Cognitive Models of Optimal Sequential Search with Recall

Sudeep Bhatia

University of Pennsylvania

Lisheng He

Shanghai International Studies University

Wenjia Joyce Zhao

The Ohio State University

Pantelis P. Analytis

University of Southern Denmark

January 4, 2021

Address correspondence to Sudeep Bhatia, Department of Psychology, University of Pennsylvania, Email: bhatiasu@sas.upenn.edu. Funding was received from the National Science Foundation grant SES-1847794.

Abstract

Many everyday decisions require sequential search, according to which available choice options are observed one at a time, with each observation involving some cost to the decision maker. In these tasks, decision makers need to trade-off the chances of finding better options with the cost of search. Optimal strategies in such tasks involve threshold decision rules, which terminate the search as soon as an option exceeding a reward value is found. Threshold rules can be seen as special cases of well-known algorithmic decision processes, such as the satisficing heuristic. Prior work has found that decision makers do use threshold rules, however the stopping thresholds observed in data are typically smaller than the (expected value maximizing) optimal threshold. We put forward an array of cognitive models and use parametric model fits on participant-level search data to examine why decision makers adopt seemingly suboptimal thresholds. We find that people's behavior is consistent with optimal search if we allow participants to display risk aversion, psychological effort cost, and decision error. Thus, decision makers appear to be able to search in a resource-rational manner that maximizes stochastic risk averse utility. Our findings shed light on the psychological factors that guide sequential decision making, and show how threshold models can be used to describe both computational and algorithmic aspects of search behavior.

Keywords: sequential search; cognitive modeling; optimal stopping; model fitting; risk preference

1. Introduction

Consider the task of finding juicy cherries at a large farmers' market. Decision makers visit stalls sequentially, tasting the cherries offered by each seller. After each sample, they must decide whether or not they want to terminate the search and purchase cherries from one of the previously encountered stalls, or continue the search by tasting the cherries offered at the next stall. Search is costly (perhaps due to a restless toddler tagging along) and thus good decision making involves ending the search after appropriately juicy cherries have been found. Although decision makers may not end up purchasing the juiciest cherries available, the optimal decision strategy would nonetheless trade off the cost of search against the chances of finding a better option, and thus maximize the decision maker's overall (expected) utility from search.

Sequential search tasks are a key feature of everyday decision making, and the task structure introduced here describes many different types of common search problems (e.g., finding a cheap product at a store or a skilled employee for a job). Sequential search is also at play in cognitive tasks that do not explicitly involve decision making (e.g., searching memory for the best answer to a question). For this reason, characterizing optimal strategies in sequential search, and understanding whether or not individuals are able to appropriately implement those strategies, has been a major focus of research in disciplines such as economics, marketing, management, and statistics, as well as in cognitive psychology and cognitive science (Simon, 1956; DeGroot, 1970; Zwick et at, 2003; Lee, 2006; Caplin et al., 2011; Guan & Lee, 2018; Baumann et al., 2020; Goldstein et al., 2020).

In settings where decision makers sample options with a known reward distribution and cost of search, and can select any of the options sampled previously (a setting known as search with recall), the optimal strategy is characterized by a simple threshold rule: Terminate the

search and choose the sampled option if and only if its value exceeds a fixed predetermined threshold, derived based on the reward distributions and costs of available options (see DeGroot, 1970; Telser, 1973). Search problems with recall are part of a larger family of optimal stopping problems in which decision makers may not be able to return to options sampled early on during search, or where they may only seek the best possible option (e.g., secretary problem, see Seale & Rapoport, 1997) or where the option attributes reveal some information about the option quality and this information can be used to decide on the search order (Bearden et al., 2005; Analytis et al., 2014). Optimal solutions in problems with such specifications vary in their complexity and behavioral plausibility (see Weitzman, 1979; Harrison & Morgan, 1990; Gabaix et al., 2006).

The threshold rules encountered in many optimal stopping problems with recall involve a fairly simple set of cognitive operations, and can thus be seen as describing not only computational but also algorithmic levels of Marr's (1982) hierarchy, as it pertains to the sequential search problem. Indeed, threshold strategies are compatible with the satisficing heuristic proposed by Simon (1955, 1982) and the notion of an aspiration level, and thus the optimal strategy in many sequential search tasks is also a behaviorally plausible one. Related threshold-based models of binary choice are also considered to be good models of neural information processes (Busemeyer et al., 2019; Gluth et al., 2012; Smith & Ratcliff, 2004). Such processes typically implement threshold firing rules in individual neurons or groups of neurons, suggesting that threshold decision making also provides an account of the implementational level of Marr's hierarchy.

Given the statistical optimality, heuristic simplicity, and neural interpretability of threshold rules, it is not surprising that prior experimental work has found that decision makers

often use such rules in sequential search tasks. Importantly, however, observed thresholds typically do not correspond to "optimal" thresholds derived based on task reward distributions and search costs. Decision makers have been found to use, on average, lower thresholds that result in fewer searches and lower expected rewards than the standard optimal search model (Rapoport & Tversky, 1970; Schotter & Braunstein, 1981; Hey, 1982, 1987; Sonnemans, 2000).

The standard optimal search model assumes expected value maximization (i.e., risk neutrality), and thus several authors have suggested risk aversion as a potential explanation for apparent undersearch (Schotter & Braunstein, 1981; Sonnemans, 1998; Schunk & Winter, 2009). Of course, it is well known that decision makers are typically risk averse, and often prefer to avoid variability in reward, even if this involves a reduction in the expected value of the reward (Arrow, 1965; see also Pedroni et al., 2017 for a recent review). Risk averse agents would be content with stopping at options with a lower expected value (Nachman, 1972) and thus risk aversion could result in lower thresholds than those obtained under a risk neutral strategy.

This is not the only possible explanation. Prior work has also assumed costless effort. Specifically, optimal thresholds have previously been derived using only the explicit (monetary) costs of search. Yet search may be cognitively taxing (e.g., involving a cost of computation) and thus also impose an implicit psychological burden on the decision maker. For example, Reutskaja et al. (2011) conducted an experiment where the explicit search cost was 0 and people had a time constraint. People terminated search before the time expired although this is clearly suboptimal if the only search cost was the explicit cost. Optimal search strategies that account for psychological effort costs would likely involve a reduction in the threshold relative to the standard model, and thus provide another explanation for apparent undersearch. Indeed, existing theories suggest that decision makers are able to optimally use limited cognitive resources, and

that many apparent deviations from optimal decision making can be interpreted as being resource rational – that is, optimal under realistic cognitive constraints (Lieder & Griffiths, 2020; Payne et al., 1993). The effort cost mechanism considered here is consistent with such a theory, and the use of psychological effort costs in the derivation of the optimal threshold can be seen as an example of resource rational analysis.

Finally, as Rapoport and Tversky (1970) already point out, people may use optimal thresholds, but occasionally make mistakes according to which they fail to terminate the search despite observing a reward that exceeds the threshold, or conversely terminate the search preemptively despite observing a reward that does not exceed the threshold. These types of errors are common in cognition and behavior, and have previously been shown to capture systematic deviations from optimality in decision making (Bhatia & Loomes, 2017; Erev et al., 1994; He et al., 2019; Regenwetter et al., 2011). In fact, models that allow for decision slips or some randomness account for behavioral data better than their errorless counterparts (Rieskamp, 2008). Thus, it may also be the case that decision error gives rise to the appearance of undersearch relative to the optimal search strategy.

Each of these three mechanisms could produce undersearching behavior relative to the optimal (expected value maximizing) thresholds. However, prior work has not attempted a quantitative parametric analysis of these mechanisms, either in isolation or in conjunction, and has not examined whether individual-level variability in these mechanisms explains the appearance of undersearch. Thus, despite the fact that many people appear to use threshold-based search strategies (Shotter & Braunstein, 1981; Hey, 1987), it is still unclear what the exact cognitive mechanisms that determine these thresholds are.

Our goal in this paper is to systematically compare an array of possible cognitive models that could account for empirical results on decision thresholds in sequential search with recall.

We test risk aversion, psychological effort cost, and decision error explanations for undersearch using quantitative model fits of threshold rule models to search data from two experiments, and estimate model parameters associated with each of these models at the individual level. Our fits reveal that most decision makers do in fact display risk aversion, effort cost, and decision error. Importantly, models equipped with these mechanisms predict observed search patterns (including undersearch relative to the predictions of the standard error-free risk neutral model without effort costs), suggesting that participants are able to set sensible thresholds in sequential search. These results bridge the gap between computational (optimal) and algorithmic (satisficing) theories of decision making, and by doing so, provide a comprehensive and cohesive account of choice behavior in sequential search.

2. Overview of Experiments and Analyses

Participants in our experiments are offered a choice between a number of different options, with each option containing a monetary reward drawn from a known distribution.

Participants are told about the underlying reward distributions and can search options sequentially, at a fixed monetary cost, in order to learn about the specific reward associated with each option. They can also select any previously sampled option when they decide to terminate search. The monetary costs of search and the reward distributions are identical across all options in a given trial, though we do vary the monetary costs of search across trials. This task is identical to the "search with recall" tasks considered by statisticians and economists (DeGroot, 1970; Telser, 1973) and has been studied experimentally by psychologists and economists (Rapoport & Tversky, 1970; Schotter & Braunstein, 1981).

