make a selection in an optimal stopping problem. Decision makers may set explicit thresholds to decide when to stop their search, or perhaps they may simply use their similar past experiences to decide when to make the purchase.

Researchers have argued that people set thresholds to make these kinds of decisions (Guan, Stokes, Vandekerckhove, & Lee, 2020; Campbell & Lee, 2006; Baumann, Singmann, Gershman, & von Helversen, 2020). When deciding when to stop, they compare the current alternative to some threshold in their mind, and if it is better than the threshold, they choose it, and if not, they continue searching. Furthermore, this literature suggests that this threshold may change throughout a sequence of alternatives, as people may start with high standards and become more accepting as the number of remaining options dwindles. Researchers have proposed that the adjustment of these thresholds may take a nonlinear or linear shape over a sequence, characterizing thresholds by how they deviate from optimal (the Bias-From-Optimal (BFO) Model; Guan et al., 2020) or based on a linear change (the Linear Threshold Model (LTM); Baumann et al., 2020). However, this research infers these thresholds and their adjustments based on observed human choices. Although researchers have found evidence that people accept less valued options later in the sequence (Lee, 2006), to our knowledge, there is no evidence that confirms that humans actually set or adjust these thresholds explicitly in optimal stopping tasks and that they make decisions based on thresholds.

Recently, Bugbee and Gonzalez (2024) proposed that humans make stopping decisions by using their previous experiences to dynamically adjust their decisions rather than setting thresholds beforehand. Their conclusion was drawn from simulations using a cognitively-grounded learning model, comparing the model's predictions to human data, rather than directly fitting it. Although their study provides indirect evidence that stopping decisions occur without predefined thresholds, no empirical research has explicitly tested whether humans set thresholds or rely solely on experience when making stopping decisions.

In this research, we aim to address this gap in the literature. We designed an experiment where participants were either asked to set a threshold or not, and when asked to set a threshold, we constrained them to a binding or non-binding setting. This design helps to determine the value of setting thresholds and whether humans adjust their thresholds with experience. It also demonstrates the predictions of an Instance-Based Learning (IBL) model for optimal stopping (Bugbee & Gonzalez, 2024), where decisions are made from experience and in the absence of pre-set thresholds.

Hypotheses

We investigate the thresholds that people explicitly set and the shape the threshold adjustment takes within a sequence. Past research has tested specific adjustment functions, such as those that are consistently biased from optimal (Guan et al., 2020) or linear (Baumann et al., 2020). We are also interested in whether these threshold values change depending

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

on whether the thresholds are binding or not, and how the set thresholds may change over time with experience.

We hypothesize that the thresholds people set will not follow the biased from optimal shape as in the BFO model (Guan et al., 2020) or the linear shape as in the LTM model (Baumann et al., 2020). Instead, as has been argued by Bugbee and Gonzalez (2024), we expect that people will learn to stop more optimally over time, by using their experiences in the task, and not by setting thresholds. That is, we expect that the thresholds set by participants will not translate into their stopping decisions naturally, and, in fact, they may hinder their accuracy and learning.

We hypothesize the following:

Hypothesis 1: Behavior will differ between conditions where thresholds are set and not set, as explicitly considering thresholds may influence decision making.

Hypothesis 2: Accuracy will be higher when not setting thresholds compared to when setting them.

Hypothesis 3: If people learn from experience when setting thresholds, the set thresholds will become more optimal over time; otherwise, they will remain stable.

Hypothesis 4: When setting non-binding thresholds, people may not adhere to them. Since thresholds may not naturally guide stopping decisions, the ones they set may not reflect their true intentions, leading to revisions of their behavior.

