

A Hierarchical Cognitive Threshold Model of Human Decision Making on Different Length Optimal Stopping Problems

Maime Guan (hongyang@uci.edu)

Michael D. Lee (mdlee@uci.edu)

Joachim Vandekerckhove (joachim@uci.edu)

Department of Cognitive Science, University of California, Irvine
Irvine, CA 92617 USA

Abstract

In optimal stopping problems, people are asked to choose the maximum out of a sequence of values, under the constraint that a number can only be chosen when it is presented. We present a series of threshold models of human decision making on optimal stopping problems, including a new hierarchical model that assumes individual differences in threshold setting are controlled by deviations or biases from optimality associated with risk propensity, and is applicable to optimal stopping problems of any length. Using Bayesian graphical modeling methods, we apply the models to previous data involving 101 participants with large individual differences who completed sets of length 5 and length 10 problems. Our results demonstrate the effectiveness of the bias-from-optimal hierarchical model, find individual differences in thresholds that people use, but also find that these individual differences are stable across the two optimal stopping tasks.

Keywords: optimal stopping; secretary problem; sequential decision-making; threshold models; hierarchical Bayesian modeling

Introduction

The optimal stopping problem, also known as the secretary problem, is a decision-making task in which people must choose the highest value out of a sequence of numbers, under the constraint that a number can only be chosen when it is presented (Ferguson, 1989; Gilbert & Mosteller, 1966). Optimal stopping problems are interesting for understanding human decision making because they have two features found in many real-world decision-making settings. The first feature is that there is *no going back*. Oftentimes, it is difficult or even impossible to decline an earlier option and then return to it later. For example, in searching for jobs, it is almost impossible to come back to a job offer that you have already rejected. The second feature is that *only the best will do*. In some real-world situations, there is only one best option and any other option is completely and equally useless. For example, trying to find the correct key out of a set to open the door to your house will only result in success if you find the lone correct house key.

In the cognitive sciences, people's decision making on a number of different versions of optimal stopping problems have been studied. One is the classic rank order version, in which only the rank of the current option relative to the options already seen is presented (e.g., Seale & Rapoport, 1997, 2000). We focus on the alternative full-information version of the task, in which people are presented with the actual continuously-scaled values of the alternatives (e.g., Lee, 2006). In the full information optimal stopping problem, the

known optimal solution is to choose the first value that is both currently maximal and above a certain mathematically derived threshold for the current position in the sequence (Gilbert & Mosteller, 1966, Table 2).

Our previous work (e.g., Guan & Lee, 2014; Lee, 2006) found evidence that people use a series of thresholds to make decisions, and that there are large individual differences in thresholds, with many people using suboptimal thresholds. In this paper, we examine decision making on two optimal stopping tasks with different lengths. In one task, people must choose the maximum out of 5 numbers, and in the other they must choose the maximum out of 10 numbers. If there are psychological components that determine the thresholds in which people use, such as risk propensity or intelligence, then we should expect behavior to be similar between the two tasks. For instance, participants who use thresholds higher than optimal in the length 5 task should also use thresholds higher than optimal in the length 10 task. A recruiter who is generally picky and willing to hold out until the perfect candidate comes along will have relatively high thresholds for job applicants whether they are choosing from 5 applicants or 10.

Our goal is to develop a hierarchical threshold model for the optimal stopping problem that can account for the deviation or bias from optimality in terms of psychological variables, and be applicable to optimal stopping problems of any length. In the following section, we describe the experiment and the data set we use. Next, we present a series of threshold models leading up to a hierarchical psychological threshold model that is applied jointly to tasks of different lengths. We then present the results from using Bayesian methods to apply these threshold models to the behavioral data, and discuss the implications of the results.

Burns, Lee, and Vickers (2006) Data

The optimal stopping data set we use is a subset of a larger data set taken from Burns et al. (2006), which includes data from various cognitive and perceptual optimization problems, as well as various standard psychometric tests of intelligence. The original data set includes a total of 101 participants recruited from the general community, with within-participants data for the set of cognitive, perceptual, and psychometric tasks. We use just the two optimal stopping tasks, which involved 40 problems of length 5 and 40 problems of length 10.

