Cognitive Science 30 (2006) 1–26

A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem

Michael D. Lee

Department of Cognitive Sciences, University of California, Irvine

Received 11 May 2005; received in revised form 9 January 2006; accepted 13 January 2006

Abstract

We consider human performance on an optimal stopping problem where people are presented with a list of numbers independently chosen from a uniform distribution. People are told how many numbers are in the list, and how they were chosen. People are then shown the numbers one at a time, and are instructed to choose the maximum, subject to the constraint that they must choose a number at the time it is presented, and any choice below the maximum is incorrect. We present empirical evidence that suggests people use threshold-based models to make decisions, choosing the first currently maximal number that exceeds a fixed threshold for that position in the list. We then develop a hierarchical generative account of this model family, and use Bayesian methods to learn about the parameters of the generative process, making inferences about the threshold decision models people use. We discuss the interesting aspects of human performance on the task, including the lack of learning, and the presence of large individual differences, and consider the possibility of extending the modeling framework to account for individual differences. We also use the modeling results to discuss the merits of hierarchical, generative and Bayesian models of cognitive processes more generally.

Keywords: Decision-making; Problem-solving; Optimal stopping; Hierarchical Bayesian modeling; Generative modeling

1. Introduction

1.1. Optimal Stopping Problems

Many real world decision-making problems are sequential in nature. A series of choices is made available over time, and it is often efficient (and sometimes even necessary) to make a selection without waiting to be presented with all of the alternatives. On long cross-country drives, for example, people refill their cars at one of a sequence of towns on the route, without knowing the price of fuel at subsequent towns. This type of sequential decision has a continu-

Correspondence concerning this article should be addressed to Michael D. Lee, Department of Cognitive Sciences, University of California, Irvine, CA 92697–5100. E-mail: mdlee@uci.edu

ous utility function. People aim to choose the cheapest price, and measure their success by how much their purchase exceeded this minimum.

Other sequential decision-making tasks have binary utility functions, where any incorrect decision is equally (and completely) incorrect. For example, consider being a witness for a police line-up, where, because of the circumstances of the case, the offender is known to be in the line-up. Police line-up policy demands that suspects are presented one at a time, may only be viewed once, and that a suspect must be identified at the time they are presented (e.g., Steblay, Deisert, Fulero, & Lindsay, 2001). The aim is to choose the offender, and any misidentification has the equally bad outcome of selecting an innocent suspect.

These real world decision-making scenarios have the same essential features as some optimal stopping problems studied in recreational mathematics (see Ferguson, 1989, for an overview). In this paper, we consider human performance on an optimal stopping problem where people are presented with a list of numbers independently chosen from a bounded uniform distribution. People are told how many numbers are in the list, and how they were chosen. People are then shown the numbers one at a time, and are instructed to choose the maximum, subject to the constraint that they must choose a number at the time it is presented, and any choice below the maximum is incorrect.

1.2. Why Study Optimal Stopping?

Problems like this occupy a useful niche in the study of human problem solving, for at least three reasons. First, the optimal stopping problem is suited to controlled laboratory study, unlike studies of expertise in 'knowledge-rich' real-world domains (e.g., Klein, 1998).

Secondly, the optimal stopping problem has features of real-world problem solving not evident in 'knowledge-lean' problems like the "Towers of Hanoi" or "Cannibals and Missionaries". Historically, most laboratory research on human problem solving has relied on these sorts of artificial problems, characterized by well-defined initial and terminating states that must be linked by a systematic, finite series of steps. Typically, these problems are deterministic, and have state spaces with combinatorially limited possibilities. Optimal stopping problems, in contrast, do not have a simple deterministic solutions and require people to reason under uncertainty. In this way, optimal stopping problems allow for the study of not only what Simon (1976) terms 'substantive' rationality— the ability of people to produce optimal final decisions—but also what he terms 'procedural' rationality—the efficiency of the processes required to make the decision.

Thirdly, the optimal stopping problem neatly complements combinatorial optimization problems, like the Traveling Salesperson Problem (TSP), for which human performance has recently begun to be studied (e.g., MacGregor & Ormerod, 1996; Vickers, Butavicius, Lee, & Medvedev, 2001; Vickers, Bovet, Lee, & Hughes, 2003). Because they are not inherently perceptual, optimal stopping problems allow consideration of whether results obtained with problems like TSPs generalize to cognitively-based problem solving. Optimal stopping problems also introduce uncertainty, and place demands on memory. While visual problems like TSPs are combinatorially large, the basic information required to solve the problem is always perceptually available in a complete and certain form to subjects. In contrast, the sequences of information in optimal stopping problems are stochastic and presented only temporarily, requiring people to deal with uncertainty and rely on their memory. Our optimal stopping problem also has the advantage of having a known optimal solution, which distinguishes it from problems like TSPs, which are NP-complete, and so have no known process for arriving at optimal solutions. This means we are able to distinguish measures of performance based on achieving optimal outcomes from those based on following optimal decision processes.

1.3. Previous Research

Previous research examing human performance on optimal stopping problems has tended to focus on versions of the problem that provide rank order information, rather than the values themselves (e.g., Dudey & Todd, 2001; Seale & Rapoport, 1997, 2000). These rank order problems are interesting for the same reasons as the problem we study, but have a very different optimal solution process. When only rank order information is available, the optimal process involves observing a fixed number of values, then choosing the first subsequently maximal one (see Gilbert & Mosteller, 1966, Table 2). When the value itself is available, the optimal process is based on a set of thresholds, one for each position in the sequence, and requires a presented value to be selected if it is currently maximal, and exceeds the threshold for its position (see Gilbert & Mosteller, 1966, Tables 7 and 8). These thresholds can naturally be conceived as the aspiration levels in Simon's seminal theory of satisficing (e.g., Simon, 1955, 1982).

Kahan, Rapoport, and Jones (1967) studied human performance on essentially the same task that we consider, using problems of length 200. Different problems involved values drawn from either a positively skewed, negatively skewed, or a uniform distribution. These authors found no evidence for the different distributions affecting the decisions made. They also compared individual and group decision-making, and found that decisions were made earlier in the sequence by individuals. Corbin, Olson, and Abbondanza (1975) considered human peformance on problems of length five, and by systematically manipulating the values presented, found sequential and contextual dependencies within problems. Other empirical studies (e.g., Kogut, 1990; Zwick, Rapoport, Lo, & Muthukrishnan, 2003) have made a large methodological departure by requiring subjects to sacrifice explicitly held resources to view additional presentations, usually because they are interested in applications to economic decision-making.

Lee, O'Connor, and Welsh (2004) presented the study that is most directly relevant to the one reported here. These authors considered human performance on problems with lengths 10, 20 and 40, and evaluated three candidate models, drawn from the previous mathematical and psychological literature, of the way people made decisions. They concluded that the best accounts were provided by 'threshold' models in which people choose by comparing the presented value to fixed thresholds. What Lee et al. (2004) observed, however, was that there seemed to be significant individual differences in the exact thresholds that people used. Some subjects behaved consistently with applying a single fixed threshold across the entire sequence. Effectively, these people chose the first number that exceeded a fixed value. Other subjects, however, behaved consistently with using thresholds that decreased as the sequence progressed, as with the optimal solution. Lee et al. (2004) concluded by arguing that shorter problem lengths needed to studied to provide the empirical data that would distinguish different threshold-based accounts.

1.4. Modeling the Decision-Making of Optimal Stopping

Much of the previous research on optimal stopping problems has proposed formal models of the decision-making process (e.g., Corbin, 1980; Dudey & Todd, 2001; Lee et al., 2004; Seale & Rapoport, 1997), although sometimes their evaluation has taken the form of simulation studies, rather than by making inferences from human data.

Even when comparison with human decisions has been made, however, there have been a number weaknesses in the modeling. Some of these weaknesses are shared with much of contempory psychological modeling in general, such as focusing on the best-fitting behavior of a flexible model as an measure of its adequacy, and ignoring the inherent complexity of the model (see Roberts & Pashler, 2000; Pitt, Myung, & Zhang, 2002). Others are more specific to the nature of the optimal stopping problem, such as failure to make probabilistic inferences about models and their parameters because the choice data are noise-free and the models are deterministic. Even Lee et al. (2004), who did consider these basic model theoretic issues, failed to present a modeling framework that addressed bigger questions, such as how the various decision processes people use to solve the problems might be generated in the first place, how they might adapt these processes, and how individual differences might be accommodated.