Optimal Stopping Task

We use an optimal stopping task introduced in Bugbee and Gonzalez (2024). In this task, participants are told that they will encounter sequences of boxes. Each sequence contains 10 boxes. The length of the sequence and the current box's position in the sequence are known to participants. Each box has a value independently sampled from a distribution with no repetitions within a sequence. The distribution for this task is normally distributed with a mean of 50 and a standard deviation of 20, bounded from 0 to 100. This distribution is unknown to the participants. The goal is to Select the box with the maximum value in the sequence. They are rewarded with bonus money only if they correctly select the maximum. There is only one correct maximum, and therefore each box has a correct action: either Pass if it is not the maximum, or Select if it is the maximum. Participants are shown a box with a value and must decide to Pass the box if they do not believe that it is the maximum. If they Pass, they move on to the next box. This continues until they Select a box, or until they reach the last box, in which case they are forced to Select it. After selecting a box, the participant is given detailed feedback following the detailed feedback condition in Bugbee and Gonzalez (2024), meaning that they are informed of the correctness of their decision, as well as which box was the correct choice and its corresponding position and value. After feedback, the participant continues to the next sequence. This repeats for 50 problems.

Experimental Methods

We pre-registered an experiment to test the hypotheses presented above in the optimal stopping task. The preregistration, data, and scripts used for the results presented in this paper are available on the Open Science Framework at `https://osf.io/zqa5r`.

Participants

Participants were recruited through Amazon Mechanical Turk. They were paid $3.00 for completing the task and $0.06 for each correct problem of 50 total problems, with the opportunity to earn up to $3.00 in bonus payment (mean bonus = $1.19, $SD = 0.55$). Participants were removed from the analysis if they attempted any part of the task multiple times or if we did not obtain complete data from them, leaving a total of 143 participants for analysis (36% women, 63% men, and 1% non-binary or preferred not to answer; median age = 41 years, $SD = 10.90$).

Experimental Design

Participants were randomly assigned to one of three between-subjects conditions, varying whether a threshold was set and whether it was binding (No Setting: $n = 52$, Non-Binding: $n = 46$, Binding: $n = 45$).

In the No Setting condition, the participants encounter a box and then decide whether to Pass or Select it without explicitly setting a threshold. This condition is a replication of a condition in a previous experiment (Bugbee & Gonzalez, 2024) and is similar to other optimal stopping tasks in the literature (Guan et al., 2020; Baumann et al., 2020).

In the Binding and Non-Binding conditions, the experiment has additional instructions and steps, shown in Figure 1. Before starting the task, participants are informed about the concept of a threshold. They then answer questions to show that they understand the concept. Once the task begins, for each box in each problem, participants must set a threshold value before the box value appears (Figure 1a). The threshold they set is the minimum value that they would select. They continue setting thresholds and making choices for every encountered box in a problem until a selection is made.

In the Non-Binding condition, the set threshold does not restrict the action that the participant can take. Participants are informed of how the box value compares to their set threshold and which action they planned to take. They can press the Pass or Select button regardless of the threshold they set (Figure 1b). Their thresholds are non-binding because the provided values do not dictate the actions taken.

In the Binding condition, participants are informed that the threshold value they set affects the action they take for that box. That is, if the box value is equal to or above their set threshold, the Pass button is disabled and they must click Select (Figure 1d). If the box value is below their set threshold, the Select button is disabled and they must Pass (Figure 1c). In this condition, their thresholds are binding because the provided value dictates the actions they can take.

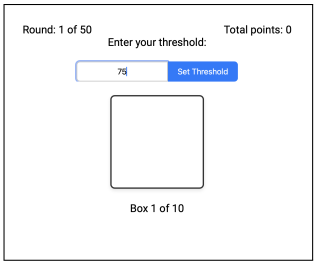

(a) Participants set a threshold for each box in the Non-Binding and Binding conditions.

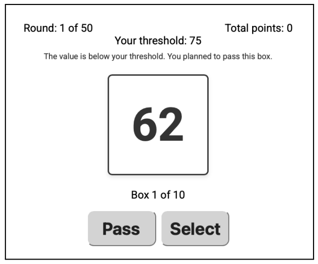

(b) In the Non-Binding condition, the participant decides which action to take, regardless of the set threshold.