Each participant completed the two sets of optimal stop-

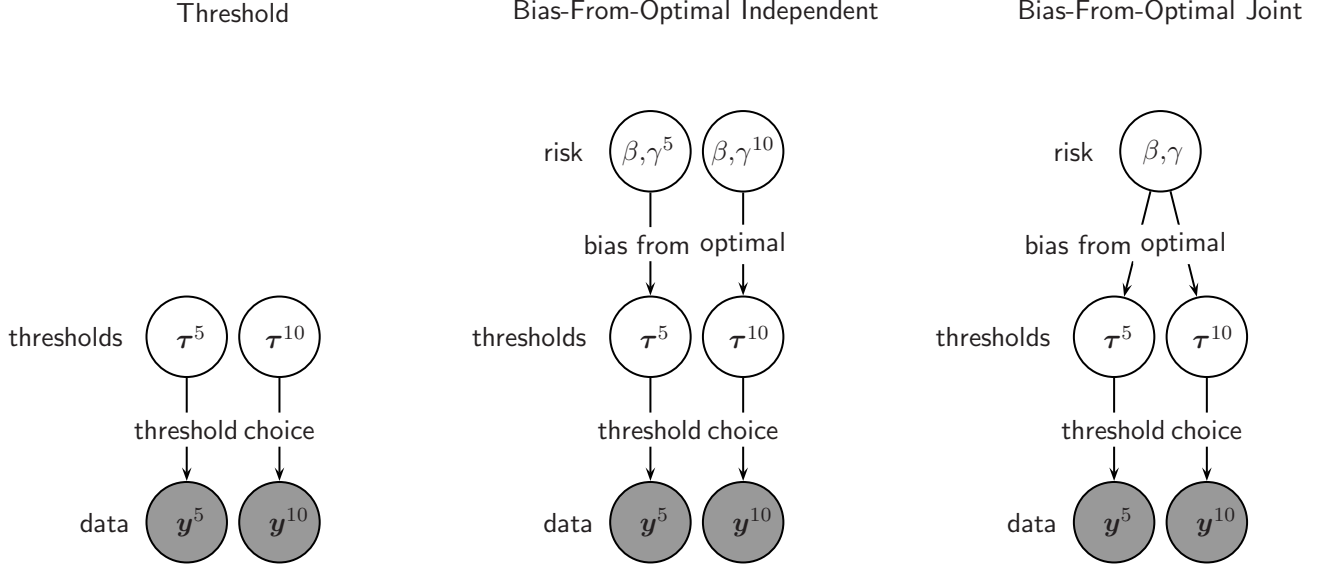


Figure 1: Conceptual overview of the threshold model, bias-from-optimal model independently applied, and bias-from-optimal model jointly applied to the optimal stopping data for problems of length 5 and length 10.

ping problems through a computer interface. All participants completed the length 5 set first, followed by the length 10 set, and completed the same 40 problems within each set. The order of problems within each set was randomized across participants.

Participants were instructed to pick the highest value out of a set of dollar amounts that ranged from 0 to 100. They were told (1) the length of the sequence, (2) that the dollar amounts were uniformly and randomly distributed, (3) that a value could only be chosen when it is presented, (4) that any value that is not the maximum is completely and equally incorrect as the others, and (5) that the last value must be chosen if no values were chosen in all previous positions. Participants indicated whether or not they chose each presented value by pressing either a “yes” or “no” button. After each problem, participants were provided with feedback on their response.

We removed 3 contaminant participants from the length 5 task, and 7 participants from the length 10 task, because they choose values that were not currently maximal on more than 10% of the problems (excepting the final value, which is a forced choice).

Overview of Models

Previous work suggests that people use threshold-based rules to make decisions in optimal stopping problems (e.g., Guan & Lee, 2014; Lee, 2006). In this section, we develop a set of three threshold models that start with the simple threshold model we have used previously, but then extend the model hierarchically to add cognitive processes and parameters accounting for how the thresholds themselves are generated. This theoretical progress is summarized in Figure 1, using a schematic form of graphical model representation.

The first threshold model consists of a set of independent thresholds, which are assumed to generate the data by a simple choice model that selects with high probability the first presented value that is above the threshold, and currently maximal. The hierarchical “bias-from-optimal” model generates thresholds based on latent psychological parameters, representing biases of deviations from suboptimality each individual has. In the bias-from-optimal independent model, we apply this model independently to both the length 5 and length 10 data. In the final bias-from-optimal joint model, however, we apply make the assumption that individual-level bias is the same for both length problems, and apply the model simultaneously to both sets of problems.