This simple search task yields an optimal search rule involving a single fixed threshold, that remains identical across all options throughout the search course within a trial. Participants who are capable of searching optimally should search options until they find one whose reward exceeds that threshold. They should then select that option and terminate search. Unlike prior modeling work on sequential search, we allow for risk aversion and psychological effort cost in the decision maker's utility function. This can yield optimal thresholds that are different from those prescribed by the risk neutral search strategy without effort cost. We also allow decision makers to make errors when implementing the search strategy.

We fit a variety of threshold models to individual-level search data from two experiments, in which sampled reward values are presented either numerically (Study 1) or graphically (Study 2). The use of two different experimental designs helps establish the robustness of our results, though due to the similarity in the underlying experimental structure we present the methods and results for these experiments concurrently. Using our best-fit models, we attempt to predict the total number of searches and the obtained reward values in each monetary cost condition and for each participant. These predictions can help us evaluate the degree to which our cognitive models (equipped with risk aversion, effort cost, and decision error) describe undersearch relative to the standard model, as well as the degree to which our models account for the effect of monetary cost on search behavior, and on individual heterogeneity in search behavior. By interpreting the best-fit parameters of our models, we also hope to better understand the degrees of risk aversion, effort cost, and decision error at play in our search task.

Our use of a fairly simple search task is motivated in part by the computational difficulty of fitting complex cognitive models to search behavior. Specifically, our models generate

parameters, and deriving these thresholds for a candidate set of parameter values is computationally very intensive. By restricting our tests to a task in which optimal thresholds remain fixed across all options throughout the trial, and are insensitive to the reward value sampled over the course of the trial, we are able to use various shortcuts and sidestep many of the computational difficulties involved in model fitting. For example, our simple search task allows us to precompute optimal thresholds for a grid of candidate parameter values, and query these precomputed thresholds during model fits. Such shortcuts would not be feasible for more complex tasks, such as search without recall, in which thresholds change dynamically as a function of the search step (e.g., Harrison and Morgan, 1990) or search with side information in which the stopping rule depends on previously sampled information (Weitzman, 1979; Analytis et al., 2014).

Note that prior work has also considered non-threshold models of sequential search. One such model is the *N-search* rule, which proposes that decision makers search a fixed *N* number of options and select the most rewarding option encountered among them. This model goes back to a seminal paper by Stigler (1961), in which he used a search model to define the economic value of information. In Stigler's problem, however, people choose once and for all the exact number of items they will search (e.g. shortlisting job candidates). Thus, to select optimally in that strict setting people need to choose the right *N* in the *N-search* rule (see Lippman & McCall, 1976), a strategy that is suboptimal in sequential search settings. Another model is the *N-bounce* rule, which proposes that decision makers search at least *N* options and then terminate search if the values of searched options are higher than previously searched option(s) (Hey, 1982, 1987). These models have been found to be poor predictors of search termination in sequential search

settings (Hey, 1982, 1987; Schotter & Braunstein, 1981; Sonnemans, 2000). Thus, we limit our tests to only thresholds models, which, as discussed above, have straightforward interpretations in terms of both optimal search and heuristic (satisficing) search, and are already considered good descriptive theories of search behavior. To allow for the occasional selection of a previously searched option, we equip our threshold models with a random termination mechanism (outlined in more detail below), but concede that such a mechanism may not provide a complete description of search and selection behavior when decision makers return to options they have sampled previously.

Overall, our work belongs in a stream of research that uses cognitive modeling to investigate search behavior in optimal stopping tasks (Lee, 2006; Guan & Lee, 2018; Guan 2019; Baumann et al., 2020; Song et al., 2019) and sequential decision making tasks more broadly (i.e., Wallsten et al., 2005; Reutskaja et al., 2011), and is the first paper to use quantitative cognitive modeling techniques in search problems with recall. Although different behavioral strategies have been suggested in previous empirical work in search with recall, the authors primarily compared the consistency of people's behavior with these strategies and did not fit models to the behavioral data (i.e., Hey, 1982; Moon & Martin 1990; Sonnemans, 1998; Schunk & Winter 2009). Our work goes above and beyond these previous studies in that it examines strategies based on precise and widely studied behavioral principles and systematically tests these accounts using model fits at the individual level. This approach allows us to estimate individual-level parameters that can account for behavioral variability in the data and can be used to directly compare the behavior of different individuals.

3. Experimental Methods

The search task was an incentivized treasure hunt game. In each trial participants were presented with boxes that contained unknown numbers of gems. They were allowed to open these boxes (at a cost) in order to observe the number of gems, before selecting a box from which they would like to collect gems. The number of gems was presented either in a numerical (Study 1) or graphical (Study 2) format. Raw data for both studies is available at https://osf.io/8rvak/.

3.1. Stimuli

There were 35 available boxes in each trial in each study. Each box contained between 100 and 999 gems, and all possible numbers were equally likely for each box. The cost associated with opening the boxes was constant within each trial. To separate the effects of monetary and psychological effort costs in searching, we manipulated the former within participants and across trials. There were eight levels of costs: 1, 2, 3, 5, 10, 15, 25, and 50 gems per search, and ten trials for each cost level. We randomly generated a set of stimuli for each participant, resulting in 80 trials for each participant in each study. Our experiments also involved a practice trial at the very start. Note that although the number of available boxes was finite, it was large enough that no participants ever clicked through all the boxes.

3.2. Procedure

Before the experiment, participants were informed of the gem number distribution for the boxes (i.e., each box contains 100 to 999 gems, and all integer numbers in this range are equally likely). At the onset of each trial, 35 unopened boxes were presented on the screen, and participants were told about the cost of opening each box in the trial (Figure 1A). To observe the number of gems, participants had to left click on the box. The number of gems in that box would then be revealed for one second. To test the robustness of our results, we implemented two

different designs to reveal the information. In Study 1, the number of gems was presented in a numerical format (Figure 1B). In Study 2, the inside view of the box was displayed in a pop-out window, in which gems were piled graphically in a 37×27 grid (Figure 1C). Participants were allowed to reopen previously opened boxes, and could open as many boxes as they wanted, in whichever order they desired. Throughout the sampling stage, the screen showed the cost of sampling (*Cost per Click* in Figure 1), the total number of boxes opened in the trial (*Click Count* in Figure 1), and the cumulative cost incurred in the trial (*Cumulative Cost* in Figure 1).

Participants indicated their choice by right clicking on their desired box. The number of gems earned in a trial was determined by the number of gems in the chosen box, minus the total (cumulative) cost incurred by opening boxes. At the end of the experiment, one of the trials was selected randomly, and gems earned in that trial were converted into a monetary amount at the rate of 1 gem = USD 0.01. This monetary amount served as a bonus payment for the participant.

3.3. Participants

We recruited 48 participants (58% female, age: M = 26.9, SD = 11.4) for Study 1 and 47 participants (53% female, age: M = 25.3, SD = 8.4) for Study 2, respectively, from a research participant pool of the University of Pennsylvania. Two participants (one in each study) always chose boxes without observing their underlying reward values. There was also one participant in Study 1 who chose without sampling the chosen box in 79 of the 80 trials. It is impossible to fit our models to these participants' data and thus we excluded them from our analysis. This left 46 participants for each experiment.

4. Models

4.1. Standard Model

The EV-optimal threshold value for an option i can be computed by finding a reward value z_i that makes the decision maker indifferent between selecting a previously sampled option offering z_i and searching option i at a cost c_i . If the decision maker chooses not to sample option i, she receives reward z_i deterministically. If the decision maker searches option i, the expected benefit of search (or more specifically, the expected value of search) is a simple function of the probability of finding a reward less than z_i and thus obtaining reward z_i , and the probability of finding a reward greater than z_i and thus obtaining that reward. This expected benefit can be written as $z_i \int_{-\infty}^{z_i} dF_i(x_i) + \int_{z_i}^{\infty} x_i dF_i(x_i)$. Here $F_i(x_i)$ is the probability distribution of the reward in option i. c_i is the cost of searching option i. The EV-optimal threshold value is the value z_i that makes the expected benefit of search equal to the deterministic reward, z_i , obtained if the decision maker chose not to search, minus the cost of search, c_i . This is the value of z_i that sets $\int_{z_i}^{\infty} (x_i - z_i) dF_i(x_i) = c_i$.

In our setting, both the rewards (which are integers drawn from the discrete uniform distribution on the interval [100, 999]) and the costs of search (which are in the set $\{1,2,3,5,10,15,25,50\}$) are identical across options. This implies that the threshold values are identical across options. We derive EV-optimal threshold values using the above formula. Note that we are not guaranteed an integer value z such that the decision maker is exactly indifferent between search and no search. Thus, we take the smallest integer z such that the net utility from search is negative. For costs 1, 2, 3, 5, 10, 15, 25 and 50, this yields optimal thresholds of 958, 940, 927, 905, 866, 836, 788 and 700 gems respectively. We refer to the model described in this section as the *standard model*.