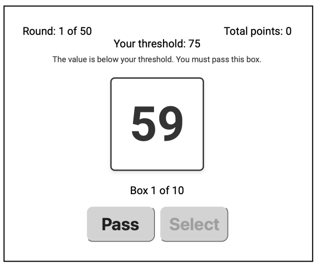

(c) In the Binding condition, if the value is below the threshold, only the Pass button is available.

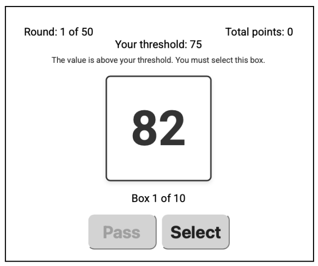

(d) In the Binding condition, if the value is above or equal to the threshold, only the Select button is available.

Figure 1: The process of setting thresholds and making the choice to Pass or Select in the Non-Binding and Binding conditions.

Models

Optimal Model

Given a particular sequence length and distribution, there is an optimal solution that involves following optimal thresholds that depend on the box position within the sequence. These thresholds can be calculated following the process proposed by Gilbert and Mosteller (1966). We calculate the optimal thresholds using the percentiles from Table 7, Column 2 of Gilbert and Mosteller (1966), in alignment with Table 1, Column 2 from Goldstein, McAfee, Suri, and Wright (2020). We simulated an *Optimal Agent* that follows the optimal thresholds for each human participant: that is, the agent compares the value of the current box to the optimal threshold for that position in the sequence, and if the value meets or exceeds the threshold and is the best so far, it selects it, otherwise it passes. Optimal agents do not learn; they deterministically decide when to stop based on the optimal thresholds.

Instance-Based Learning Model

We built an Instance-Based Learning (IBL) model to make simulated predictions of what humans would do in this experiment. The model makes sequential decisions according to the general theory of decisions from experience, Instance-Based Learning Theory (IBLT) (Gonzalez, Lerch, & Lebiere, 2003), which assumes that decisions are made from experience rather than by using thresholds. IBL models have been shown to be predictive in several sequential decision tasks, including optimal stopping tasks (Bugbee, McDonald, & Gonzalez, 2022; Bugbee & Gonzalez, 2024) and risky choice tasks (Bugbee & Gonzalez, 2022b). We briefly summarize the theory, and refer the reader to previous publications in which the algorithm and mathematical equations of IBL models have been reported (see Nguyen, Phan, & Gonzalez, 2022; Gonzalez, 2023).

IBLT proposes that learning occurs through the accumulation of memory units called instances. Each instance represents a potential decision or a decision made, and each instance has an activation value that represents the ease of retrieval of that information from memory, according to their similarity to the current situation, their frequency of occurrence, and their recency (see ACT-R equation by Anderson and Lebiere (2014)). A blended value (BV) is a form of expected utility, calculated as the sum of the product of the probability of retrieval of each instance and its utility. For each decision, the IBL agent chooses the alternative with the highest BV. When the agent receives feedback, this is stored as the utility for that instance.

The instance structure for this model is the same representation used in previous versions of this model for the same optimal stopping task (Bugbee & Gonzalez, 2024). The state consists of the value of the box and the number of boxes remaining in the sequence after the current box; the action is to Select or Pass; and the utility is binary for the correctness of the decision. We use linear similarity to compare each of the attributes of the current state and the state of past instances, and the decay and noise parameters are set to the ACT-R default values of $d = 0.5$ and $\sigma = 0.25$ respectively. The model also uses credit assignment (Nguyen et al., 2022), assigning an expected value, which is the blended value for passing at that position, to all pass actions until a selection is made, and then all decisions are updated to the outcome obtained.

Before starting the task, we provide agents with prior knowledge in the form of prepopulated instances. We populate all agents with instances giving a utility of 2 to selecting 100 in the first box and a 0 to passing it, and a 0 to selecting 0 in the first box and 2 for passing it. This is to represent knowledge people may have, such as knowing that higher values are better than lower ones. The utilities are higher than what can possibly be obtained to encourage exploration of both actions.