In this paper, we try to address some of these modeling issues, studying human performance on problems with length five. We first provide basic analyses that confirm the plausibility of threshold models. We then develop a new generative account, cast in a Hierarchical Bayesian framework, describing how threshold models may be generated and applied to the optimal stopping problem. Applying Bayesian and Minimum Description Length statistical methods—including model averaging, model selection and 'entropification' methods—we learn about the parameters of the generative process, make inferences about the threshold decision models people use, and demonstrate the predictive capability of our approach.

2. Experiments

We consider data collated from two separate experiments.[1] Both experiments were identical methodologically, but involved different subjects, and different problems.

2.1. Method

2.1.1. Subjects

In the first experiment there were 50 subjects (26 males and 24 females), with a mean age of 19.8 years, most of whom received course credit. In the second experiment, there were 97 subjects (44 males and 43 females), with a mean age of 25.3 years, who were paid for participating.

2.1.2. Procedure

Each experiment involved a different set of 40 randomly generated problems, with numbers between 0 and 100, defined to two decimal places. Each subject completed all 40 problems in their experiment in a random order. For each problem, subjects were told the length of sequence was five, and were instructed to choose the maximum value. It was emphasized that (a) the values were uniformly and randomly distributed between 0.00 and 100.00, (b) a value could only be chosen at the time it was presented, (c) the goal was to select the maximum value, with any selection below the maximum being completely incorrect, and (d) if no choice had been made when the last value was presented, they would be forced to choose this value. As each value was presented, its position in the sequence was shown, together with 'yes' and 'no' response buttons. When a value was chosen, subjects rated their confidence in the decision on a nine point scale ranging from "completely incorrect" to "completely correct".

2.2. Results

All of the decision data are summarized in Figure 1. The top half of this Figure contains 40 panels, corresponding to the 40 problems in the first experiment. The second half of the Figure contains 40 panels corresponding to the problems from the second experiment. Within each panel, the five values for that problem are shown in sequence from left to right by circles. Values that were chosen by at least one subject are shown as filled circles, and the area of filled circles is proportional to the relative frequency with which that choice was made. The panels are ordered according to where the maximum value is in the sequence so that, for example, the first row contains problems where the maximum value is the first one presented.

Our basic analysis of these decision data takes the form of identifying five empirical regularities, each of which informs the development of models to account for the way people solved the problems. We present each regularity in turn, and explain the way it informs subsequent modeling.

2.2.1. Individual Differences

Figure 2 shows the variation in performance across all subjects. The left-hand panel shows the relationship between decision accuracy, as measured by the proportion of times each subject chose the maximum value, and mean confidence across all decisions. The right-hand graph shows how decision accuracy, as measured by the proportion of times each subject chose in accordance with the prediction of the optimal decision rule, relates to mean confidence. It is clear that there is considerable individual variation in both measures of decision performance, and in confidence, and a positive relationship between the variables. In particular, we note that while some subjects choose consistently with the optimal decision rule only about half of the time, others follow it almost perfectly. The message we draw from this empirical regularity is that it is important to have a model capable of accommodating individual differences.

2.2.2. Choosing the Current Maximum

A second regularity is evident by counting how often, across all subjects and problems, the value chosen was the maximum value encountered at that stage. For the first through fifth position, respectively, this was true in 100% (inevitably), 98.5%, 96.4%, 97.7%, and 35.6% of cases. Thus, it is clear that people almost always choose a currently maximal value, until the final position where forced choices are made. The message we draw from this empirical regularity is that a model of human performance needs to keep track of the maximum value encountered at any stage in a problem sequence, and only choose currently maximal values.

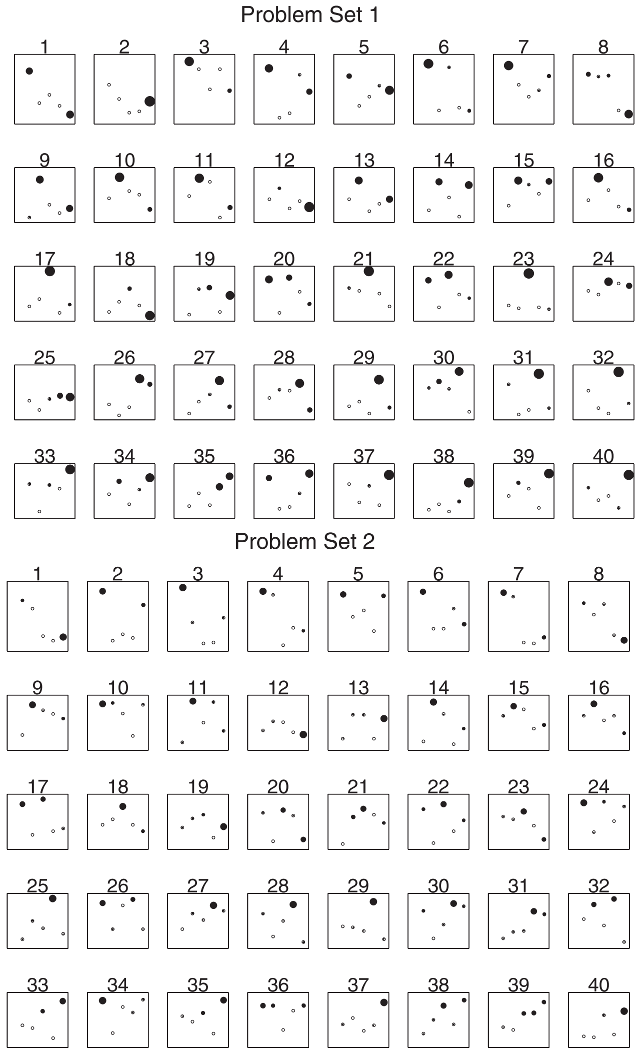

Fig. 1. Summary of the problems and decisions for both problem sets. Each panel, corresponds to a problem, with the five values for that problem shown in sequence from left to right by circles. Values that were chosen by at least one subject are shown as filled circles, and the area of filled circles is proportional to the relative frequency with which that value was chosen.

2.2.3. Using Different Thresholds

Figure 3 shows how often various values were chosen, across all problems and subjects, in terms of their position in the problem sequence. Visually, there seems to be a trend for successively lower values to be more likely to be chosen at later positions. This trend is confirmed in the analysis presented in Figure 4, which shows the proportion of values in successive ranges of length ten (0 to 10, 10 to 20, ..., 90 to 100) that were chosen when presented, for each of the five positions in the sequence. Starting in the 40s, for example, a greater proportion of values

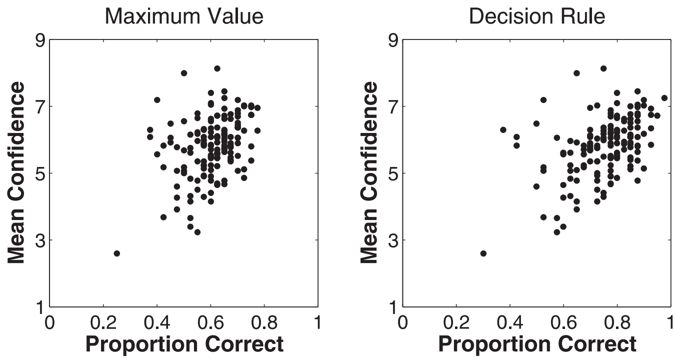

Fig. 2. The relationship between decision accuracy and mean confidence, with decision accuracy assessed in terms of choosing the maximum value (left-hand panel) and choosing consistent with the optimal decision rule (right-hand panel). Each point corresponds to an individual subject.