4.2. Alternate Models

The above model has been the basis of a number of experiments on optimal search, and has typically predicted thresholds that are higher than those used by human decision makers. This does not necessarily suggest that people use suboptimal thresholds as the standard model involves a number of restrictions designed to simplify the calculation of the optimal threshold. These restrictions include risk neutrality (according to which the benefit of search is simply the expected value of search) and costless effort (according to which the cost of search is simply the monetary cost). Here we relax these restrictions in order to consider more reasonable variants of the standard model.

The first such variant allows decision makers to display risk aversion or risk seeking, by transforming reward values in potentially non-linear manner. We assume a power function for reward value, such that the value of a reward x is given by $v(x) = x^{\alpha}$. α is a flexible parameter inferred from data, with $\alpha < 1$ resulting in risk aversion, $\alpha > 1$ resulting in risk seeking, and $\alpha = 1$ resulting in risk neutrality (see O'Donoghue and Somerville, 2018). This model satisfies all the prerequisites for the optimal search model. However, it uses expected utility, $v(z_i) \int_{-\infty}^{z_i} dF_i(x_i) + \int_{z_i}^{\infty} v(x_i) dF_i(x_i)$, rather than expected value to calculate the expected benefit of search. This model also uses the non-linear value function to measure both the value of the deterministic reward if the decision maker chooses not to search, $v(z_i)$, as well as the cost of search, $v(c_i)$. We refer to this model as the *risk preference model*.

Our second variant allows decision makers to incur an additional psychological effort cost for search. We assume that this is a fixed positive cost that combines additively with the monetary search cost to determine the overall cost of search. We write this effort cost as γ and assume that γ is the same for each option and remains constant throughout search (thus the total effort cost is linear in the number of options searched). γ is a flexible parameter inferred from

data. This model again satisfies all the prerequisites of the optimal search model except that it uses the combined monetary cost of search plus effort cost of search, $c_i + \gamma$, to calculate the threshold value. Note that this specification of effort cost implies that γ is interpretable in terms of a monetary value – specifically, an implicit monetary amount that an additional search operation costs the decision maker. We refer to this model as the *effort cost model*.

Our third variant combines the above two assumptions into a single model which assumes that the decision maker displays both a risk aversion or risk seeking, as well as a psychological effort cost. Here we again use expected utility (with parameter α) rather than expected value to measure the value of rewards and costs, and combine the non-linear valued cost with the effort cost, $v(c_i) + \gamma$, to determine the expected benefit of search and the optimal threshold. Both α and γ are flexible parameters inferred from data. We refer to this model as the *full model*.

We also consider two additional non-optimal models which assume that thresholds are not dependent on monetary costs. The first of these models posits a fixed threshold value τ and merely checks whether or not an option's reward exceeds τ . If it does, then the decision maker ends search and chooses that option. τ is a flexible parameter inferred from data. We refer to this model as the *fixed threshold model*. The second of these models starts off with a threshold value τ but then linearly changes τ with each search step. After t searches, the threshold value used is τ - δ ·t. Again, the decision maker ends search and chooses an option as soon as its value exceeds τ - δ ·t. We refer to this model as the *time-dependent threshold* model. When δ > 1, this model resembles the collapsing threshold model previously proposed for two-alternative-forced-choice (Hawkins, et al., 2015). Collapsing thresholds are optimal in search tasks without recall (see Degroot, 1970; Harrison and Morgan, 1990), and when the cost of search is increasing in time.

An increase in the psychological cost of search could be induced by impatience or boredom, in which case it would be adaptive for decision makers to use a collapsing threshold.

Note that the full model nests the risk preference and effort costs models, which can be obtained by setting $\alpha=1$ or $\gamma=0$ respectively. The full model, the risk preference model, and the effort cost model nest the standard model, which can be obtained by setting the free parameters of these models to $\alpha=1$ and/or $\gamma=0$. The fixed threshold and time-dependent threshold models are not nested by any of the four optimal threshold models.

4.3. Error Specifications

The above models are deterministic, whereas individual choices are typically stochastic. Thus, we also assume some degree of noise or error that leads to occasional deviations from the intended choice. We consider two types of error. The first follows a Gaussian distribution with a zero-mean and standard deviation σ . According to this specification, the decision maker probabilistically chooses a sampled option if its reward value x exceeds the derived threshold z plus error ε . As ε is distributed normally, the probability of choosing an option with a sampled value of x is given simply by G(x-z), where G is the cumulative distribution function for the centered Gaussian distribution with standard deviation σ .

The second error specification involves tremble noise. According to this specification, the decision maker has a fixed probability, ρ , of making an error. Thus with probability $1-\rho$, the decision maker chooses the sampled option if $x \ge z$, and chooses to continue search is x < z. However, with probability ρ the decision maker does the inverse of this, that is she chooses the sampled option if x < z, and chooses to continue search if $x \ge z$.

Both the above specifications allow for the imperfect execution of the potentially optimal search strategy. The Gaussian noise specification assumes that the probability of this error

depends on the size difference between the sampled reward x, and the threshold z, whereas the tremble specification assumes that this probability is independent of the rewards and thresholds. In both cases, the probability of error is assumed to be independent across searches within a given trial (as well as across trials). Finally, both error specifications use the thresholds specified by the models present in the section above. For example, the Gaussian error specification applied to the full model would predict that decision makers derive optimal thresholds according to a power value function with effort cost, but may occasionally make mistakes in using these thresholds in each search. It is also useful to note that both types do not involve a bias favoring undersearch or oversearch relative to the threshold.

In addition to the Gaussian and tremble specifications, we need one additional source of error according to which the decision maker can terminate the decision and select an option that isn't the most recently searched option. Such a *random exit* mechanism will not alter the relative performance of our various models but will allow for the occasional deviations from these models according to which participants choose an option that was searched early on in the trial, or an option that has not even been searched at all. We also specify this random exit mechanism with a trembling noise, which assumes that, with some probability π , the decision maker decides to terminate search and choose one of the options at random. We assume that the random exit decision is made prior to each search operation.

4.4. Model Fitting

With the above assumptions, we obtain a model that makes multiple search (and search termination) decisions prior to eventually selecting an option. These decisions depend on the model's thresholds (which are specified by each of our six models) as well as on stochasticity in executing the threshold-based choice rules (which is specified by our two error specifications),

and on a random search termination mechanism (which does not depend on our models or error specifications). Pseudo-code for implementing the search process for our full model combined with the normally distributed error specification is as follows:

```
Step 1: Compute the optimal threshold, which is a function of the risk aversion parameter, \alpha, the effort cost parameter, \gamma, and the cost of search in the trial, c.
```

Step 2a: Determine with probability π whether you want to terminate search and choose an option at random.

Step 2b: If search is not terminated, search one of the unsearched options.

Step 2c: If the reward in the option is higher than the optimal threshold from step 1, plus random error (determined by parameter σ), then terminate the search and select the option.

Step 2d: If search is not terminated, repeat steps 2a-2d.

A similar process holds for the remaining five models (which vary only in terms of how the threshold is set, in step 1) and the tremble noise specification (which varies in terms of how error is combined with the sampled reward to determine choice, in step 2c).

The six models and two error specifications yield twelve distinct model and error combinations, which we fit on the individual-level using maximum likelihood estimation. This was implemented in Python using the Nelder-Mead routine which minimized the negative log-likelihood of the participant's observations. The likelihood functions used in this estimation took into account all of the information observed from the participant, that is, the sequence of options searched by the participant, whether they terminated after a given search and if so, the box they chose. For example, if a decision maker observed reward values $x_1, x_2, \dots x_n$, after which they terminated the search and chose the nth option, the likelihood (given a set of model parameters)

would be determined by the predicted probability of not terminating search randomly prior to each of the searches, the probability of not terminating the search after having observed the first n-1 reward values, and the probability of terminating the search after having observed the nth reward value.

4.5. Model Properties

Before continuing to the results, it is useful to analyze the thresholds generated by our various models. These can be observed in Figure 2A which plots the relationship between the threshold and the power function parameter, α , and the cost of search c. The predicted thresholds for the standard (risk neutral) model are shown using a white line. The optimal thresholds reduce as the cost of search increases for all risk preference levels. Additionally, optimal thresholds reduce as risk aversion increases, for all cost levels. Intuitively, risk averse decision makers obtain greater utility from the deterministic reward offered by a known (opened) option than from searching novel options, which offer potentially higher but uncertain rewards at an additional cost. Finally, the effect of the psychological effort cost of search can be understood using the cost axis in Figure 2A. Higher effort costs would shift the entire curve to the right on this axis. Intuitively, decision makers who find search to be extremely effortful are likely to use smaller thresholds for all risk preference levels.

Of course the thresholds used by decision makers are not directly observable in the data. Rather these thresholds must be inferred from the reward values selected by the decision maker. In the absence of decision error, these chosen reward values will typically be a bit higher than the underlying threshold. Intuitively, decision makers in a noiseless model never select options with reward values lower than the threshold, which implies that chosen reward values are distributed between the threshold value and the maximum possible reward. Thus, while the thresholds

generated by the noiseless standard model for costs 1, 2, 3, 5, 10, 15, 25 and 50, are 958, 940, 927, 905, 866, 836, 788 and 700 respectively, the expected chosen reward values for these costs are 978.55, 969.57, 962.99, 962.33, 932.94, 917.92, 892.90, and 848.00 respectively.