Threshold Model

The threshold model has independent threshold parameters $\tau_{i1}^m, \dots, \tau_{i(m-1)}^m$ for the i th participant in each of the first $m-1$ positions. Since the last value in a sequence must be chosen, the threshold τ_{im}^m for that position is always 0. These threshold parameters are unconstrained, with the same uniform prior probability on the subspace of $(0, 1)^{m-1}$. According to the threshold choice model, the probability the i th participant will choose the value they are presented in the k th position on their j th problem is

$$\theta_{ijk}^m = \begin{cases} \alpha_i^m & \text{if } v_{ijk}^m > \tau_{ik}^m \text{ \& } v_{ijk}^m = \max \{v_{ij1}^m, \dots, v_{ijk}^m\} \\ \frac{1-\alpha_i^m}{m} & \text{otherwise} \end{cases}$$

for the first four positions and $\theta_{ijm}^m = 1 - \sum_{k=1}^{m-1} \theta_{ijk}^m$ for the last position, where $\alpha_i^m \sim \text{Uniform}(0, 1)$ is the individual-level “accuracy of execution” parameter that describes how often the deterministic threshold model is followed (Guan & Lee,




Figure 2: The behavior of the bias-from-optimal threshold model under different parameterizations.

2014). The threshold model is completed by the observed data being distributed according to the choice probabilities, so that

$$y_{ij}^m \sim \text{Categorical}(\theta_{ij1}^m, \dots, \theta_{ij5}^m).$$

Bias-From-Optimal Independent Model

In the threshold model, the thresholds are free parameters and can consequently take any shape across positions. The bias-from-optimal model constrains the relationship between the thresholds by modeling each participant’s thresholds in terms of how they deviate from optimality.

We denote the optimal thresholds as $\tilde{\tau}_1, \dots, \tilde{\tau}_m$ for a problem of length m (Gilbert & Mosteller, 1966, Table 2). The i th participant’s thresholds now depend on a parameter $\beta_i^m \sim \text{Gaussian}(0, 1)$ that determines how far above or below their threshold is from optimal, and a parameter $\gamma_i^m \sim \text{Gaussian}(0, 1)$ that determines how much their bias increases or decreases as the sequence progresses. Formally, the i th participant’s thresholds for a problem of length m is

$$\tau_{ik}^m = \Phi(\Phi^{-1}(\tilde{\tau}_k) + \beta_i^m + \frac{k}{m}\gamma_i^m)$$

for the first $m - 1$ positions, and $\tau_{im}^m = 0$ for the last. The link functions Φ and Φ^{-1} are the Gaussian CDF and inverse CDF, respectively. The remainder of the bias-from-optimal model is identical to the threshold model, completed by the choice probabilities θ_{ijk}^m determined by the i th participant’s accuracy of execution α_i^m for a task of length m , and the problem values.

Figure 2 shows how the shape of threshold functions changes with different values of β and γ . The optimal threshold corresponds to the case with $\beta = 0$ and $\gamma = 0$, and is shown in bold. The β parameter represents a shifting bias from this optimal curve, with positive values resulting in thresholds that are above optimal, and negative values resulting in thresholds that are below optimal. The γ parameter represents

how quickly thresholds are reduced throughout the problem sequence, relative to the optimal rate of reduction. Positive values of γ produce thresholds that drop too slowly, while negative values of γ produce thresholds that drop too quickly.

The middle panel of Figure 1 provides an overview of the bias-from-optimal model. The β_i and γ_i parameters now generate the thresholds τ_i^m that the participant uses, and the same threshold choice process is then assumed to generate the observed behavioral data.

Bias-From-Optimal Joint Model

The bias-from-optimal model generates thresholds for problems of any length based on the β and γ parameters. This means it can be applied jointly to both the length 5 and length 10 tasks in our data set. The right panel in Figure 1 shows the hierarchical graphical model that achieves this simultaneous application. There are now single β_i and γ_i parameters for the i th participant that generate predictions about decisions on both tasks

We implemented all of our models as graphical models using JAGS (Plummer, 2003), to facilitate MCMC-based computational Bayesian inference (Lee & Wagenmakers, 2013). Figure 3 shows this final bias-from-optimal joint model in the graphical modeling formalism. The latent parameters corresponding to the thresholds τ_i^m and accuracy of execution α_i^m are represented by unshaded and circular nodes, since they are unobserved and continuous. The values v_j^m presented on the j th problem, standardized to lie between 0 and 1, instead of the 0 to 100 scale used in the experiment, are shown as a shaded node, since they are observed and continuous. Together, the parameters and problem values determine the probabilities θ_{ij}^m for each possible decision, shown as a double-bordered node since it is a deterministic function of its parents in the graphical model. The decision y_{ij}^m is shown as a shaded and square node, since it is observed and discrete. Encompassing plates for participants and problems indicate independent replications of the graph structure in the model.

Modeling Results

In this section, we apply all three models to all of the data from both optimal stopping tasks. We first examine the descriptive adequacy of each model, and the thresholds they infer. We then present a generalization test of the joint model, in which the data from one task are withheld. All of our modeling results are based on four chains of 2000 samples each, collected after 1000 discarded burn-in samples. The chains were verified for convergence using the standard \hat{R} statistic (Brooks & Gelman, 1997).