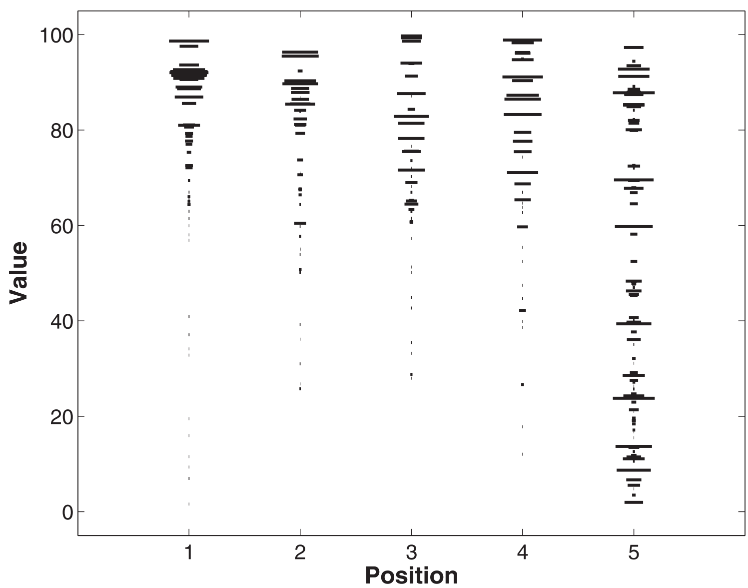

Fig. 3. The relative frequency with which values, across all problems and subjects, were chosen, as a function of their position in the problem sequence. Each line corresponds to a value at a position in a problem, with the length of the line indicating how often that value was chosen.

are chosen when they are presented in position four than in earlier positions. A similar effect is evident starting in the 60s for position three. In general the ordered and monotic rise of the proportions as values increase suggests that people lower their threshold value for choice as the problem sequence progresses. The message we draw from this empirical regularity is that any threshold model needs to have the flexibility to apply successively lower thresholds at later positions in the problem sequence.

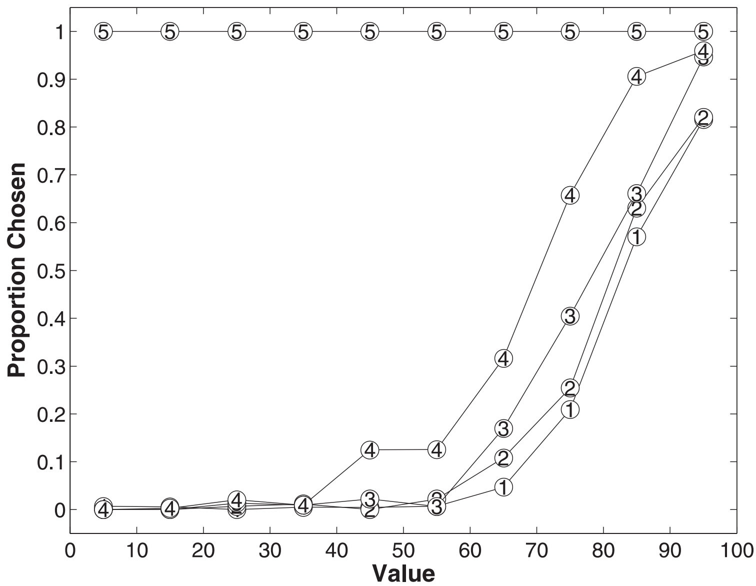

Fig. 4. The proportion of presented values in ranges 0 to 10, 10 to 20, and so on, chosen across all problems and participants, as a function of the position in the sequence for the presented value.

2.2.4. *Insensitivity to Previous Values*

Figure 5 shows the relationship between values that occur in the position before a choice is made, and those that precede values that are also not chosen. There is no obvious difference between the distributions for values at the same position in the problem sequence. This conclusion is supported by the analysis presented in Figure 6, which shows as scatteplots the relationship between a value that was not chosen, and the proportion of times the subsequent value was chosen. There is no strong relationship between these variables for any of the first three positions in a problem sequence, with correlation measures corresponding to less than 10% of the variance being explained in each case. This means that knowing the preceding value in a problem provides little information about whether the subsequent one will be chosen. Although not presented in detail here, other analyses considering earlier values than the immediately preceding one in the sequence, and also the averages and maxima and minima of a sequence of earlier values, also failed to find any strong relationship. The message we draw from these findings is that, as a first approximation, beyond keeping track of maximality, we do not need models that are sensitive to previous values in a problem sequence to explain human performance.

2.2.5. *Lack of Learning*

Figure 7 shows the pattern of change in decision accuracy across all subjects and problem sets, in the order in which they were completed. It can be seen there is no evidence for either the maximum value (left-hand panel) or decision rule (right-hand panel) measure of accuracy that performance improves as additional problems are completed. The message we draw from this empirical regularity is that, as a first approximation—at least for this version of the task, which does not include any feedback—we do not need learning models to explain human performance.

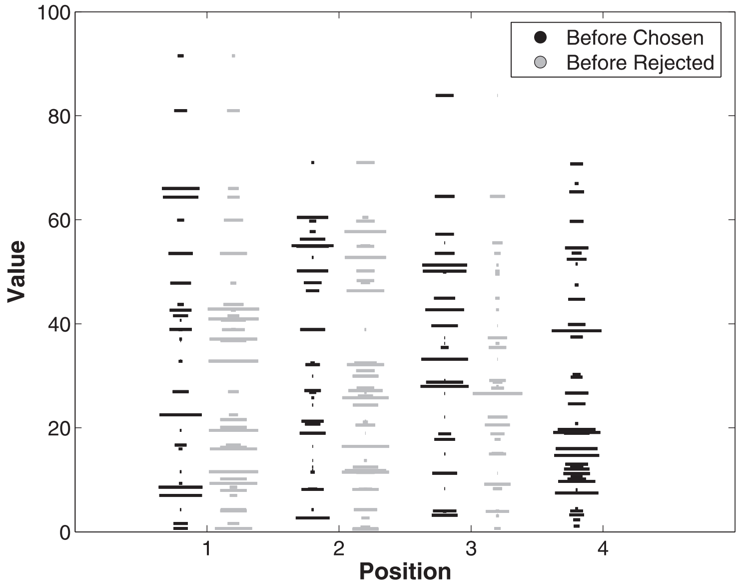

Fig. 5. Distributions of values in positions immediately before chosen and rejected values across all problems and subjects. Each line corresponds to a value at a position in a problem, with black lines representing values before a choice, and grey lines representing values before a rejected value. The length of lines indicates how often that value was before a chosen or rejected value.

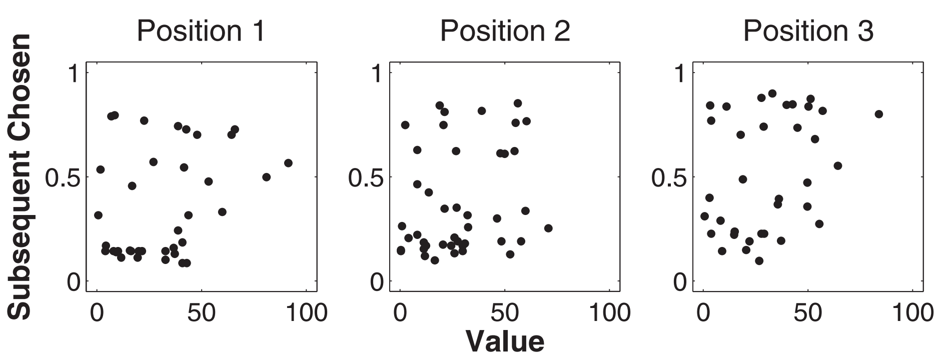

Fig. 6. Scatterplots showing the relationship between a presented value that was not chosen, and the proportion of times the subsequent value was chosen, for the first three positions in a problem sequence.

2.2.6. *Consolidated Conclusion*

Considering these five empirical regularities collectively suggests that human performance on the task might be modeled using a flexible threshold approach, where the first currently maximal value above a position-dependent threshold is chosen. These thresholds need to be specified, in the first instance, at the level of individual subjects, but do not need to depend on previous values in a problem sequence, nor be subject to learning across successive problems.

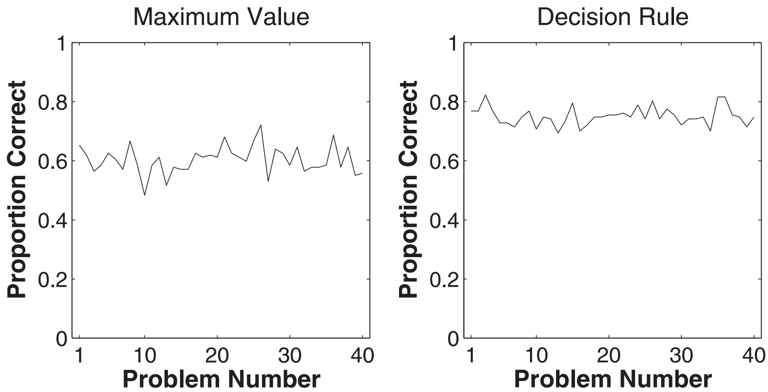

Fig. 7. The pattern of change in decision accuracy across problems in the order they were completed, as measured by choosing the maximum value (left-hand panel) and choosing consistent with the optimal decision rule (right-hand panel).