The deviation between the chosen reward value and the underlying threshold tends to diminish, and eventually reverse, in the presence of decision error. While noiseless decision makers only select an option if its reward value is higher than the threshold, noisy decision makers can occasionally make mistakes and select options with reward values lower than the threshold. The probability of making such a mistake at a given point in the trial is equivalent to the probability of making the opposite mistake (i.e., not selecting an option whose reward value is higher than the threshold), as both Gaussian and tremble noise are unsystematic. Nonetheless, the fact that search is terminated if a threshold is crossed but not terminated if it isn't implies that noisy decision makers are more likely to terminate search with the selection of a suboptimally small reward value.

As an illustration, consider a simple tremble noise specification with probability $\rho = 0.25$ of making an error, applied to a model with threshold z = 900. In our experiment, in which reward values x are integers sampled from the discrete uniform distribution on the interval [100,999], there is a probability of 1/9 that $x \ge z$ and a probability of 8/9 that x < z. Subsequently there is a probability of 8/9*0.25 = 0.22 that the decision maker samples and incorrectly selects an undesirable reward (x < z), and a probability of only 1/9*0.75 = 0.08 that the decision maker samples and correctly selects a desirable reward $(x \ge z)$, after the first option is searched. With probability 0.70, the search continues, and the same odds are generated for the second searched box. With the parameters used here, the probability of exiting with a suboptimally small reward is fairly high, so that decision makers end up with rewards below their thresholds. Indeed, the

average reward value obtained by the decision maker in this setting is only 499.66, which is much lower than the threshold value of 900.

The properties discussed here yield analogous patterns for the expected number of searches performed by the decision maker. Overall, the expected number of searches for the noiseless standard model for costs 1, 2, 3, 5, 10, 15, 25 and 50, are 22.14, 15.69, 12.66, 12.55, 6.73, 5.62, 4.32, and 2.99 respectively. Decision error, which results in the selection of suboptimally small rewards and the premature termination of search, yields fewer expected searches and can give rise to the appearance of undersearch relative to the standard model. Of course, higher risk aversion and (effort or monetary) cost, which result in lower thresholds, also yield fewer expected searches, and thus a high proportion of trials in which the decision maker undersearches relative to the standard model.

5. Results

5.1. Behavioral Patterns

Participants opened an average of 3.90 options in Study 1 (SD = 5.69) and 3.44 options in Study 2 (SD = 3.91), and the chosen options had an average reward value of 818.74 (SD = 173.02) gems in Study 1 and 803.61 (SD = 184.80) gems in Study 2 (resulting in an average earning of 782.51 gems in Study 1 and 772.79 gems in Study 2). The standard (EV-optimal) model without any error would open an average of 10.32 options and choose options with an average reward value of 933.15 (earning an average payment of 862.72), suggesting that participants did not search in accordance with the noiseless standard model in our experiments.

Indeed, participants undersearched (selected options with rewards lower than the threshold specified by the expected value maximizing optimal thresholds) in 53% of trials in Study 1 and 60% of trials in Study 2. In contrast participants oversearched in only 3% of trials in

Study 1 and only 2% of trials in Study 2. Thus, they terminated search in accordance with the noiseless standard model in only 44% of trials in Study 1 and 38% of trials in Study 1. Overall, 59% of participants undersearched in at least half the trials in Study 1, and 65% of participants undersearched in at least half the trials in Study 2. These results are largely consistent with prior work documenting undersearch relative to the EV-optimal model in sequential search tasks.

In Figures 3A and B, we provide summary statistics for the obtained reward values and number of searches for each cost condition separately, aggregated over participants, for each experiment. We also present the predictions made by the noiseless standard model for these variables. Here we can see that although participants are sensitive to cost, and thus open more boxes and obtain higher rewards in lower cost conditions, there is substantial undersearch for all cost conditions in both experiments.

The above analysis pertains to aggregate search data. However, there is substantial heterogeneity in individual search and decision behavior. This is shown in Figures 4A and B which display the distribution of the average chosen reward values of participants and the average number of options opened by each participant for the two studies. Most participants search less and obtain lower rewards than predicted by the noiseless standard model, though the degree of undersearch varies substantially across participants, with some participants searching as few as 0.31 boxes and others searching as many as 7.87 boxes, on average.

Although participants may not always have been using thresholds based on the standard model, they do appear to have been using some kind of threshold rule. Particularly, in line with the predictions of threshold models, 86% of trials ended with the selection of the last sampled option. Alternate non-threshold decision rules would predict that the last sampled option would be chosen much less frequently. For example, an *N-search* rule, which proposes that decision

makers search a fixed N number of options before selecting the most rewarding option encountered in the search, predicts that the last sampled option would only be chosen 1/N times (Stigler, 1961; Lippman & McCall, 1976). Thus, in order to capture the fact that participants open 3.90 options in a given trial, such models would have to predict that the last sampled option is chosen only $1/3.90 \approx 25\%$ of the time.

Again, our finding that people appear to be using some kind of threshold rule (rather than non-threshold rules such as N-search) is aligned with prior work, which has found that threshold models are consistent with the behavior of a large proportion of participants in search studies (Hey, 1982; Sonemanns, 1998, 2000; Schunk & Winter, 2009). Also note that although all the models that we consider in this paper are threshold models, they still predict that searches will occasionally terminate without the selection of the last sampled option, which is consistent with previous findings in the empirical literature (i.e., Kogut, 1992; Reutskaja et al., 2011). This is due to the random exit mechanism, which terminates the decision and selects an option at random with probability π . Thus our threshold models can capture observed (infrequent) deviations from threshold decision making.

5.2. Model Fits

Our goal is to understand the ways in which participants set thresholds, and to evaluate whether or not optimal threshold models (equipped with risk aversion, psychological effort cost, and decision noise) account for observed patterns of undersearch. Thus we fit the twelve models (composed of six threshold rules with two error specifications) described in the prior section to participant-level data. The results of these fits are shown in Table 1, which displays Akaike Information Criterion (AIC, Akaike, 1974) values of the models in the two experiments. The lower the AIC values, the better the models fit the individual-level search data with model

complexity accounted for by means of the number of free parameters. These values are obtained by summing the AICs of the models over all the participants, and thus describe the aggregate fits of the models on the data. Table 1 also shows the proportion of participants that are best described by a given model (based on that model's participant-specific AIC).

Our fits reveal that both the Gaussian and tremble standard models, which assume expected value maximization without psychological effort cost (but with either Gaussian or tremble noise), perform very poorly in the two studies. These models have the highest AIC values of all models considered and cumulatively best describe 0% of participants in Study 1, and 4% of participants in Study 2. Thus participants in our study do not appear to be behaving in accordance with the strict notion of (expected value maximizing) optimality considered in prior work, even if this notion of optimality is equipped with random error. However, participants also appear to be sensitive to the costs of search, as revealed by the poor fits of the fixed threshold model. The Gaussian and tremble noise specifications of this model have very high AIC values, which are larger than those for all other models except for the standard model. Collectively the Gaussian and tremble noise specifications of the fixed threshold model best describe only 7% of participants in Study 1 and 2% of participants in Study 2. Note that the timedependent threshold variant of the fixed threshold model (which assumes that thresholds are insensitive to cost, but sensitive to the time spent searching) does not do much better. This model does have a lower AIC value than the fixed threshold model, but has larger AIC values than the remaining models. The Gaussian and tremble noise specifications of the time-dependent threshold model best describe only 9% of participants in Study 1 and 11% of participants in Study 2.

Overall, the model that performs the best in both studies, based on aggregate AIC, is the full model with the Gaussian noise specification. This is followed by the risk preference model and effort cost model with the Gaussian error specifications, which display fairly similar performance. In terms of individual model fits, we find that the Gaussian risk preference model best describes the most participants (43%) in Study 1, followed by the Gaussian full model (26%). In Study 2, we find that the Gaussian effort cost model best describes the most participants (28%) followed by the Gaussian full model (26%). The tremble versions of these three models describe a small fraction of participants (between 0% and 2%) in the two studies.

Table 1 also shows that all Gaussian noise models outperform their corresponding tremble noise models. Overall, tremble models best describe fewer than 10% of participants in Study 1 and Study 2. Thus, on the basis of these results it seems that Gaussian noise, in combination with risk preference and effort cost, best describes participant search behavior.

Finally, note that AIC assumes that the penalty for flexibility is linear in the number of parameters. This provides a very rough approximation of the models' true flexibilities. It is thus necessary to validate the results in Table 1 using a second metric of model fit. For this we can exploit the fact that all our optimal models nest the standard model as a special case, and that it is possible to recover the standard model if we constrain the parameter values of our full, risk preference, and effort cost models by setting $\alpha = 1$ and/or $\gamma = 0$. This means that we can test whether the additional assumptions of these models improve model performance using the likelihood ratio test (LRT). We implemented the LRT on the individual-level for each pair of nested and non-nested optimal models in Study 1 and Study 2. The results of this test are shown in Table 2. Here we can see that the Gaussian risk preference, effort cost, and full models significantly (with p < 0.05) outperformed the Gaussian standard model for more than 80% of

participants, and that the tremble risk preference, effort cost, and full models significantly outperformed the tremble standard model for more than 90% of participants. The full models also significantly outperformed their corresponding risk preference and effort cost models for a sizable subset of participants. These tests indicate that assuming risk preference and effort cost significantly improves model performance.