For the No Setting condition, the agent encounters each box and is free to decide to either Pass it and move on the next box, or Select it and conclude search for that problem. For the Non-Binding and Binding conditions, the agent generates a threshold and sets it before seeing the box's value.

The thresholds are generated according to Algorithm 1 for both the Non-Binding and Binding conditions. The goal is to determine the box value where the blended value of the Pass action is equal to that of the Select action. This implies that the agent is indifferent between passing or selecting, so this translates to the threshold. To find this threshold for a position, we use binary search. The value bounds are set to 0 and

100 and the midpoint of the range is calculated. The blended values for Pass and Select for the midpoint value at that position are calculated. If the difference between the blended values is greater than 0, the BV for Select is greater than that for Pass, meaning the agent prefers to Select, and the threshold should be lowered by setting the upper bound to the midpoint and repeating; if below 0, the BV for Pass is greater than for Select so the threshold should be raised by setting the lower bound to the midpoint and repeating. This iteratively narrows down the range of possible threshold values until the range is within tolerance $\varepsilon = 1$. The process is repeated for every position in the sequence, resulting in a set of n thresholds. This is also repeated for every sequence, allowing the agent to incorporate newly gained experiences into the blended values, resulting in changes in the set thresholds over time.

Algorithm 1 IBL Agent Threshold Generation

Input: Sequence length n, value bounds $(lower, upper)$, tolerance ε
Output: Thresholds $\{v_1^*, v_2^*, \ldots, v_n^*\}$
Initialize $thresholds \leftarrow []$
for each position p in 1 to n **do**
    Define BV_difference$(v) = BV_{\text{Select}}(v, p) - BV_{\text{Pass}}(v, p)$

    Perform binary search to find v^*:
    while $upper - lower > \varepsilon$ **do**
        $mid \leftarrow \frac{lower+upper}{2}$
        Update bounds based on BV_difference(mid):
            If BV_difference$(mid) > 0$, set $upper \leftarrow mid$
            Else, set $lower \leftarrow mid$
    Add $\lceil v^* \rceil$ to $thresholds$, where $v^* \leftarrow \frac{lower+upper}{2}$
Return: $thresholds$

In the Non-Binding condition, the agent determines its threshold to set but is free to decide to Pass or Select, as in the No Setting condition, regardless of its set threshold. Therefore, the only difference between the No Setting and the Non-Binding conditions is that the agent goes through the process of determining its own threshold. In the Binding condition, the agent's threshold dictates which action is taken. If the box's value is equal to or above the set threshold, the agent must Select it; if below, the agent must Pass.

Simulation Predictions from IBL Model

We simulated an Optimal Agent and an IBL Agent for each human participant in the experiment. These agents encountered the same sequences of box values in the same order as the human, under the same experimental condition. The number of agents is the same as the number of participants in each condition, but no human data is used to make the predictions.

Simulation Results

With the IBL model, we predict the behavior of the participants in this task. We averaged across the individual agent predictions for these figures. In Figure 2 (top), we see the proportion of selections of the IBL agents that involve correctly choosing the maximum value box, making an under-exploration error and stopping too early, and making an over-exploration error and stopping too late. For comparison, an agent choosing randomly from the 10 boxes would achieve 10% accuracy. The IBL model predicts that participants in the No Setting condition will have the highest accuracy, slightly below optimal. Both Non-Binding and Binding conditions are predicted to result in substantially lower proportions of correct decisions, with slightly higher accuracy in the Binding compared to the Non-Binding condition.