3. A Generative Model

In this section, we develop a generative account of the threshold-based model family suggested by the empirical regularities. First we specify how threshold models may be generated, and then how these models make decisions for a specific set of problems. Having specified the generative process, we then describe the reverse statistical process by which the parameters of the model can be inferred from behavioral data.

3.1. Model Family Level

Any threshold model implicitly uses a threshold of zero for the final position, since it is a forced choice. Accordingly, our account of generation of threshold models for problems of length five describes how the first four thresholds are determined. We begin by assuming the threshold for the first position is in place, and then describe the processes that give rise to the second, third and fourth thresholds. Moving from the first to the second threshold presents two reasonable possibilities: fixing at the first threshold, or moving the threshold downwards by some amount. Moving to the third and fourth thresholds then involves three possibilities: fixing at the previous threshold, moving downward by the previously used amount of decrease, or moving downward by a new amount.

Overall, we (crudely) represent this generative process as a multinomial $\theta = (\theta_F, \theta_D, \theta_N)$, where θ_F gives the probability of fixing the threshold, θ_D gives the probability of moving down by the same amount, and θ_N gives the probability of moving down by a new amount. This parameterization is shown at the top of Figure 8, with the multinomial θ represented in a triangle that has the F, D, and N events as vertices.

For problems of length five, the model family encompasses the 14 different models shown in the middle of Figure 8. Each of the models is labeled according to the sequence of fixed (F), down (D), and new (N) transitions by which it was generated. Continuing horizontal lines indi-

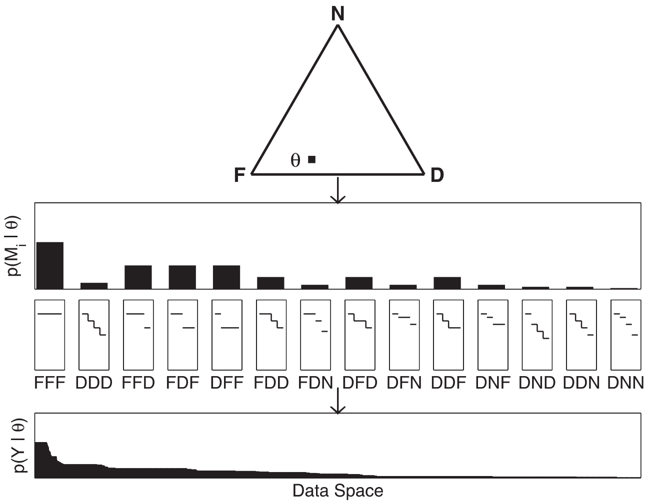

Fig. 8. A generative account of using threshold models to make decisions on the optimal stopping problem.

cate a fix transition, while vertical lines indicate down transitions. The first FFF model simply uses a fixed threshold across the problem sequence. The next DDD model decreases the threshold by the same amount as positions progress through the sequence, and so on. The final DNN model is the most flexible one possible, since it allows any threshold in any position, subject to the constraint that thresholds do not increase.

A particular parameterization θ of the generative process corresponds to a probability distribution over the 14 models, determined in the obvious way. The probability of the FFF model is θ_F^3, the probability of the DDD model is θ_D^3, the probability of the FFD model is $\theta_F^2\theta_D$, and so on. For the particular choice of θ shown at the top of Figure 8, the probability distribution $p\,(M \mid \theta)$ over the 14 models is shown immediately below.

3.2. Model Level

Each of the 14 models involves one or more parameters. For the FFF model, there is a single parameter, that determines the fixed threshold. For the DDD model, there are two parameters, determining the initial threshold, and the steady rate of decrease in subsequent thresholds, and so on. For the final DNN model, there are four parameters, setting the thresholds at each of the four positions in the problem sequence. At a specific parameterization, a model is applied to a problem by having it choose the first currently maximal value above its threshold for that position in the sequence.

The models make different decisions across a set of problems as their parameters are varied. Since we are dealing with sets containing 40 problems, the data space contains the set of all 40-vectors containing the values 1, 2, . . . , 5 in each position, with the values giving the choice made for each problem. Effectively, therefore, over all parameterizations, models index a distribution across the data space, indicating which sets of choices they are able to make. We denote these sets of decisions indexed by the models as Y. We note that different threshold models have different complexities, in the sense that they index different numbers of sets of decisions. For example, for our first problem set, the simple fixed-threshold FFF model indexes only $||Y|| = 79$ different sets of decisions, while the fully flexible DNN model indexes $||Y|| = 3, 121$. Other models have intermediate complexity. A similar state of affairs is true for the second problem set in this study, with the FFF model indexing $||Y|| = 81$ sets of decisions, and the DNN model indexing $||Y|| = 2, 959$.

3.3. Data Level

Using the distribution over the models arising from the model family level, the collection of decisions sets indexed by individual models are combined in proportion to their probabilities. This mixture $p\,(Y \mid \theta)$ is the final result of the generative model: a single parameterization at the mode family level produces a distribution in the data space, as mediated by the family of 14 possible threshold models. The bottom of Figure 8 shows this distribution, ranking the indexed decision sets from most to least probable, given the initial parameterization θ of the generative process.

4. Model Inference

We place a uniform prior over θ at the generative level, so that $p\,(\theta) \propto 1$. Given decision data D, the posterior for the generative process parameter θ,

$$p\,(\theta \mid D) \propto p\,(D \mid \theta)\, p\,(\theta)\,,$$

is found by assuming a uniform prior $p\,(\theta) \propto 1$, and integrating across the 14 models,

$$p(D|\theta) = \sum_i p(D|M_i)\, p(M_i|\theta),$$

where M_i denotes the ith model, and $p\,(M_i \mid \theta)$ is found in the obvious way described earlier, with $p\,(M_{FFF} \mid \theta) = \theta_F^3$, and so on.

The final quantities required to make inferences from data, $p\,(D \mid M_i)$, present an interesting challenge. The threshold models are deterministic, in the sense that a particular model at a particular parameterization makes exactly one set of decisions for a problem set, and so indexes a single point in the data space with probability one. If the set of decisions observed empirically do not coincide with this point, there is effectively zero probability at that parameterization. If the model does not index the empirical decisions at any parameterization, then there will be no probability as a whole. This means that standard Bayesian methods, which rely on integrating the likelihoods of models across the parameter space, are not applicable for the inferences we require.

One way to address this problem would be to introduce an error theory to the model. An alternative approach, that is in some ways more principled and satisfying, was developed by Lee (2004). This approach uses an information theoretical method, based on Minimum Description Length (MDL) methods for model selection, called 'entropification' (Grünwald, 1998, 1999;

Myung, Pitt, & Kim, 2005). Entropification provides a principled technique for associating deterministic models with probability distributions, allowing inferences to be made that are 'safe', in the sense of minimizing the expected worse case errors. Intuitively, entropification introduces a principled conservative error theory to the process of inference, so that empirical data is able to be brought into contact with deterministic models.

A detailed and general account of applying entropification under 0–1 loss is provided in the Appendix. Basically, we consider the conditional probability

$$p\left(D \mid Y, w, M\right) = \frac{e^{-wm}}{\sum_{x=0}^{40} \binom{40}{x} \left(k-1\right)^{x} e^{-wx}},$$

where $k = 5$ is the length of the problem, w is a positive scalar, and $f_{0-1}(D, Y) = m$ is the 0–1 loss function that indicates the model M makes m different decisions in its indexed decision set Y from the empirical data D of a subject. The entropification method then requires finding

$$p^{*} = \max_{(Y,w)} p\left(D \mid Y, w, M\right),$$

giving the MDL value

$$MDL = -\log p^{*} + \log \left\| Y \right\|,$$

where $\|Y\|$ is the total number of predictions indexed by the model. This MDL value can then be used to give the required probability as

$$p\left(D \mid M\right) \approx e^{-MDL}.$$

In addition, the entropification process identifies a 'best' prediction, Y^*, which is naturally associated with a parameterization of the model, in the form of a set of thresholds.[2] This set of thresholds can be regarded as the result of 'fitting' each model to the data.

4.1. A Concrete Example

We now give a concrete example of inference across the generative model. Following the observation of potentially large and meaningful individual differences, inference is always done here at the level of individual data. In particular, we focus on a single subject from the second experiment, whose analysis contains clear examples of most of the interesting and important features of the inference process.