5.3. Best-fit Parameters

Best-fit parameters associated with the above model fits can shed light on the participant preferences that give rise to observed search patterns in the two studies. The median values of these parameters are shown in Table 1. Here we can see that on aggregate, participants were risk averse, with the median participant having an α of 0.89 in the full model with Gaussian error in Study 1 and 0.88 in Study 2. Overall, 72% of participants in Study 1 and 76% of participants in Study 2 had best-fit α < 1 according to this model. Participants also displayed positive effort cost, with the median participant having a γ of 2.59 in the full model with Gaussian error in Study 1 and 5.12 in Study 2. γ is directly interpretable in terms of a monetary value, and these fits indicate that implicitly the cost of search is 2.59 gems (or USD 0.0259) and 5.12 gems (or USD 0.0512), higher than the actual monetary cost in many of the cost conditions.

Note that the degrees of risk aversion and effort cost observed in our data change slightly with the use of the tremble noise specification: models fit with this type of error typically have slightly lower risk aversion and higher effort costs than their Gaussian counterparts. The degree of risk aversion and effort cost also changes with the core model. For example, the Gaussian risk preference model (which does not have effort cost) displays a higher degree of risk aversion than the full Gaussian model (which has both risk preference and effort cost). Similarly, the Gaussian effort cost model (which does not have risk preference) displays a higher degree of effort cost

than the full Gaussian model. Both risk aversion and effort cost lead to undersearch, so that a model that turns off one mechanism compensates by increasing the strength of the other.

What kinds of decision thresholds do our best-fit parameters generate? This is shown in Figures 2B, which plots the median optimal thresholds according to participant best-fit risk aversion and effort cost parameter values for five out of our six models with the Gaussian error specification (we do not plot the time-dependent thresholds as these vary based on the number of searches made in a given trial). Here we see that all models, except the fixed threshold model, have smaller thresholds than the standard model in all cost conditions (the fixed threshold model has a lower threshold for smaller cost conditions but a higher threshold for the higher cost conditions). All models, except the fixed threshold model, also display cost sensitivity, with lower thresholds for higher costs. The risk preference and effort cost models do generate similar thresholds, though the risk preference model has relatively lower thresholds for larger costs, whereas the effort cost model has relatively lower thresholds for smaller costs. Intuitively, risk averse decision makers tend to be less likely to search as the cost of search (and subsequently, the loss from not finding a desirable reward) increases. This yields relatively lower thresholds for larger cost conditions. In contrast, the effort cost model imposes a fixed additive increase to the disutility from search, which has a proportionally larger impact when the monetary cost of search is small. This yields relatively lower thresholds for lower cost conditions. Note that the full model, which permits both mechanisms, generates intermediate sized thresholds which lie between the thresholds of the risk preference and effort cost models.

5.4. Capturing Behavioral Patterns

Although Table 1 shows that our assumptions of risk preference and effort cost in the full model improve our ability to describe behavior, it does not provide an intuitive interpretation of

how accurate our predictions are. Instead, the full model's ability to capture participants' behavioral patterns is shown in Figures 3-5. In Figures 3A-B we plot the average number of searches and reward values predicted by the best-fit full model with Gaussian noise, for each cost condition. We can see that the full model closely approximates search behavior in both studies: It captures both undersearch relative to the noiseless standard model and reduction in search with increases in monetary search cost. Overall, the mean absolute error of the full model's predictions of the number of searches is 0.34 in Study 1 and 0.47 in Study 2. This is much smaller than the analogous error generated by the noiseless standard model, which is 6.43 in Study 1 and 6.88 in Study 2. Likewise, the mean absolute error of the full model's predictions of the obtained reward values is 18.38 in Study 1 and 22.77 in Study 2. This is again much smaller than the analogous error generated by the noiseless standard model, which is 114.62 in Study 1 and 129.82 in Study 2. Note that the full model does somewhat under predict the number searches and the chosen rewards for very small cost values. This could be due to a number of different additional psychological mechanisms not contained in our models; a point we elaborate on in the discussion section below.

Figures 4A-B show that we are also able to describe participant heterogeneity in search. Here we display the average number of searches and reward values predicted by the best-fit full model with Gaussian noise, for each participant. These are very close to actual participant data. Overall, the correlation between the observed and predicted number of searches is 0.97 in Study 1 and 0.98 in Study 2. Likewise, the correlation between observed and predicted reward values is 0.95 in Study 1 and 0.96 in Study 2. Note that the standard model is unable to accommodate individual-level variability as it assumes that all participants are risk neutral and have zero effort costs.

The full model also predicts search termination accurately. We illustrate this in Figure 5, which presents predicted probabilities for terminating search both when participants terminated search (which we refer to as the "hit rate") and when they did not (the "false alarm rate"), for each participant. To obtain these values, we first calculated the best-fit full model's predicted probability of search termination after each search made by each participant in each trial. This probability would merely be G(x-z), where G is the CDF for the Gaussian distribution, x is the observed value and z is the optimal threshold for that participant and cost condition generated by the best-fit full model. We then averaged these probabilities for each participant to get participant-specific predicted probabilities of terminating search, when search was actually terminated by the participant and when it was not respectively, that is the hit rate and the false alarm rate for the participant. The average predicted probability of terminating search when search was actually terminated (hit rate) is 80% in Study 1 and 82% in Study 2. Similarly, the average predicted probability of terminating search when search was not terminated (false alarm rate) is 8% in both Study 1 and Study 2, indicating that the best-fitting full model can predict what decision makers will do with high accuracy.

Figure 5 also shows heterogeneity in our ability to predict participant data. Although the majority of participants are in the top left corner of the scatterplots, indicating that hit rates are high and false alarm rates are low, there is a small minority of participants for whom our predictions are less accurate. These participants are typically those that are poorly fit by the full model. Indeed, we find that there is a correlation of 0.15 and 0.28 between the full model's AIC value and the false alarm rate in Study 1 and 2 respectively. Likewise, there is a correlation of -0.66 and -0.40 between the full model's AIC value and the hit rate in Study 1 and 2 respectively.

6. Discussion

6.1. Redefining Optimal Search

We studied people's behavior in a classic dynamic decision making problem, referred to as optimal stopping with recall. In this problem, decision makers can choose between a large number of options, and have to sequentially sample the options to learn their exact utility by paying a non-negligible search cost. This problem is characterized by a clear-cut optimal solution, which is defined by a predetermined stopping threshold: Decision makers should continue sampling options until they encounter one with a utility higher than the stopping threshold. Prior research has found that participants appear to undersearch relative to the "optimal" model in sequential search tasks (Rapoport & Tversky, 1970; Shotter & Braunstein, 1981; Kogut, 1992). This work has used a narrow definition of optimality, which assumes risk neutrality, costless psychological effort, and error-free decision making. Consistent with these existing results we have found that participants in our experiments do in fact search fewer options and obtain lower reward values than the risk neutral, effortless and error-free optimal model.

However, these results do not provide conclusive evidence for suboptimal thresholds. Rather, we find that participants behave in accordance with an optimal threshold model equipped with risk aversion, psychological effort cost, and decision error. Such a model achieves the best-fits to participant-level search data, and can accurately predict search behavior, including undersearch relative to the standard (risk neutral, effortless, error-free) model. It also captures the effect of monetary cost on the number of searches and obtained reward values, as well as participant heterogeneity in the extent of search and obtained reward. Best-fit parameter values of this model indicate that most participants are in fact risk averse and are subject to substantial effort costs (equivalent to about USD 0.02 to 0.05 per search, for the median participant) and

Gaussian (rather than tremble) decision noise. The decision thresholds generated by the risk averse model with costly psychological effort are lower than those prescribed by the standard model, but are still rational given an expanded notion of optimality that allows for risk preferences and cognitive constraints.

6.2. Algorithmic Models of Search

The threshold rule models considered in this paper are not only optimal strategies for making sequential search decisions. They are also an instantiation of the satisficing heuristic proposed by Simon (1955). Threshold rules in optimal stopping problems with recall are modest in terms of computational complexity and it is possible to implement them using simple cognitive architectures. In our work we show that people's undersearch can be interpreted as the product of a boundedly rational decision strategy. This result bridges the gap between computational and algorithmic theories of search behavior. By explaining search behavior using an optimal model our paper attempts a form of rational analysis (Anderson, 1991; Chater & Oaksford, 1999). This analysis reveals that participants can be described as stochastic risk averse utility maximizers that are able to tradeoff potential monetary rewards against both monetary and psychological search costs. The inclusion of psychological effort cost within a utility maximizing function additionally indicates that participants are displaying resource rationality (Lieder & Griffiths, 2020; Payne et al., 1993), which involves optimizing behavior in the face of cognitive constraints.