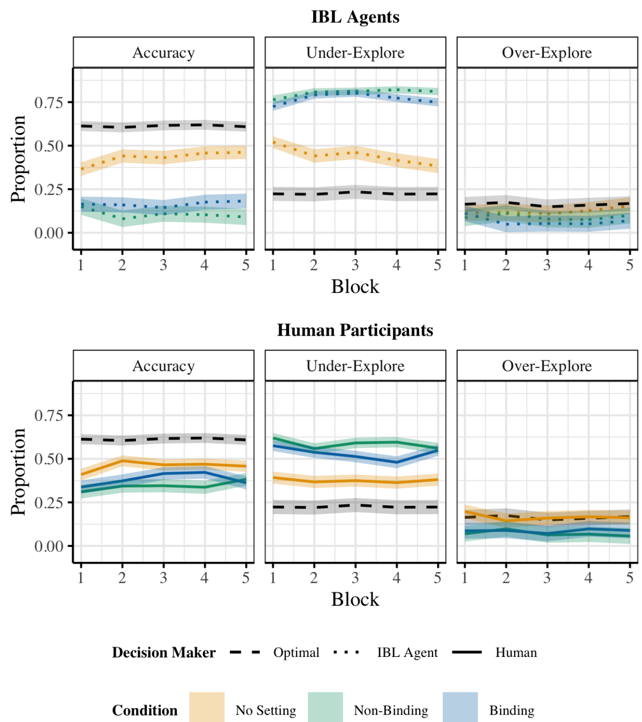

Figure 2: Accuracy, under-explore errors, and over-explore errors for optimal agents and IBL agents (top), as well as human participants (bottom), over blocks of 10 problems by condition. Error bars represent standard errors.

IBL agents in all conditions make more errors of under-exploration than is optimal, though they make fewer errors in the No Setting condition relative to the threshold setting conditions. Agents in the No Setting condition have close to the optimal errors of over-exploration, while the Threshold conditions have slightly fewer over-exploration errors.

Figure 3 (left) shows the optimal search length and the search length predicted by the IBL agents over blocks of 10 problems. IBL agents predict that participants will search more in the No Setting condition, although less than optimal. The model also predicts a short search length and minimal difference in the threshold setting conditions.

In the Non-Binding condition, IBL agents predict that participants will adhere to their set thresholds on approximately 25.6% (SE = 0.005) of the boxes. Therefore, the agents predicted that participants will not follow their set thresholds. The agents predict that participants would plan to Select but choose to Pass for approximately 57.0% ($SE = 0.006$) of the boxes, and would plan to Pass but choose to Select for approximately 17.4% ($SE = 0.005$).

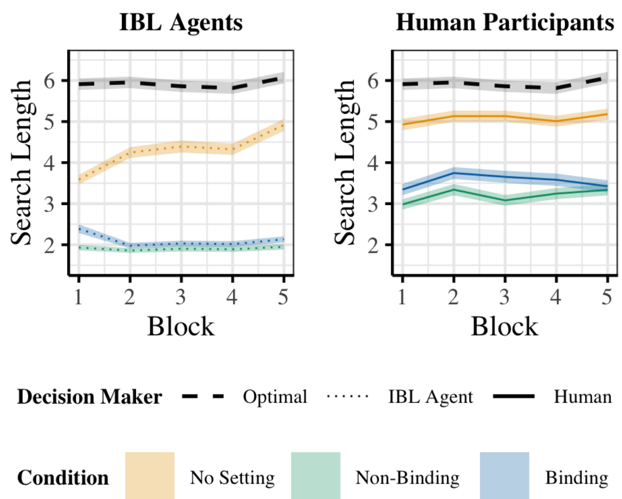

Figure 3: Search length for IBL agents (left) and human participants (right) alongside optimal agents, over blocks of 10 problems by condition. Error bars represent standard errors for the mean.

Experimental Results

Figure 2 (bottom) shows the accuracy (proportion of problems in which the decision maker stopped at the correct box) and the proportion of under-exploration and over-exploration selections relative to the correct box over blocks of 10 problems. Table 1 reports the results of repeated-measures ANOVAs predicting accuracy, under-explore errors, over-explore errors, and search length. We perform t-tests with Bonferroni adjustments for pairwise comparisons.