Inference for this subject is summarized graphically in Figure 9. Their data (i.e., the decisions they made across all 40 problems) are shown in the bottom panel. Each of the 5 × 8 boxes represents a problem, with circles from left to right corresponding to the values presented in the sequence, and the subjects choice indicated by the filled circle. Thus, for this subject, their set of decisions is represented by the 40-vector $D = [5, 1, 1, 1, 5, \ldots, 5]$.

Applying the entropification method for the FFF model against these data involves comparing every set of decisions the model can make by varying its single threshold parameter. The

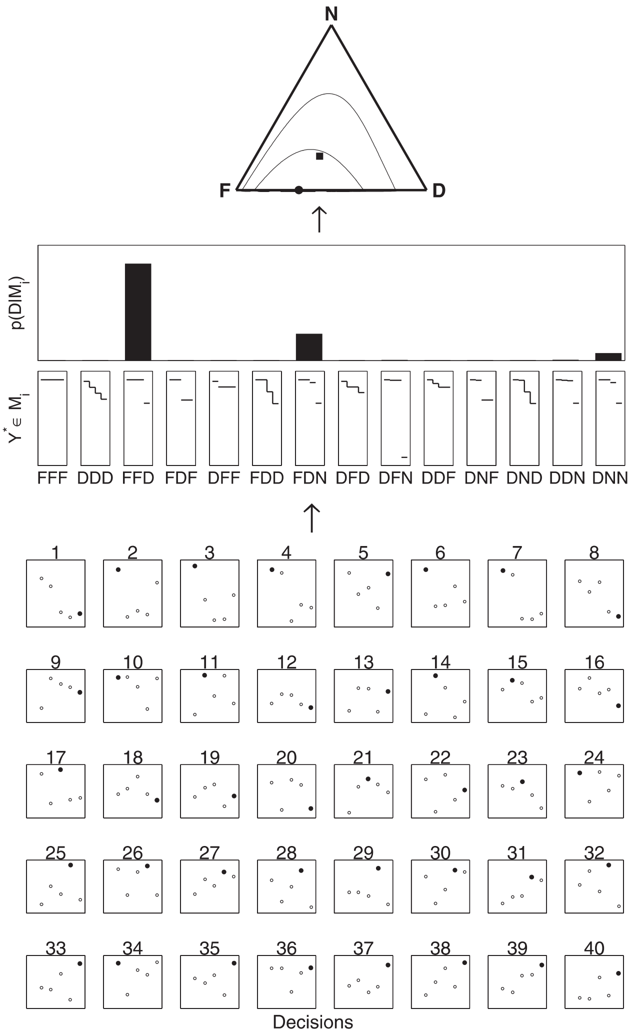

Fig. 9. The results of making inference across the generative model for one subject in the second experiment. The decisions made by the subject are shown in the bottom panels, from which best indexed sets of decisions for each model are inferred, and the marginal densities of the models evaluated, as represented in the middle panels. These densities, in turn, are used to infer the posterior over the model generating parameters, summarized by the mode (circle), mean (square) and 50% and 95% credible intervals of the distribution in the top panel.

threshold value corresponding to the best indexed set of decisions $Y^* \in M_{\mathrm{FFF}}$ is shown in the left-most panel immediately above the data. This is the best single threshold model for the subject's data. Similarly, the best parameterizations of the remaining models are shown in the other panels immediately above the data.

Each of these best models has an MDL value, which is naturally associated with the probability the subject's data would have arisen under the model $p\,(D \mid M)$. These marginal densities are shown by the bar graph above the panels for each model. It can be seen that only three of the 14 models—the FFD, FDN, and DNN models—have any appreciable density. This means these models, at their best setting of thresholds, are the best able to match the decisions made by the subject across the 40 problems.

Finally, on the basis of the different marginal densities found for each of the models, an inference is made about the generative process the subject used to construct candidate models. This corresponds to the rate at which they use fixed, downwards, and new transitions moving from one threshold to another as the problem progresses. The range of values for these rates that can be inferred from the subject's data are summarized at the top of Figure 9, with the mode indicated by a circle, the mean by a square, and 50% and 95% credible intervals drawn. It can be seen that, consistent with the strong evidence in favor of the FFD model, rates are most likely that give greatest probability to fixed transitions, then downwards transitions, then new transitions. The wide credible intervals show, however, that considerable uncertainty about the rates remains. At this level of abstraction, the data do not provide a strong constraint on what can be inferred about the subject's cognitive processes.

4.2. Posterior Prediction

Having made inferences about the generative process from observing a subject's data, it is possible to specify a posterior predictive model. This is essentially a set of thresholds that represent the account our modeling framework provides of how that subject made their decisions. It defines a threshold model for that subject that can be applied to other problems, furnishing a prediction about how that subject will behave in the future.

We consider three approaches to constructing a posterior predictive model. The first approach simply uses the preferred threshold parameterizations in the maximally flexible DNN model. The DNN model allows for any combination of thresholds, and the entropification process identifies from the subject's data which of these combinations is the most likely. In effect, this approach corresponds to choosing the 'maximum likelihood' model.

The second approach uses the preferred threshold parameterization for the model with the best MDL value across all 14 models. These MDL values are sensitive to the various complexities of the models, as measured by the number of data points they index. Accordingly, the set of thresholds selected will be sensitive to both data-fit and complexity, with a focus on choosing the simplest sufficiently accurate model. In this sense, the approach corresponds to choosing the 'maximum a posteriori' model.

The third approach uses the preferred parameterization of all of the 14 models, by averaging across them in proportion to the posterior probability of each model, as given by

$$p\left(M_i \middle| D\right) = \frac{p\left(D \middle| M_i\right) p\left(M_i\right)}{\sum_j p\left(D \middle| M_j\right) p\left(M_j\right)}.$$

That is, the threshold for the first position is defined as the average across the best threshold for each of the 14 models, weighted by the posterior of the model given the observed. This approach corresponds to 'model averaging'.

For the particular subject whose inference was demonstrated in Figure 9, the three predictive models found by these approaches are shown in Figure 10. The maximum likelihood ac-

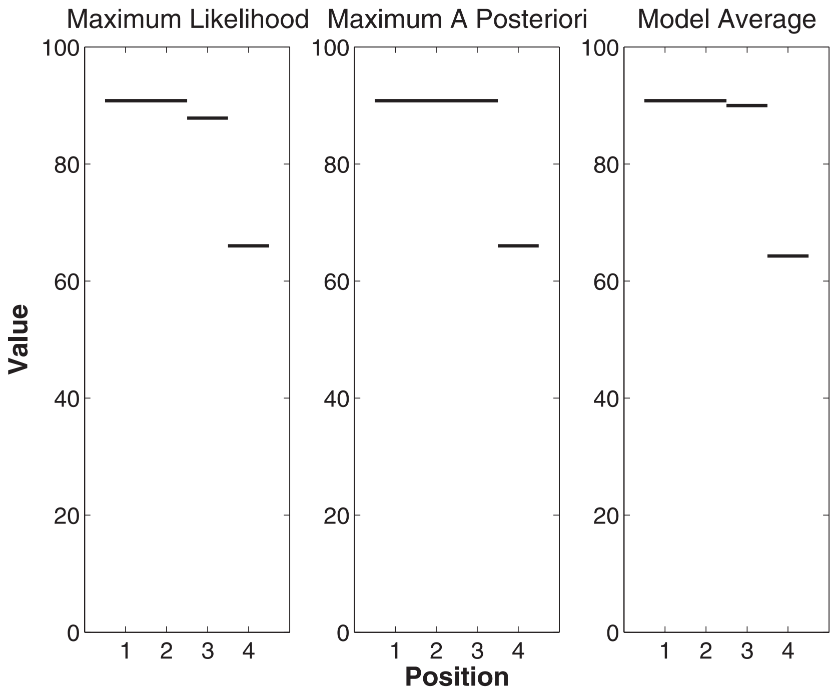

Fig. 10. The predictive models inferred from the decisions made by a subject in the second experiment, using the maximum likelihood (left panel), maximum a posteriori (middle panel) and model averaging (right panel) methods.

count corresponds to the best thresholds inferred within the fully flexible DNN model. The maximum a posteriori account corresponds to the best thresholds inferred within the simpler FFD model, which has the greatest marginal density, and so provides the best balance between data-fit and model complexity. The model average account essentially combines the best thresholds inferred within the FFD, FDN and DNN models, since these are the only three with significant marginal density.