6.3. Theories of Error and Behavior

The study of perceptual and decision errors has been one of the main motivating forces in cognitive psychology and judgment and decision making (Erev et al., 1994; Loomes, 2015; Reason, 1990). Already in the first studies of sequential search, Rapoport and Tversky (1970)

pointed out the possibility of decision errors in the tasks. They labelled failing to stop on a satisficing offer as Type I error and early stopping as Type II error. They also noted that although people can commit a Type I error more than once in any given trial, a Type II error is irrevocable. An error theory may thus provide an explanation for why people, on average, stop early, yet such a theory has not been tested concurrently with alternative accounts such as risk aversion. Here we contrast two potential theories of error that are commonly used in the behavioral sciences. In the first, we assume Gaussian noise in the way people experience the magnitudes of rewards. This theory or error is commonplace in psychophysics and goes at least back to Thurstone's law of comparative judgment (1927). The second theory assumes a tremble error --- a common assumption in many game theoretical and cognitive models (e.g., Selten, 1975). Note that theories of error are not only prevalent in the behavioral sciences, but also necessary for the design of plausible cognitive models. Even in very simple choice tasks people's behavior is known to be stochastic (see Loomes, 2015). As a result, deterministic models fit behavioral data rather poorly (Rieskamp, 2008). Our analyses suggest that error theories operate in conjunction with other behavioral mechanisms such as risk aversion and effort cost. Further, our findings show that the Gaussian error specification explains people's behavior better than trembling hand error when combined with the risk preference and effort cost mechanisms.

6.4. Sequential Search, Dynamic Decision Making and Risk Aversion

Our experiments also show that risk aversion is one of the main factors guiding people's behavior in sequential search tasks. Although risk preferences have been studied extensively in static tasks using binary or multi-alternative gambles, there is little work on transferring theoretical and empirical insights obtained from static gambles to dynamic decision making problems. There are a few notable exceptions, such as Wallsten et al.'s (2005) study on people's

behavior in the Balloon Analogue Response Task, which finds that accounting for risk preferences increases the explanatory power of cognitive models.

Why is there a scarcity of cognitive models that take into account risk preferences in dynamic decision making tasks? In many dynamic settings, computational complexity makes calculating the expected value optimal strategy computationally challenging or intractable. Calculating an expected utility optimal strategy is even more complex and often cognitively implausible. Therefore, simpler strategies need to be used to solve the tasks (see Analytis et al, 2019; also see Wilson et al. 2014 and Speekenbrink & Konstantinidis, 2015 for a similar discussion for bandit tasks). Although optimal stopping tasks with recall have a dynamic sequential structure, calculating the optimal and risk averse solution is computationally tractable and psychologically plausible. Therefore, building cognitive models of people's behavior in these tasks could provide a useful domain for testing people risk preferences in dynamic search problems more generally.

Over the last decades, several authors working on optimal stopping tasks have proposed risk aversion as a potential explanation for undersearch (i.e., Shotter & Braunstein, 1981, Sonnemans, 1998). Nonetheless, only one study by Schunk and Winter (2009) has explicitly explored the role of risk aversion in the task. The authors ran a search experiment with recall and also presented the participants with (i) a lottery task whereby they could elicit risk preferences from choices and (ii) a risk taking questionnaire. Schunk and Winter found no relationship between the risk aversion parameters elicited from the choice task and search behavior, and only a weak relationship between the risk taking measure from the questionnaire and search length. If people's risk taking behavior was consistent across decision making tasks these results would refute risk aversion as an explanation of undersearch. However, in recent experimental work,

Pedroni et al., (2017) have showed that people do not have consistent risk preferences across different behavioral decision making tasks and self-report measures. Thus, it is possible that people are indeed risk averse in sequential search tasks, as shown by the very good model fits in our paper, but their risk preferences in search tasks are only weakly correlated, if at all, with risk aversion in lottery tasks. These findings also highlight why building cognitive models in dynamic decision making tasks where risk preferences are expected to play a role is crucial. If risk preferences are domain specific rather than domain general, we need to find ways of retrieving them directly from the relevant tasks.

6.5. Relative Influences of Mechanisms

Our model fits have found evidence for both risk aversion and psychological effort cost. However, we are unable to accurately measure the relative influence of these two mechanisms, as they both lead to lower thresholds and undersearch relative to the standard model. This is why, for example, the risk preference only and effort cost only models perform very similarly in our two experiments. This is also why estimated parameter values of risk aversion and psychological effort cost vary based on the model specification that is fit to the data (the risk preference only model has a higher degree of risk aversion, and the effort cost only model has a higher degree of effort cost, than the full model; as both risk aversion and effort cost lead to undersearch, a model that turns off one mechanism compensates by increasing the strength of the other). That said, these two mechanisms are not unidentifiable, as risk aversion yields relatively smaller thresholds for higher search costs whereas effort cost yields relatively smaller thresholds for smaller search costs. Future work can use these differences to design a more precise experiment and subsequently measure the relative influence of risk aversion and effort cost with increased accuracy.

6.6. Inter-individual Differences

Consistent with previous work on sequential search, we have found substantial heterogeneity in people's search behaviors (Hey, 1987; Cox & Oaxaca, 1989; Schunk, 2009). One of the strengths of our approach is that we can directly compare the best-fit parameters of our models for different individuals or groups of individuals, rather than merely comparing people's average search length or how consistent they are with the optimal rule. Although we find that most participants are risk averse and experience effort costs and Gaussian decision noise, there is substantial variability across individuals. For example, roughly 80% of participants are risk averse, whereas 20% are risk seeking. This distribution is roughly consistent with results obtained when eliciting risk preferences from choices (e.g., Abdellaoui et al., 2000) or when recovering risk aversion parameters from other sequential decision tasks (e.g., Wallsten et al., 2005). Further, people differ in their psychological cost parameters. Indeed, people's opportunity cost of time is an idiosyncratic variable that we cannot control in the experiment, but can retrieve through cognitive modeling and model fitting. Last but not least, there is variability in people's susceptibility to error, and some individuals are more prone to error than others. This might be the cause of the discrepancy between the average utility of the options selected by our participants and those selected by the full model for high costs of search (see Figure 3).

6.7. Additional Models

In this paper we tested an array of cognitive models of sequential search. We selected these models by looking at previous theoretical work, the empirical regularities in experiments on sequential search with recall and considering the computational resources necessary to fit different models. Still, the list of plausible cognitive models is larger and it is possible that we excluded models that can perform as well as, or even better than the models that we included.

One family of candidate models concerns agents that learn and adapt the threshold dynamically in a trial or across trials (Conlisk, 2003; Goldstein et al., 2020). Such models are particularly valuable in cases where the distribution is not known before-hand, and consequently people lack the information necessary to find the optimal threshold. The uniform distribution that we used in our experiments, however, is easily comprehensible. Therefore, in our tasks people have adequate information and cognitive resources to devise good threshold strategies ex-ante. In fact, Hey (1987) found that in sequential search problems where the distribution is known beforehand most people are described by threshold rules, and only a modest minority by simpler adaptive rules. In addition, Conlisk (2003) has shown that simple learning models eventually converge to the optimal threshold policy. Therefore, even if learning played a role in our task, learning models would converge to threshold models.

A second family of models concerns the sunk cost family: Kogut (1990) designed an experiment in which he introduced a first step with a much larger search cost and showed that people may commit the sunk cost fallacy in search tasks. The sunk-cost fallacy, however, would alter the optimal strategy of the problem (see Schunk & Winter, 2009), which would then need to be calculated by backward induction. It is already known that people are not very good at using backward induction for more than a few time steps (Huys et al., 2015), and therefore we consider this possibility unlikely unless combined with further behavioral mechanisms.

It is also possible that risk aversion captures another underlying mental process (in the same way that risk aversion and psychological cost compensate for each other) that we have not properly identified. For example, Schunk and Winter (2009) found that the loss aversion parameters they elicited from choices correlated with search length. In the past, models based on loss aversion have successfully explained risk averse behaviors in market contexts, such as the

equity premium puzzle (Thaler et al., 1997). In a similar vein, loss aversion might be the cause of apparent risk aversion in our task.

It may also be necessary to modify our implementation of the value function to allow for a peanuts effect, according to which very small monetary amounts are given disproportionately small utilities or disutilities. The peanuts effect has been previously documented in risky decision making (Weber & Chapman, 2005; also see Markovitz, 1952), where it predicts an increase in risk seeking for small-scale gambles as well as a neglect for small losses and costs. In our experiment, a peanuts effect could be responsible for oversearch relative to the full model in small costs conditions (see Figure 3). Here, the disutility of incurring a very small cost of search (e.g. 1 gem, which is to equivalent to USD 0.01) may be smaller than that predicted by the (concave) power value function used for our full model, resulting in the full model under predicting the number of searches and chosen rewards when the cost is very small.

Finally, future work could modify our models to attempt to better predict the (small minority of) trials in which participants selected a previously sampled option. Although our random exit mechanism does permit such selections, it may not characterize the richness in participant behavior. For example, this mechanism predicts that people err systematically by uniformly selecting among previously sampled options, but it is much more likely that people choose among previously sampled options in a more systematic manner, for example, by selecting the most desirable option previously sampled, or by testing a softmax rule where the probability of choosing previously sampled options directly depends on their utility. It could be possible to augment the threshold mechanism with the ability to select a previously sampled option if the recalled reward from such an option (plus noise) surpasses the threshold. We could also utilize collapsing thresholds to allow such a model to return to previously sampled options if

recently sampled options are undesirable. In preliminary tests we did try to fit such models, but found that they overpredicted return (and more generally, overpredicted search termination). However, future work could attempt alternate implementations of these ideas, examining more sophisticated mechanisms for deriving a collapsing threshold, revisiting older modeling ideas such as N-search and N-bounce rules (Stigler, 1961; Hey, 1982) or exploring novel mechanisms. Although beyond the scope of this paper, we hope that future work will consider a larger set of modeling accounts that could be directly assessed against or combined with those that we put forward in our contribution.