Supporting *Hypothesis 1*, we observe that behavior differs greatly between conditions where thresholds are set and not set. As predicted by the IBL model, we observe that participants who do not set a threshold are more accurate ($M = 0.46$, $SE = 0.013$) than those who set thresholds, regardless of whether they are binding ($M = 0.38$, $SE = 0.016$; $t(458.95) = 3.70$, $p < 0.001$) or not ($M = 0.34$, $SE = 0.015$; $t(471.64) = 5.67$, $p < 0.001$). This provides support for *Hypothesis 2*. Accuracy was indistinguishable for binding and non-binding ($t(452.16) = 1.77$, $p = 0.211$). We also observe a significant main effect of block, indicating that participants learn to increase their accuracy over time.

We observe a significant main effect of condition for both under-explore and over-explore errors. Participants who were not asked to set thresholds had substantially fewer under-explore errors ($M = 0.38$, $SE = 0.013$) than the binding ($M = 0.53$, $SE = 0.019$; $t(436.17) = 9.78$, $p < 0.001$) and non-binding conditions ($M = 0.38$, $SE = 0.013$; $t(436.17) = 9.78$, $p < 0.001$). They also had fewer over-explore errors ($M = 0.17$, $SE = 0.009$) than the Binding ($M = 0.09$, $SE = 0.007$; $t(475.82) = 8.38$, $p < 0.001$) and Non-Binding conditions ($M = 0.07$, $SE = 0.007$; $t(467.19) = 7.09$, $p < 0.001$). There was no significant difference between the Binding and Non-Binding conditions in terms of under-explore errors ($t(447.41) = 2.12$, $p = 0.065$) or over-explore errors ($t(452.7) = 1.60$, $p = 0.49$).

Figure 3 (right) shows the search length of participants for each of the experimental conditions over blocks of 10 problems, and Table 1 shows the results of a repeated-measures ANOVA. Again, similar to the predictions of the IBL model, participants explored less than optimal (for participants, $M = 3.99$, $SE = 0.037$; for optimal agents, $M = 5.92$, $SE = 0.036$; $t(14281) = 37.52$, $p < 0.001$). The No Setting condition was closer to the optimal search length ($M = 5.08$, $SE = 0.060$), while the Binding ($M = 3.55$, $SE = 0.065$), Non-Binding ($M = 3.20$, $SE = 0.059$) were significantly lower than optimal. Clearly, as predicted by the IBL model, participants who did not set a threshold searched longer than participants who set thresholds, both Binding ($t(4749) = 17.24$, $p < 0.001$) and Non-Binding ($t(4889) = 22.23$, $p < 0.001$). The model did not predict differences in the search lengths of the Binding and Non-Binding conditions, but human data showed a longer search in the Binding condition than in the Non-Binding condition ($t(4501) = 4.02$, $p < 0.001$).

Participant Threshold Setting

Figure 4 shows the average thresholds set by participants in the Binding and Non-Binding conditions at each position of the sequence of 10 boxes, compared to the theoretical optimal thresholds, over blocks of 10 problems. It is clear that the thresholds set by the participants do not align with the optimal thresholds. Participants in the Binding condition ($M = 66.20$, $SE = 0.32$) set higher thresholds on average than participants in the Non-Binding condition ($M = 59.06$, $SE = 0.46$; $t(13490) = 12.74$, $p < 0.001$), exclusive of three boxes where participants set thresholds over 1,000. A repeated-measures ANOVA indicates that there is no significant difference in the set thresholds by block, position, or condition. Regarding *Hypothesis 3*, this indicates that people are not learning to improve their thresholds with experience, as they do not become closer to optimal. Instead, they remain stable, despite participants learning to increase their accuracy over time, as indicated by the main effect of block.