Descriptive Adequacy

We first measured the ability of the models to describe the behavioral data, using a standard Bayesian approach based on posterior predictive checking (Gelman, Carlin, Stern, & Rubin, 2004). Specifically, we measured the agreement between each model’s modal posterior prediction and the observed




Figure 3: Graphical model for the joint application of the bias-from-optimal threshold model to both length 5 and length 10 optimal stopping tasks.

decision for each participant on each problem. The posterior predictive agreement for the threshold model is 87% and 89%, for the bias-from-optimal model applied independently it is 86% and 84%, and for the bias-from-optimal model applied jointly it is 81% and 82%, for the length 5 and length 10 problems, respectively. Given that the base-rate or chance level of agreement are 20% and 10%, we conclude all three models provide reasonable accounts of participants' behavior.

Inferred Thresholds

Figure 4 shows the inferred thresholds, under all three models on both the length 5 and length 10 problems, for four representative participants. These participants were chosen because they span the sorts of individual differences seen across all participants in the data set.¹ The third participant, for example, has higher starting thresholds than the first participant, but drops their threshold more quickly, consistent with a positive β and negative γ parameterization. The fourth participant also has higher starting thresholds than the first participant but barely drops their threshold as position increases, consistent with positive β and γ .

Figure 4 shows close agreement between the inferred thresholds for all three models, on both problem lengths, for all four representative participants. The agreement between the threshold and bias-from-optimal model indicates that *the cognitive model we developed is a useful one*. The threshold model is free to find whatever thresholds are likely given the data. The bias-from-optimal model is simpler and more constrained, yet infers very similar thresholds for all four representative participants (and the vast majority of all participants) for both problem lengths. Given its ability to generate appropriate thresholds, the bias-from-optimal model has

¹The results for all participants can be found at <http://osf.io/vga6n>.

$$\begin{aligned} \beta_i &\sim \text{Gaussian}(0, 1) \\ \gamma_i &\sim \text{Gaussian}(0, 1) \\ \alpha_i &\sim \text{Uniform}(0, 1) \\ \tau_{ik}^m &= \Phi\left(\Phi^{-1}(\tilde{\tau}_k^m) + \beta_i + \frac{k}{m}\gamma_i\right) \\ \tau_{im}^m &= 0 \\ \theta_{ijk}^m &= \begin{cases} \alpha_i & \text{if } v_{ijk}^m > \tau_{ik}^m \text{ \& } v_{ijk}^m = \max\{v_{ij1}^m, \dots, v_{ijk}^m\} \\ \frac{1-\alpha_i}{m} & \text{otherwise} \end{cases} \\ \theta_{ijm}^m &= 1 - \sum_{k=1}^{m-1} \theta_{ijk}^m \\ y_{ij}^m &\sim \text{Categorical}(\theta_{ij}^m) \end{aligned}$$

a number of important advantages. One is that it is parameterized in terms of psychologically interpretable deviations from optimality, rather than simple thresholds. Another is that its simplicity—coming from the stronger theoretical assumptions it formalizes—means it requires fewer data to infer thresholds. This advantage can be seen clearly in the inferred thresholds for the length 10 problem for the second participant in Figure 4. Because this participant used relatively low thresholds, they rarely progressed far in problems for that task, and there are few decisions that inform the threshold model for later positions in the sequence. As a result, the inferences of the threshold model are much less constrained or informed than for the bias-from-optimal model.

Figure 4 also shows close agreement between the bias-from-optimal model applied independently and jointly to the two tasks. This result suggests that *the individual differences across problem lengths are stable*. That is, the same deviations from optimality parameterized by β and γ generate appropriate thresholds for a participant for both the length 5 and length 10 problems. This is clear in the representative participants with, for example, the second participant using thresholds that start low and decrease quickly for both problem lengths, while the fourth participant uses thresholds that start high and decrease slowly.

Individual Differences

Figure 5 summarizes the individual differences across all participants, using the inferences of the bias-from-optimal model applied jointly to both problem lengths. The posterior expected means of the β and γ for all participants are shown as a scatterplot in the main panel, with their marginal distributions shown as histograms. The dotted lines represent the no bias values for both parameters, corresponding where they meet to optimal thresholds. The range of individual differences is apparent, with both β and γ varying from positive to negative




Figure 4: Inferred thresholds for four representative participants for both length 5 and length 10 problems, based on the threshold model, bias-from-optimal model applied independently, and bias-from optimal model applied jointly.

biases, and all four quadrants around optimality populated by participants. There does appear to be, however, a relationship between the two bias parameters, with negative values of β often paired with positive values of γ , and vice versa. This suggests that participants who bias their thresholds to be too high also decrease them faster than is optimal, while participants who set their thresholds too low tend to decrease them more slowly than is optimal.