6.8. Additional Tasks

We have obtained nearly identical results in our two experiments, which vary in terms of whether reward information is presented numerically or graphically to participants. This suggests that our results are fairly robust. That said, both our experiments involve a simple search task, in which each option has the same reward distribution and cost of search, and participants can select previously sampled options when they wish to terminate search. We have chosen such a task due to the computational complexity of fitting optimal models to individual-level search data. The optimal model in our experiments involves a single threshold that is identical for all options and remains fixed throughout the trial. This allows for computational shortcuts such as precomputed optimal threshold values over a feasible parameter space. That said, with advances in computing technology more intensive fits may soon become feasible. If so, future work should attempt to replicate our analysis for complex search tasks.

For example, we could use insights from this study to predict people's behavior in sequential search tasks with recall where the distribution of rewards is not known to the participants (Rothschild, 1974; Hey,1982,1987) or in minimal exploration-exploitation tasks

(Song et al., 2019; Sang et al.,2020). Similarly, we could contrast our cognitive models with the best fitting models in search tasks without recall, where participants can only select the most recently searched choice option (Harrison & Morgan, 1990; Shu, 2008; Baumann et al., 2020) or in search tasks where people search through the options in a specific order (i.e., Caplin et al., 2011). Although it is likely that risk aversion, psychological effort cost, and decision error play some role in these more complex tasks, it is also possible that behavior deviates from optimality in the presence of complexity. Further exploring this possibility is an important topic for future work.

7. Conclusion

We formulated an array of cognitive models that can capture people's behavior in sequential search problems with recall, and assessed their quality by fitting the models to individual-level behavioral data from two experiments. Our results suggest that most people display risk aversion, psychological cost and Gaussian errors, and that individual-level differences in the parameters corresponding to these psychological variables predict differences in the extent of search. Thus, our results resolve the undersearch conundrum in sequential search with recall by expanding the bounds of rationality, and open new directions for future research.

Acknowledgements

We would like to thank Alexandros Gelastopoulos for his insightful comments.

References

- Abdellaoui, M. (2000). Parameter-free elicitation of utility and probability weighting functions. *Management Science*, 46(11), 1497-1512.
- Analytis, P. P., Kothiyal, A., & Katsikopoulos, K. V. (2014). Multi-attribute utility models as cognitive search engines. *Judgment and Decision Making*, *9*(5), 403-419.
- Analytis, P. P., Wu, C. M., & Gelastopoulos, A. (2019). Make-or-Break: Chasing Risky Goals or Settling for Safe Rewards?. *Cognitive science*, *43*(7), e12743.
- Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14(3), 471-485.
- Arrow, K. J. (1965). Aspects of the theory of risk bearing. *The Theory of Risk Aversion*. Helsinki: Yrjo Jahnssonin Saatio.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716-723.
- 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.
- Bearden, J. N., Murphy, R. O., & Rapoport, A. (2005). A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology*, 49(5), 410-422.
- Bhatia, S., & Loomes, G. (2017). Noisy preferences in risky choice: A cautionary note. *Psychological Review*, 124(5), 678.

- Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, 23(3), 251-263
- Caplin, A., Dean, M., & Martin, D. (2011). Search and satisficing. *American Economic Review*, 101(7), 2899-2922.
- Chater, N., and Oaksford, M., (1999), Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences* 3(2), 57-65.
- Conlisk, J. (2003). A note on convergence of adaptive satisficing to optimal stopping. *Journal of Political Economy*, 111(6), 1353-1360.
- 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
- DeGroot, M. H. (1970). Optimal Statistical Decisions. John Wiley & Sons.
- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and underconfidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519-527.
- Harrison, G. W., & Morgan, P. (1990). Search intensity in experiments. *The Economic Journal*, 100(401), 478-486.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015).

 Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision making. *Journal of Neuroscience*, 35(6), 2476-2484.
- He, L., Golman, R., & Bhatia, S. (2019). Variable time preference. *Cognitive Psychology*, 111, 53-79.
- Hey, J. D. (1981). Are optimal search rules reasonable? And vice versa?(And does it matter anyway?). *Journal of Economic Behavior & Organization*, 2(1), 47-70.

- Hey, J. D. (1982). Search for rules for search. *Journal of Economic Behavior & Organization*, 3(1), 65-81.
- Hey, J. D. (1987). Still searching. Journal of Economic Behavior & Organization, 8(1), 137-144.
- Huys, Q. J., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., ... & Roiser, J. P.(2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences*, 112(10), 3098-3103.
- Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4), 1043-1068.
- Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, 32(31), 10686-10698.
- Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339.
- Guan, M. (2019). A Cognitive Modeling Analysis of Risk in Sequential Choice Tasks (Doctoral dissertation, UC Irvine).
- Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, 66(3), 1375-1394.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, *14*(3), 381-392.
- Kogut, C. A. (1992). Recall in consumer search. *Journal of Economic Behavior & Organization*, 17(1), 141-151.

- Lee, M.D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science* 30 (3), 1-26.
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, 1-60.
- Lippman, S. A., & McCall, J. J. (1976). The economics of job search: A survey. *Economic Inquiry*, 14(2), 155-189.
- Loomes, G. (2015). Variability, noise, and error in decision making under risk. *The Wiley Blackwell handbook of judgment and decision making*, 2, 658-695.
- Markowitz H. (1952). The utility of wealth. *Journal of Political Economy*. 60(2), 151–158.
- Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. New York, NY, USA: Henry Holt and Co., Inc.
- Moon, P., & Martin, A. (1990). Better heuristics for economic search—experimental and simulation evidence. *Journal of Behavioral Decision Making*, *3*(3), 175-193.
- Nachman, D. C. (1972). On risk aversion and optimal stopping. *unpublished paper,*Northwestern University.
- O'Donoghue, T., & Somerville, J. (2018). Modeling risk aversion in economics. *Journal of Economic Perspectives*, 32(2), 91-114.
- 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.

- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task.

 Organizational Behavior and Human Performance, 5(2), 105-120.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333.
- Reason, J. (1990). Human error. Cambridge University Press.
- Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review*, 118(1), 42-56.
- 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.
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1446-1465.
- Rothschild, M. (1974). Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4), 689-711.
- 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).
- 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.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1), 25-55.

- Shu, S. B. (2008). Future-biased search: the quest for the ideal. *Journal of Behavioral Decision Making*, 21(4), 352-377.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99-118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Simon, H. A. (1982). Models of Bounded Rationality: Empirically Grounded Economic Reason.

 MIT Press.
- Song, M., Bnaya, Z., & Ma, W. J. (2019). Sources of suboptimality in a minimalistic explore–exploit task. *Nature Human Behaviour*, *3*(4), 361-368.
- Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, 35(3), 309-332.
- Sonnemans, J. (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, 21(1), 91-102.
- 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.
- Schunk, D., & Winter, J. (2009). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization*, 71(2), 347-360.
- Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2), 351-367.

- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213-225.
- Telser, L. G. (1973). Searching for the lowest price. *The American Economic Review*, 63(2), 40-49.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112(2), 647-661.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34(4), 273-286.
- 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 B. J. and Chapman G. B. (2005). Playing for peanuts: Why is risk seeking more common for low-stakes gambles? *Organizational Behavior and Human Decision Processes*, 97, 31–46.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641-654.
- 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.

Table 1. Summary statistics of model fits. AIC corresponds to the summed AIC values across participants. % Best corresponds to the proportion of participants for whom a given model has the lowest AIC value. The lowest AIC and highest % Best values are bolded. Parameter values presented correspond to the median best-fit parameter values across participants. Note that parameters that are constrained in a given model are indicated with squared brackets.

Study 1									
Model	σ	ρ	α	γ	τ	δ	π	AIC	% Best
Fixed threshold G	111.86	-	-	-	811.38	[0]	0.03	15,679	0.07
Time-dependent threshold G	108.08	-	-	-	797.75	-1.53	0.03	15,562	0.09
Standard Model G	158.01	-	[1]	[0]	-	-	0.03	16,846	0
Risk Preference G	102.98	-	0.81	[0]	-	-	0.03	14,560	0.43
Effort Cost G	93.64	-	[1]	11.49	-	-	0.03	14,552	0.11
Full Model G	89.85	-	0.89	2.59	-	-	0.03	14,207	0.26
Fixed threshold T	-	0.08	-	-	839.27	[0]	0.03	17,256	0
Time-dependent threshold T	-	-	-	-	836.18	0.87	0.03	17,128	0
Standard Model T	-	0.12	[1]	[0]	-	-	0.03	19,118	0
Risk Preference T	-	0.07	0.81	[0]	-	-	0.03	16,400	0.02
Effort Cost T	-	0.07	[1]	7.26	-	-	0.03	16,495	0
Full Model T	-	0.06	0.97	7.73	-	-	0.03	15,861	0.02

Study 2									
Model	σ	ρ	α	γ	τ	δ	π	AIC	% Best
Fixed threshold G	102.66	-	-	-	797.66	[0]	0.05	16,778	0
Time-dependent threshold G	97.71	-	-	-	783.67	-2.81	0.05	16,589	0.09
Standard Model G	178.27	-	[1]	[0]	-	-	0.05	17,406	0.04
Risk Preference G	94.38	-	0.76	[0]	-	-	0.05	15,896	0.24
Effort Cost G	90.62	-	[1]	15.10	-	-	0.05	15,735	0.28
Full Model G	84.46	-	0.88	5.12	-	-	0.05	15,558	0.26
Fixed threshold T	-	0.07	-	-	798.23	[0]	0.05	18,194	0.02
Time-dependent threshold T	-	0.08	-	-	797.67	0.84	0.05	18,165	0.02
Standard Model T	-	0.12	[1]	[0]	-	-	0.05	20,776	0
Risk Preference T	-	0.06	0.77	[0]	-	-	0.05	17,532	0.02
Effort Cost T	-	0.07	[1]	12.81	-	-	0.05	17,499	0.02
Full Model T	-	0.05	0.93	9.22	-	-	0.05	16,959	0

Note. G = Gaussian error, T = Trembling error.