Both the first two approaches are frequently used in cognitive modeling. The maximum likelihood approach essentially corresponds to selecting the best-fitting model, while the maximum a posteriori method essentially corresponds to Bayesian model selection, preferring the model that best fits on average across all possible parameterizations.

Despite strong evidence it leads to better predictions (see Hoeting, Madigan, Raftery, & Volinsky, 1999, for a review), the model averaging approach is rarely applied in cognitive modeling, because (at least in part) it often leads to problems with interpretability. A blended model that is nine parts Model A to one part Model B, where each competing model provides a very different interpretation of a cognitive process, uses different parameters, and so on, is unlikely to be easy to understand and motivate theoretically. The threshold family of models we are considering is exceptional in this regard, because it is closed under model averaging. That is, the combination of any set of particular threshold model is itself a threshold model, and so is readily amenable to exactly the same interpretation as an account of the cognitive decision-making process involved.

5. Modeling Results

In this section, we provide an account of making inferences for all 147 subjects using the generative model. At the model family and model levels, our results are largely descriptive and exploratory, summarizing the range of parameterizations of the basic cognitive decision-making process that was observed. At the level of data, we provide a much stronger evaluation, by testing the ability of the inferred predictive models to account for human decision-making.

5.1. Model Family Level

Figure 11 shows the means of the inferred posterior for the generative parameters at the model family level. Each of the 147 subjects across the two experiment is shown by a point. The inference that would have been made if a subject had adhered perfectly to the optimal decision rule for the two problems sets are shown by large circles.

There is some significant level of variation in these means, particularly with respect to the emphasis given to the fix and down transformations. There is also some evidence of clustering, suggesting the possibility of different 'strategies' in solving the problem, although this speculation needs to be tempered by level of uncertainty indicated by the wide credible intervals in Figure 9, which are typical of those found for all of the subjects. It is also interesting to note that many subjects have means extremely similar to those corresponding to optimal performance, while others deviate either by giving relatively greater emphasis to the fix or the down transition.

5.2. Model Level

Figure 12 shows the predictive models inferred using the model averaging approach for each of the 50 subjects from the first experiment. The predictive models are shown as bold

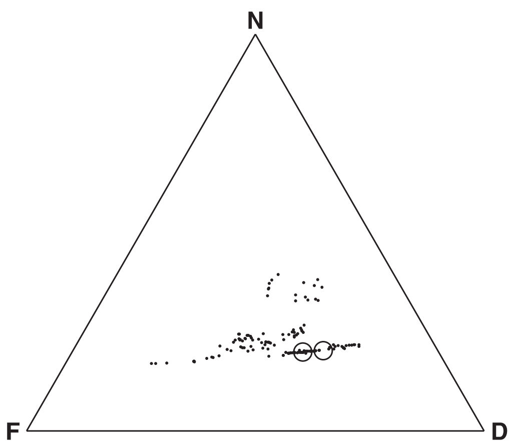

Fig. 11. Means of the posterior of the generative process $p(\theta \mid D)$, for each subject (shown as points), and for data sets corresponding to using the optimal decision rule for both problem sets (large circles).

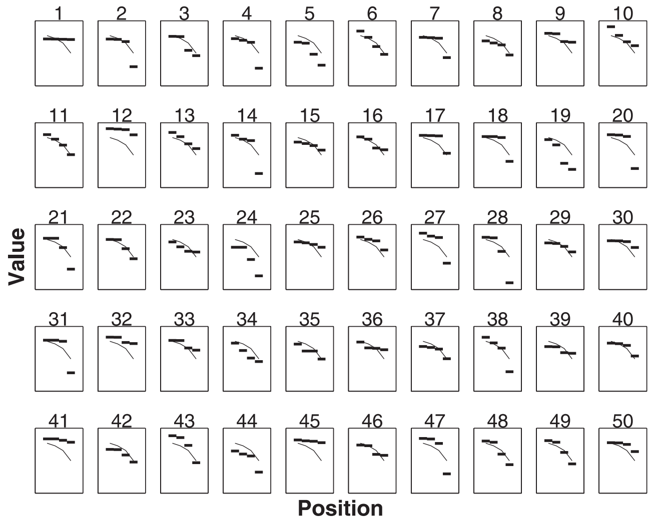

Fig. 12. Predictive models (bold lines) inferred using the model averaging methods for each of the 50 subjects in the first experiment, superimposed on the thresholds for the optimal decision rule.

lines, superimposed over the optimal decision rule thresholds. These predictive models are entirely representative of those inferred for the 97 subjects in the second experiment.

There is clear variation across subjects at this level of analysis. For example, a visual comparison of subjects 1, 11, 18 and 21 reveals four quite different threshold models. Subject 11 is close to optimal, while subject 1 relies on a single fixed threshold, and subjects 18 and 21 use non-optimal decreasing thresholds that accelerate downwards at very different rates.

5.3. *Predictive Accuracy*

A very direct and practical evaluation of the usefulness of our generative model is provided by examining its ability to predict decision-making behavior. We do this using cross-validation, in which a subset of each subjects' decisions are used for inference, and the resultant predictive models are then evaluated against the unseen data. This evaluation simply measures what proportion of the subjects' decisions were correctly chosen by the predictive model.

Specifically, we consider training sets with random samples of 2, 3, 4, 5, 10, 20 and 30 problems. For each of these sizes, 100 different training sets were used. Figure 13 shows the pattern of change in predictive accuracy for the maximum likelihood, maximum a posteriori and model average methods, as the number of training problems increases. Mean performance is shown, together with one standard error in each direction. With the maximum available information in this analysis (30 problems), the methods appear to be approaching about 75% accuracy. This compares well with chance performance, which obviously corresponds to 20% accuracy.

It is interesting to note that the model averaging method is superior to the other methods, particularly when the number of training problems used for inference is small. The maximum a posteriori method also seems to outperform the maximum likelihood method for small num-

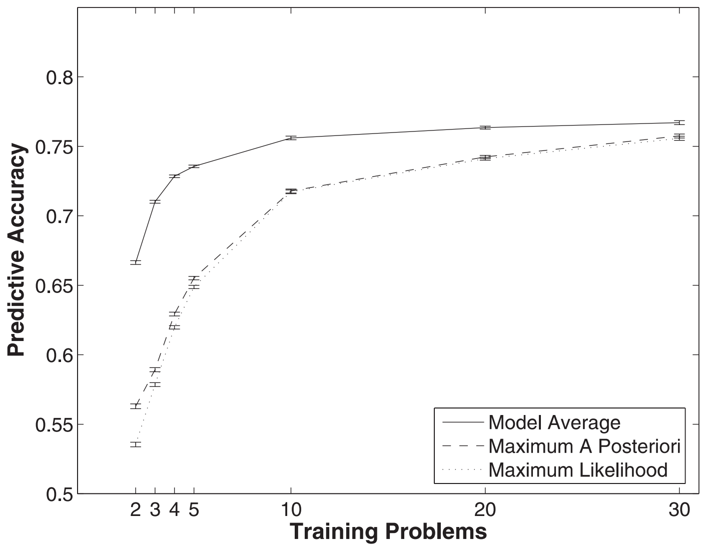

Fig. 13. Cross-validation results showing the change in predictive accuracy as a function of training set size, for the maximum likelihood, maximum a posteriori and model averaging methods. Error bars represent one standard error about the mean in each direction.

bers of training problems. In absolute terms, the model average method performs impressively. On average, it defines a predictive model from observing only three problems that agrees with more than 70% of the 37 other decisions made by subjects.

6. Discussion

Our discussion considers both methodological issues associated with the modeling approach presented here, and psychological issues concerning the observed decision-making behavior of people on the optimal stopping task.

6.1. Hierarchical, Generative and Bayesian Modeling

Most popular models of cognitive processes can be characterized as low-dimensional parametric models. These models describe a formal relationship between a small set of parameters and the data that are observed in a cognitive task. The modeling approach we have adopted is different, and is inspired by some recent generative and hierarchical models of cognition (e.g., Kemp, Perfors, & Tenenbaum, 2004; Griffiths, Baraff, & Tenenbaum, 2004). Accordingly, it is worth discussing what we view as the benefits of our modeling approach.

6.1.1. Hierarchical Modeling

The strength of hierarchical models is that they are able to represent knowledge at different levels of abstraction. This is a key feature of our account of making decisions for the optimal stopping problem. At the lowest level, we are able to think about the observed behavioral data. At the next level of abstraction, we are able to make inferences about specific models (like the fixed threshold model) that have parameterizations (in the form of thresholds, or differences between thresholds) that produce behavioral data when brought into contact with a problem stimulus. At our highest level of abstraction, we can make inferences about parameters that describe which specific models might be used.