Generalization Performance

In our jointly applied bias-from-optimal model, the β and γ parameters are stable for the same person on problems of the two different lengths. This means that observing behavior in one task should allow the joint model to make useful predictions about behavior in the other, at the level of individual participants. To examine this possibility, we conducted two generalization tests (Busemeyer & Wang, 2000). In the first, we withheld the decision data of all participants from the length 10 task. We then used the observed decisions from the length 5 problem to predict the withheld length 10 problem decisions, based on the mode of the posterior predictive distribution for each participant on each problem. The overall proportion of agreement was 69%, which can be compared to a random-choice base-rate of 10%. In the second generalization test, we withheld the decision data of all participants from the length 5 task and used the observed decisions from the length 10 problem to predict the withheld length 5 problem decisions. The overall proportion of agreement was 74%, which can be compared to the base-rate of 20%. We think both generalization tests show impressive performance, and highlight the advantage of hierarchical models that are able to make predictions about tasks for which they have not observed data.

Discussion

Optimal stopping problems are interesting in that they have two features often found in the real world: there is no going back, and only the best will do. The full-information version

has a known optimal solution to which human performance can be compared, providing a benchmark for the study of optimality and bias. In this study, we examined performance on optimal stopping problems on two different length optimal stopping problems. We developed a hierarchical cognitive model that conceived of individuals generating thresholds based on two sorts of biases or deviations from optimality. The bias parameter β reflects how far above or below an individual's threshold strays from the optimal, the γ parameter reflects how slowly or quickly their threshold drops as position in the sequence increases. We found that the thresholds generated by our bias-from-optimal model agreed closely with the thresholds independently estimated in a non-hierarchical way, suggesting it is a useful model of the cognitive process of threshold generation. Moreover, the bias parameters were stable across the two task, suggesting that there could be common latent psychological components that help determine the thresholds people use. For example, we might expect that relatively high risk propensity in individuals would be reflected in relatively greater deviation from the optimal threshold.

A natural advantage of this hierarchical approach is that it becomes possible to apply the model simultaneously to optimal stopping problems of any length. The use of the same latent variables to explain observations in multiple tasks is a hallmark of good modeling throughout the empirical sciences (Lee, 2011), but the hierarchical use of psychological variables to make simultaneous predictions about behavior in different tasks is not widely seen in the cognitive sciences (Lee & Sarnecka, 2011; Vandekerckhove, 2014). Our demonstration of the stability of individual differences, and the ability to make accurate generalization predictions, provides a good example of how hierarchical models can be tested and applied across multiple tasks.

Natural future directions would be to understand the basis of these individual differences in bias or risk, and relate this measure to other psychometric measures of intelligence or personality, as well as to other decision-making tasks that involve decision making under uncertainty.




Figure 5: Joint and marginal distributions of the posterior means for the β and γ bias-from-optimal parameters over all participants. The four representative participants used in Figure 4 are highlighted.

Acknowledgments

MG acknowledges support from the National Science Foundation—Graduate Research Fellowship Program (grant #DGE-1321846). JV was supported by grant #1230118 from NSF’s Methods, Measurements, and Statistics panel and grant #48192 from the John Templeton Foundation.

References

Brooks, S. P., & Gelman, A. (1997). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7, 434–455.

Burns, N. R., Lee, M. D., & 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), 3.

Busemeyer, J. R., & Wang, Y. M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, 44(1), 171–189.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, 4, 282–296.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian Data Analysis* (Second ed.). Boca Raton, FL: Chapman & Hall/CRC.

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

Guan, M., & Lee, M. D. (2014). Threshold models of human decision making on optimal stopping problems in different environments. In P. Bello, M. Mcshane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual*

Conference of the Cognitive Science Society (pp. 553–558). Austin, TX: Cognitive Science Society.

Lee, M. D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, 30, 555–580.

Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55(1), 1–7.

Lee, M. D., & Sarnecka, B. W. (2011). Number knower-levels in young children: Insights from a Bayesian model. *Cognition*, 120, 391–402.

Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd international workshop on distributed statistical computing*. Vienna, Austria.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the “Secretary Problem”. *Organizational Behavior and Human Decision Processes*, 69, 221–236.

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks. *Journal of Behavioral Decision Making*, 13, 391–411.

Vandekerckhove, J. (2014). A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology*, 60, 58–71.