RP

EC

70%

80%

Table 2. Proportion of participants whose data are significantly (with p < 0.05) better fit by an encompassing model relative to a nested model as evaluated by a likelihood ratio test. Here encompassing models are shown in columns and nested models are shown in rows, and each cell captures the proportion of participants for whom the encompassing model has a significantly higher likelihood than the nested model for that model pair. RP refers to the risk preference model and EC refers to the effort cost model.

Study 1										
Gaussian					Tremble					
	RP	EC	Full		RP	EC	Full			
Standard	89%	80%	89%	Standard	97%	98%	97%			
RP	-	-	39%	RP	-	-	61%			
EC	-	-	57%	EC	-	-	78%			
Study 2										
Gaussian					Tremble					
	RP	EC	Full		RP	EC	Full			
Standard	89%	85%	96%	Standard	98%	93%	96%			

RP

EC

56%

46%

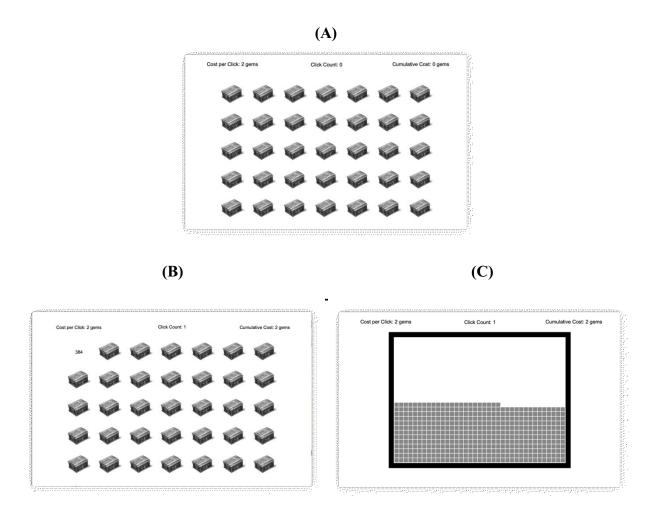
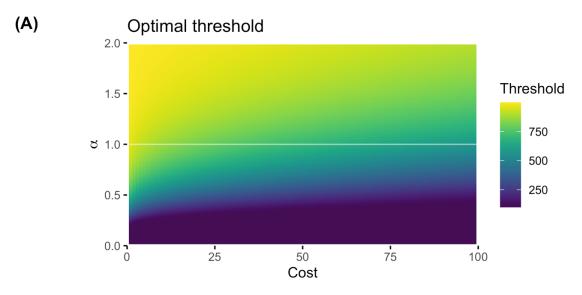


Figure 1. Stimuli in the two studies. (**A**) There are 35 available boxes in each trial, each of which contains some unknown number of gems. The cost associated with opening a box is shown at the top left corner of the screen. To the right of it, trackers are displayed for the total number of boxes opened in a trial, as well as the cumulative search cost incurred in that trial. (**B**) In Study 1, whenever a box is opened, the number of gems in the box is revealed in a numerical format. (**C**) In Study 2, the gems in a box are displayed in a graphical format using pop-out window with a 37×27 grid. Participants use left click to open boxes and right click to indicate their final choice.



(B) Model-based estimation of threshold

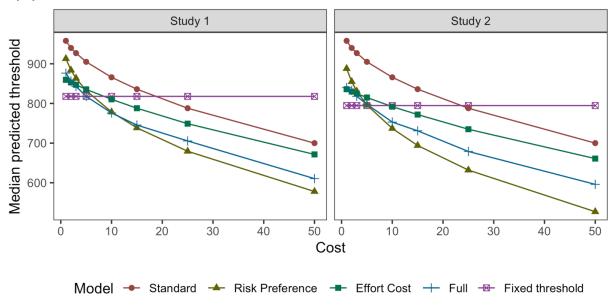
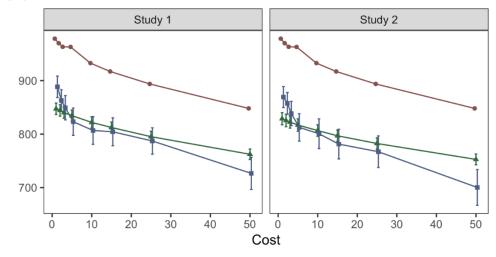


Figure 2. The relationship between optimal thresholds and various parameter values. (**A**) Optimal thresholds as a function of risk preference (α) and search cost parameters; (**B**) Predicted optimal thresholds for different search cost levels based on the various threshold models using Gaussian error. Model predictions are based on median participant-level best-fit parameters. Note the standard model in **B** corresponds to the white horizontal line (where $\alpha = 1$) in **A**.

(A) Chosen rewards



(B) Number of searches

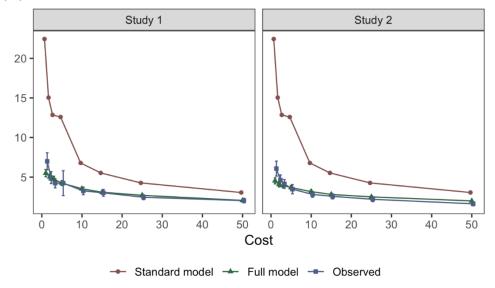


Figure 3. Observed and predicted data for each cost level. Predictions are obtained from the standard model and the full Gaussian model (averaging across participant-level predictions based on their best-fit parameters). (**A**) The reward values at different cost levels; (**B**) The number of searches (i.e., boxes opened) at different cost levels. Error bars represent standard errors.

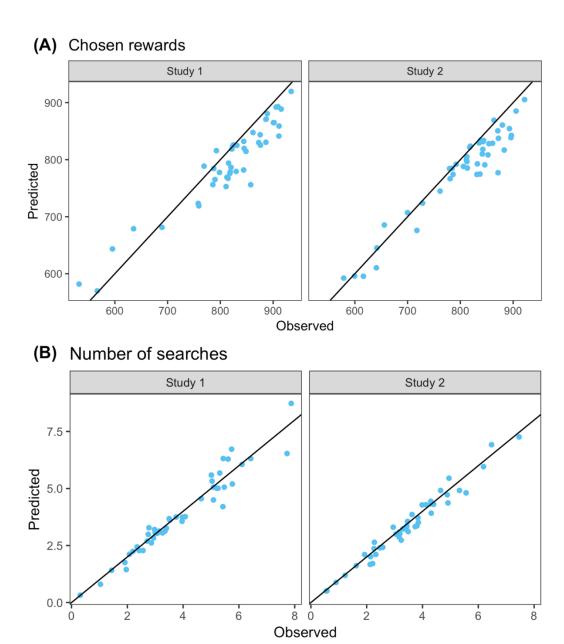


Figure 4. Predicted and observed data at the individual level. Predictions are obtained using the full Gaussian model (using participant-level best-fit parameters). (A) The average reward value each participant chose; (B) The average number of searches each participant made. Each point represents a participant's aggregate data across different cost levels.

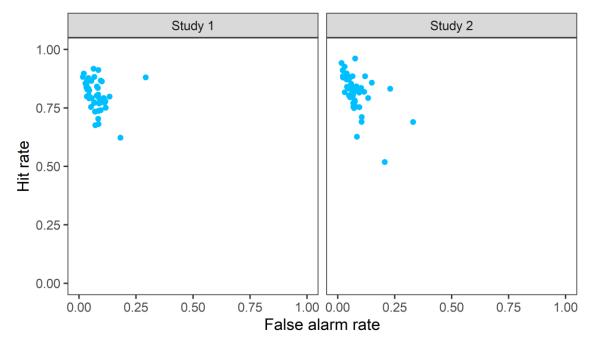


Figure 5. Scatterplots of the full Gaussian model's average predicted probabilities of search termination after each search in each trial. Each participant's data were categorized into two sets based on whether or not the participant chose to terminate search at the given point in the trial. We then calculated the full model's predictions for that point in the trial, both when search was terminated (i.e., the hit rate on the *y*-axis) and when it was not terminated (i.e., the false alarm rate on the *x*-axis). Each point in the scatterplot represents the average probability for a given participant.