The ability to represent information and uncertainty at different levels in these ways, and to make inferences within and between these levels, allows us to give a more complete and coherent account of decision-making for this task than would otherwise be possible. Low-dimensional parametric models usually focus on one level of explanation, at least in a formal sense. This means that 'meta-cognition' becomes a separate line of inquiry, and issues of parametric self-regulation and adaptation (rather than just estimation) are difficult to address. In a more practical way, a consequence of this narrowness in the level of abstraction is that most cognitive models are not amenable to model averaging (even if there are no in-principle problems with interpretability), because there is no over-arching account that unifies competing models. Thus, at least in part, the impressive performance of our predictive model using the model averaging method stems from the hierarchical nature of our account.

6.1.2. Generative Modeling

The strength of generative models is that they provide some account of how relatively specific cognitive processes might be instantiated and bounded. Having defined a plausible generative mechanism for the model family, the specification of individual models becomes a formal deduction, rather than an act of creation to be given plausible justification. Of course, other reasonable assumptions at our most abstract level would have possible, but at least the generative process is explicit. Its outcomes are also non-trivial, in the sense that there are many threshold-based models could have been proposed that are not in the current model family[3].

More practically, we believe that theory and model development is often streamlined by adopting a generative perspective. All of the theoretical and empirical insights a researcher has about a cognitive process can contribute to an account of how different data would be generated, without being encumbered by questions of inference. Once a full generative account is in place, inference involves the (conceptually) simple idea of reversing the process, and finding the most likely account for the data that have actually been observed.

6.1.3. Bayesian Modeling

The strength of Bayesian models is that they adopt a coherent method for statistical inference, founded on probability theory (Cox, 1961; Jaynes, 2003). Given a generative process across a hierarchical model structure, it is simple both to generate data, and making inferences about the parameters of the processes from data. Both questions correspond to logically deduced probability statements, rather than requiring ad hoc specification of heuristic devices. Indeed, the modeling reported here probably represents a worst-case example of the application of Bayesian inference, since the use of deterministic models necessitated a detour into (closely related) MDL methods. Even this worst case, however, contrasts favorably with the prospect of having to make inferences about our model using sampling distribution methods, defining estimators and making hypothesis tests. A little reflection suggests that sampling distribution methods simply were not developed to make inferences across structured generative accounts of cognitive processes. This provides another reason to stop using them, as if that were needed (see Jaynes, 2003, ch. 17).

6.2. Psychological Issues

6.2.1. Lack of Sequential Effects

We found, as summarized in Figure 5, that previous values in a problem did not exert a large influence on the value chosen. This conflicts with the findings of Corbin et al. (1975), who did find evidence that the sequence of values presented affected choice, with the same value more likely to be chosen after a series of relatively low values. Theoretically, the definition of the optimal stopping problem makes it clear that previous values are irrelevant, except for the need to remember the current maximal value. Experimentally, there are several possible reasons for the discrepancy in findings. Most fundamentally, Corbin et al. (1975) studied something closer to the rank order version of the problem, because they did not restrict their values to a known distribution[4]. Methodologically, Corbin et al. (1975) differed from our approach by always leaving visible all of the previous numbers in a sequence, and by creating a contrived problem set to test their specific research hypotheses. If people were sensitive to the constraints involved in generating the Corbin et al. (1975) problems, it is possible that their decision-making involving additional inferences about the generating process, requiring complexities in modeling that are not required to account for the current data.

All of that said, we do view the assumption that previous values do not affect decisions as only a first approximation. It would be an interesting and worthwhile exercise to extend our model to use previous values in some way, and examine to what extent this information can improve the predictive capabilities of the model. This modeling extension ought to be coupled with more ecologically representative assumptions about how environments generate sequences of choice values. While independence is one possibility, it seems like many real-world sequential decision-making problems will naturally have richer structures, to which people may well be sensitive.

6.2.2. Lack of Learning

Perhaps our most surprising empirical finding is the lack of learning, as summarized in Figure 7. One obvious possible reason involves the lack of feedback or financial incentive given to subjects. It would be interesting to observe what, if any, changes in adherence to the optimal decision rule, arise from providing various sorts of corrective feedback. We note, however, that even though we did not explicitly provide any feedback, the nature of the problem means that often subjects will have known, or be able to make a good inference about, whether their choice was correct. Every time the last value in the sequence is selected as a forced-choice, knowledge of the previous four values indicates whether or not the decision was correct. Less exactly, but still very informatively, the choice of a very large value at any position in the sequence is almost certainly correct. Presumably, on these problems, subjects are able to use this information to affirm their decision making in not choosing lesser, but relatively high, earlier values. Given the availability of this sort of corrective information, it is not obvious that the provision of explicit feedback will lead to large learning effects.

More speculatively, we think it is possible that the mode of stimulus presentation plays a role in preventing learning. Because the values are presented as typed four digit numbers, it seems likely that any explicit rule people could apply and adapt would be inherently language-based. For example, the decision thresholds inferred for Subject 1 in Figure 12 seems well expressed by a rule like "choose the first value above 80". It seems unlikely, however, that verbal rules of this type are easily adjusted in the ways required for incremental learning, because it is an effortful cognitive process to map words like "eighty" and "seventy-nine" on to the magnitudes they represent. It would be extremely interesting to examine human peformance on the same problems used here, but presented in a continuous and spatial form that allowed more naturally adjusted decision rules. For example, it is possible that being required to choose the longest of five perceptually presented lines, with lengths ranging from 0 to 100 units, would facilitate learning.

6.2.3. Individual Differences

The results presented here dealt with individual differences at both the model family level and the model level in largely descriptive and exploratory ways. This is deficiency arose because our generative process described only how an individual produced threshold models and made decisions, and so all modeling was done at the level of individual subjects. While this is not inappropriate, given the empirical observation of clear individual differences, the structure in individual results at the model family level and model level, as summarized in Figures 11 and 12, suggests that greater sophistication could be rewarded.

What is needed is an extended generative account that describes the processes that lead to individuals being different and the same as they develop and apply threshold models to optimal stopping problems. The formal modeling of individual differences in cognitive processes has been a focus of recent research (e.g., Lee & Webb, 2005; Navarro, Griffiths, Steyvers, & Lee, in press; Rouder, Sun, Speckman, Lu, & Zhou, 2003), and provides several ideas that might be applied to the current problem. For example, to the extent that there are different strategies evident at the model family level in Figure 11, it might be worthwhile identifying clusters of subjects in terms of the θ parameter (as per Lee & Webb, in press), and seeking some meaningful interpretation of the differences.

7. Conclusion

Optimal stopping problems like the one considered here provide an interesting and useful window onto human decision-making and problem-solving. The task is representative of many real world problems, in its requirement to reason under uncertainty, yet is easily described and controlled in the laboratory. The task is amenable to a formal solution, yet the optimal decision rule is beyond the reach of cognitive capabilities, and so requires people to employ heuristic solutions.

In this paper, we developed a generative model in which people use a series of thresholds to make decisions, choosing the first value presented that is currently maximal and above the threshold for that position in the sequence. Our main evaluation of the model focused on its predictive capability, demonstrating that it is able to predict the future decisions of subjects with considerable accuracy based on observing them solve just a few problems. In a more exploratory way, modeling results suggest interesting patterns of individual differences, both in terms of the types of threshold models they employ, and the values of the thresholds they use. Future research ought to extend the hierarchical model to embody a psychological theory of these individual differences.

Notes

1. The raw data are available as an on-line appendix on the Cognitive Science Society's website at http://www.cognitivesciencesociety.org/supplements/.
2. Actually, a small range of parameterizations, over which we average.
3. For example, a model where thesholds are allowed to rise. Or, perhaps more plausibly, one with all the thresholds being different, but the decrease from the first to the second being equal to the decrease from the third to the fourth, and this decrease being different from that from the second to third.
4. We doubt it is identical to the rank order version as they claim, because it seems likely subjects will use prior assumptions about likely distributions for the presented numbers in the absence of any relevant instructions, and so will not internally represent the numbers simply according to their rank order.