Threshold Adherence Proportion for Non-Binding

Very close to the proportion predicted by the IBL agents ($M = 25.6\%$, $SE = 0.005$), participants in the experiment adhered to the thresholds set in the Non-Binding condition only 20.7% ($SE = 0.005$) of the time. That is, they did not follow their own set thresholds, providing evidence for *Hypothesis 4*. In more than half of the boxes ($M = 55.8\%$, $SE = 0.006$), participants planned to Select but chose to Pass, indicating that they set their threshold too low for their action. In about a

Table 1: Repeated-measures ANOVA results predicting accuracy, under-explore errors, over-explore errors, and search length.

	Df	Accuracy		Under-Explore Errors		Over-Explore Errors		Search Length	
Source		F value	$Pr(>F)$	F value	$Pr(>F)$	F value	$Pr(>F)$	F value	$Pr(>F)$
Error: Participant									
Condition	2	5.27	0.006**	20.24	<0.001***	14.04	<0.001***	23.81	<0.001***
Residuals	140								
Error: Participant:Block									
Block	4	3.39	0.009**	0.92	0.453	1.81	0.126	1.85	0.118
Condition:Block	8	1.08	0.378	1.31	0.234	0.88	0.536	0.50	0.857
Residuals	560								

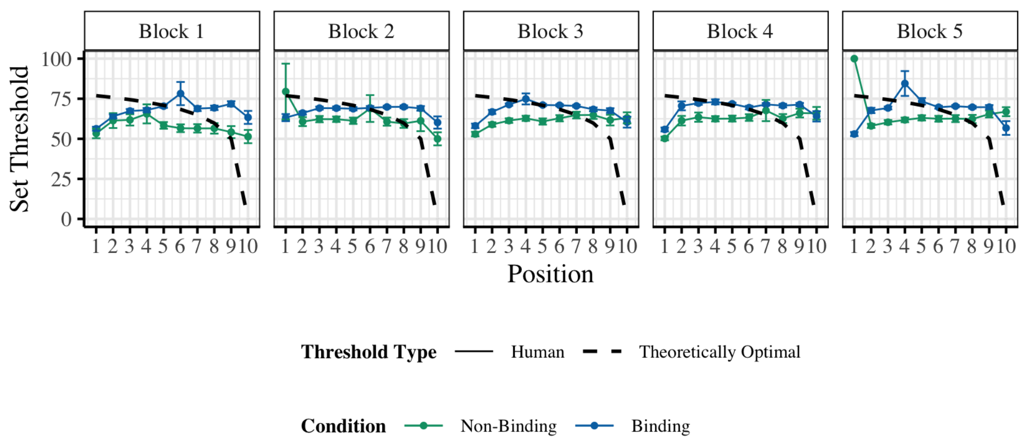

Figure 4: Set thresholds for each position by condition for human participants relative to theoretically optimal thresholds. Error bars correspond to standard errors.

quarter of the boxes ($M = 23.5\%$, $SE = 0.005$), participants planned to Pass but chose to Select, meaning that they set their threshold too high.

Higher accuracy was achieved when participants did not adhere to their previously set thresholds ($M = 0.40$, $SE = 0.012$) compared to when they did adhere ($M = 0.18$, $SE = 0.016$; $t(1217.1) = 10.72$, $p < 0.001$). Therefore, they are often correct to not follow their own set thresholds, indicating that they are setting their thresholds poorly with respect to both their own intrinsic preferences and the correct selection. As shown in Figure 4 and previously discussed, they do not learn to set improved thresholds over time.

Discussion

Our research addresses an important gap in optimal stopping research. We experimentally evaluate the effect of setting explicit thresholds and investigate human stopping decisions depending on whether those thresholds are binding or not. Although past research suggests that people set thresholds for optimal stopping tasks (Lee, 2006; Guan et al., 2020; Baumann et al., 2020), recent work argues that decisions emerge dynamically from experience rather than rigid thresholds (Bugbee & Gonzalez, 2022a, 2024).