Acknowledgments

This paper was much improved by two rounds of thoughtful and provocative reviews, for which I thank Josh Tenenbaum, Peter Todd, and the two anonymous reviewers. I also thank Nick Burns, Gary Ewing, Tess Gregory, Peter Grünwald, Jay Myung, Dan Navarro, Michael Paradowski, Doug Vickers, Matt Welsh, and Robyn Whibley. The first experiment was conducted in 2003 as part of Chloë Mount's Honours thesis in Psychology at the University of Adelaide. The second experiment was also conducted by Chloë Mount, supported by a Faculty of Health Sciences grant from the University of Adelaide. I also acknowledge the support of Australian Research Council grant DP0451793.

References

Corbin, R. M. (1980). The secretary problem as a model of choice. *Journal of Mathematical, 21,* 1–29.

Corbin, R. M., Olson, C. L., & Abbondanza, M. (1975). Contest effects in optimal stopping rules. *Organizational Behavior and Human Performance, 14,* 207–216.

Cox, R. T. (1961). *The algebra of probable inference*. Baltimore, MD: Johns Hopkins University Press.

Dudey, T., & Todd, P. M. (2001). Making good decisions with minimal information: Simultaneous and sequential choice. *Journal of Bioeconomics*, *3* (2–3), 195–215.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4* (3), 282–296.

Forsythe, G. E., Malcolm, M. A., & Moler, C. B. (1976). *Computer methods for mathematical computations*. New York: Prentice-Hall.

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

Griffiths, T. L., Baraff, E. R., & Tenenbaum, J. B. (2004). Using physical theories to infer hidden causal structure. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 500–505). Mahwah, NJ: Erlbaum.

Grünwald, P. D. (1998). *The minimum description length principle and reasoning under uncertainty.* University of Amsterdam: Institute for Logic, Language and Computation.

Grünwald, P. D. (1999). Viewing all models as 'probabilistic'. In *Proceedings of the twelfth annual conference on computational learning theory (COLT' 99).* Santa Cruz: ACM Press.

Hoeting, J., Madigan, D., Raftery, A., & Volinsky, C. (1999). Bayesian model averaging. *Statistical Science, 14,* 382–401.

Jaynes, E. T. (2003). *Probability theory: The logic of science* (G. L. Bretthorst, Ed.). New York: Cambridge University Press.

Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception & Psychophysics*, *2* (8), 374–376.

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2004). Learning domain structures. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 720–725). Mahwah, NJ: Erlbaum.

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization, 14,* 381–392.

Lee, M. D. (2004). An efficient method for the minimum description length evaluation of cognitive models. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 807–812). Mahwah, NJ: Erlbaum.

Lee, M. D., O'Connor, T. A., & Welsh, M. B. (2004). Human decision making on the full-information secretary probllem. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 819–824). Mahwah, NJ: Erlbaum.

Lee, M. D., & Webb, M. R. (2005). Modeling individual differences in cognition. *Psychonomic Bulletin & Review, 12,* 605–621.

MacGregor, J. N., & Ormerod, T. C. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics, 58,* 527–539.

Myung, I. J., Pitt, M. A., & Kim, W. J. (2005). Model evaluation, testing and selection. In K. Lambert & R. Goldstone (Eds.), *Handbook of cognition* (pp. 422–436). Thousand Oaks, CA: Sage.

Navarro, D. J., Griffiths, T. L., Steyvers, M., & Lee, M. D. (in press). Modeling individual differences using Dirichlet processes. *Journal of Mathematical Psychology.*

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, *109* (3), 472–491.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107* (2), 358–367.

Rouder, J. N., Sun, D., Speckman, P. L., Lu, J., & Zhou, D. (2003). A hierarchical Bayesian statistical framework for response time distributions. *Psychometrika*, *68* (4), 589–606.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the "Secretary Problem". *Organizational Behavior and Human Decision Processes*, *69* (3), 221–236.

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

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Simon, H. A. (1976). From substantive to procedural rationality. In S. J. Latsis (Ed.), *Method and appraisal in economics* (pp. 129–148). London: Cambridge University Press.

Simon, H. A. (1982). *Models of bounded rationality.* Cambridge, MA: MIT Press.

Steblay, N. M., Deisert, J., Fulero, S., & Lindsay, R. C. L. (2001). Eyewitness accuracy rates in sequential and simultaneous lineup presentations: A meta-analytic comparison. *Law and Human Behavior, 25,* 459–474.

Vickers, D., Bovet, P., Lee, M. D., & Hughes, P. (2003). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean traveling salesperson problems. *Perception*, *32* (7), 871–886.

Vickers, D., Butavicius, M. A., Lee, M. D., & Medvedev, A. (2001). Human performance on visually presented traveling salesman problems. *Psychological Research, 65,* 34–45.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22* (4), 503–519.

Appendix
Entropification Under 0–1 Loss

Suppose a deterministic model M is being evaluated using a dataset D that has n observations, $D = [d_1, \ldots, d_n]$. Each of the observed data are discrete, and can assume only k different values. The model uses p parameters $\theta = (\theta_1, \ldots, \theta_P)$ to make predictions $Y = [y_1, \ldots, y_n]$. To evaluate any prediction made by the model, a 0–1 loss function is defined as $f(D,Y) = \sum_{i=1}^{n} \gamma_i$, where $\gamma_i = 0$ if $d_i = y_i$ and $\gamma_i = 1$ otherwise. By considering all possible parameterizations, the model makes a total of N different predictions. In other words, there are N different predictions, $Y_1, \ldots, Y_N$, the model is able to make about the data by choosing different parameter values. In general, the relationship between parameterizations and predictions will be many-to-one. This means that every unique model prediction is naturally associated with one or more parameterizations of the model.

Under these assumptions, Grünwald (1999) shows that using entropification the model making prediction Y can be associated with a probability distribution, parameterized by the scalar w, as follows:

$$p(D|M,Y,w) = \frac{e^{-wf(D,Y)}}{\sum_{x_1=1}^{k} \cdots \sum_{x_n=1}^{k} e^{-wf(D,[x_1,\ldots,x_n])}}.$$

Determining the MDL criterion for the model requires finding the model predictions Y^* and scalar w^* that *jointly* maximize $p\,(D \mid M, \theta, w)$ to give the value p^*. Once this is achieved the MDL criterion for the model is given simply by MDL = $-\ln p^* + \ln N$.

There is an obvious difficulty in maximizing $p\,(D \mid M, \theta, w)$. The problem is that the denominator given by $Z = \sum_{x_1=0}^{k} \cdots \sum_{x_n=0}^{k} e^{-wf(D,[x_1,\ldots,x_n])}$ involves considering every possible data set that could be observed, which involves a total of k_n terms. In cognitive science, where it is possible for a deterministic model to be evaluated using many data points, each of which can assume many values, the repeated calculation of Z may be too computationally demanding to be practical.

Lee (2004) derived a simpler form for Z can be derived by noting that $f\,(D, Y)$ can only take the values $0, \ldots, n$, in accordance with how many of the model predictions agree with the data. Since Z considers all possible data sets, the number of times $n - x$ matches (i.e., x mismatches) will occur is $\binom{n}{x}(k-1)^x$. For a prediction Y that has $n - m$ matches with the data (i.e., there are m mismatches and $f\,(D, Y) = m$), this leads to the simplification

$$p(D|M,Y,w) = \frac{e^{-wm}}{\sum_{x=0}^{n} \binom{n}{x} (k-1)^x e^{-wx}},$$

which has a denominator that sums $n + 1$ rather than k^n terms.

The computational efficiency offered by this reformulation means it will generally be possible to find the w_i^* that maximizes $p(D \mid M, Y_i, w_i)$, giving p_i^*, for all N model predictions. The p^* required for MDL calculation is then just the maximum of $p_1^*, \ldots, p_N^*$.

Finding each w_i^* can also be done efficiently by observing that

$$\partial p / \partial w = \frac{1}{Z^2} e^{-wm} \sum_{x=0}^{n} \binom{n}{x} (k-1)^x (x-m) e^{-wx}.$$

This derivative is clearly always positive if $m = 0$ and always negative if $m = n$.

This means, if a model predicts all of the data correctly, $w_i^* \rightarrow \infty$, and if a model fails to predict any of the data correctly $w_i^* \rightarrow -\infty$. Otherwise, if $0 < m < n$, the substitution $u = e^{-w}$ allows w_i^* to be found from the positive real roots of the degree n polynomial

$$\sum_{x=0}^{n} \binom{n}{x} (k-1)^x (x-m) u^x.$$

by standard numerical methods (e.g., Forsythe, Malcolm, & Moler, 1976).