IBLT (Gonzalez & Aggarwal, 2023; Gonzalez et al., 2003) helps to explain how stopping decisions emerge from experience and can predict experimental results from theoretical simulations. Using the value observed in each box in a sequence and the knowledge of the number of boxes remaining in the sequence, the IBL model predicts the value of the Select and Pass choice options and selects the option that provides the highest expected value from experience. With this process, using the cognitive and memory mechanisms in IBLT, the model is able to predict the participants' stopping behavior and their errors of under- and over-exploration.

As predicted, behavior differed based on whether thresholds were set (*Hypothesis 1*), and accuracy was higher for participants who did not set thresholds (*Hypothesis 2*). The thresholds did not align with the optimal thresholds and no evidence of learning to improve threshold setting was found (*Hypothesis 3*). Participants in the Non-Binding condition adhered to their thresholds only 20.7% of the time, often setting them too low (*Hypothesis 4*). Importantly, accuracy was higher when they ignored their thresholds.

These findings challenge the assumption that humans rely on explicit thresholds for stopping decisions. Instead, they dynamically adjust based on experience, suggesting that threshold setting is an artificial constraint rather than a natural decision-making strategy.

Acknowledgments

This paper is based upon work supported by the AI Research Institutes Program funded by the National Science Foundation under AI Institute for Societal Decision Making (AI-SDM), Award No. 2229881.

References

Anderson, J. R., & Lebiere, C. J. (2014). *The atomic components of thought*. Psychology Press.

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*, *117*(23), 12750–12755. doi: 10.1073/pnas.2002312117

Bugbee, E. H., & Gonzalez, C. (2022a). Deciding when to stop: Cognitive models of sequential decisions in optimal stopping tasks. doi: 10.1184/R1/20492877.V2

Bugbee, E. H., & Gonzalez, C. (2022b). Making predictions without data: How an instance-based learning model predicts sequential decisions in the balloon analog risk task. In *Proceedings of the annual meeting of the cognitive science society* (p. 3167-3174).

Bugbee, E. H., & Gonzalez, C. (2024). Feedback promotes learning and knowledge of the distribution of values hinders exploration in an optimal stopping task. In *Proceedings of the annual meeting of the cognitive science society* (p. 4618-4624).

Bugbee, E. H., McDonald, C., & Gonzalez, C. (2022). Leveraging cognitive models for the wisdom of crowds in sequential decision tasks. In *Paper presented at Virtual MathPsych/ICCM 2022* (p. 7). Retrieved from `mathpsych.org/presentation/751`

Campbell, J., & Lee, M. D. (2006). The effect of feedback and financial reward on human performance solving 'secretary' problems. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 28).

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. doi: 10.1080/01621459.1966.10502008

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

Gonzalez, C. (2023). Building human-like artificial agents: A general cognitive algorithm for emulating human decision-making in dynamic environments. *Perspectives on Psychological Science*, 17456916231196766. doi: 10.1177/17456916231196766

Gonzalez, C., & Aggarwal, P. (2023, June). Sequential decisions from sampling:: inductive generation of stopping decisions using instance-based learning theory. In K. Fiedler, P. Juslin, & J. Denrell (Eds.), *Sampling in Judgment and Decision Making* (1st ed., pp. 289–310). Cambridge University Press. Retrieved 2024-02-08, from `https://www.cambridge.org/core/product/identifier/9781009002042%23CN-bp-13/type/book_part` doi: 10.1017/9781009002042.017

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591–635. doi: 10.1207/s15516709cog2704_2

Guan, M., Stokes, R., Vandekerckhove, J., & Lee, M. D. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision Making*, *15*(5), 823–850. doi: 10.31234/osf.io/evzp9

Lee, M. D. (2006, May). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26. Retrieved 2022-03-04, from `http://doi.wiley.com/10.1207/s15516709cog0000_69` doi: 10.1207/s15516709cog0000_69

Nguyen, T. N., Phan, D. N., & Gonzalez, C. (2022). SpeedyIBL: A comprehensive, precise, and fast implementation of instance-based learning theory. *Behavior Research Methods*. doi: 10.3758/s13428-022-01848-x