Title
Modeling choice and search in decisions from experience: A sequential sampling approach

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors
Markant, Douglas
Pleskac, Timothy J
Diederich, Adele
et al.

Publication Date
2015

Peer reviewed

Modeling choice *and* search in decisions from experience: A sequential sampling approach

Douglas Markant (markant@mpib-berlin.mpg.de)
Timothy J. Pleskac (pleskac@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Center for Adaptive Rationality,
Lentzeallee 94, 14195 Berlin, Germany

Adele Diederich (a.diederich@jacobs-university.de)
Department of Psychology, Jacobs University Bremen gGmbH
P.O. Box 750 561, 28725 Bremen, Germany

Thorsten Pachur (pachur@mpib-berlin.mpg.de)
Ralph Hertwig (hertwig@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Center for Adaptive Rationality,
Lentzeallee 94, 14195 Berlin, Germany

Abstract

In decisions from experience (DFE), people sample from two or more lotteries prior to making a consequential choice. Although existing models can account for how sampled experiences relate to choice, they don't explain decisions about how to search (in particular, when to stop sampling information). We propose that both choice and search behavior in this context can be understood as a sequential sampling process whereby decision makers sequentially accumulate outcome information from each option to form a preference for one alternative over the other. We formalize this process in a new model, Choice from Accumulated Samples of Experience (CHASE). The model provides a good account of choice behavior and goes beyond existing models by explaining variations in sample size under different task conditions. This approach offers a process-level framework for understanding how interactions between the choice environment and properties of the decision maker give rise to decisions from experience. **Keywords:** decisions from experience; sequential sampling; decision field theory

How do people decide between options whose values are uncertain? One common recourse is to obtain more information prior to making a choice. For example, before buying a car, one might test drive vehicles from different manufacturers, consult friends for their opinions, or even lease a car for a year. The individual decision maker can often control both the *amount* and *source* of experiences with available options prior to making a choice. When people play such an active role in information gathering, their knowledge of alternative options, and their ultimate choice between them, directly depends on how they explored during learning. As a result, a long-standing question for researchers, policy-makers, and marketers alike has been whether people collect the right amount and kinds of information in order to make these decisions from experience (*DFE*).

A common experimental tool for studying DFE is the *sampling paradigm*, in which people decide how many outcomes to observe from a set of gambles of uncertain value before making a choice between them (Hertwig & Erev, 2009). Participants collect information one sample at a time, where a single sample is an outcome generated probabilistically from an option according to some underlying distribution. For example, one option may produce a reward of $3 every time it is sampled, whereas a second, more risky, option produces $4 in 80% of draws and $0 otherwise. In contrast to decisions from description, in which the possible outcomes and their probabilities are provided upfront to the decision maker, in the sampling paradigm the participant learns about the options in a self-directed manner, deciding from which options to sample and how many samples to draw.

A key question in DFE is how much information people sample and how that experience influences their ultimate choice. A primary measure of interest is the *sample size*, or the number of draws taken prior to making a final choice. Previous work suggests that people tend to draw relatively few samples, which can be an adaptive strategy given the structure of the learning problem (Hertwig & Pleskac, 2010). However, small samples can also lead to predictable distortions in the perceived value of options, such that rare outcomes are underweighted relative to their objective probability (a reversal of the overweighting typically observed in decisions from description, resulting in what has been termed a *description–experience gap*).

A number of models have been proposed to account for DFE in the sampling paradigm (Erev et al., 2010; Gonzalez & Dutt, 2011; Hau, Pleskac, Kiefer, & Hertwig, 2008). Although these models are often quite successful at predicting choices, a major limitation is that they do not account for how people decide to terminate sampling. For example, a common assumption is that the number of samples collected is determined by a single distribution that is independent of the problem (e.g., Gonzalez & Dutt, 2011). However, in the sampling paradigm and similar tasks, people adjust the amount of information they collect depending on their goals or their experience with the task. For example, sample size increases when people experience higher variability in the outcomes of an option (Lejarraga, Hertwig, & Gonzalez, 2012; Pachur & Scheibehenne, 2012), when the decisions involve higher stakes (Hau et al., 2008), and when the cost of collecting information is low (Rapoport, 1969).

In this paper, we propose a new model — Choice from Accumulated Samples of Experience (CHASE) — that models the sampling and choice process in DFE as a sequential sampling process. Sequential sampling models have been widely applied to decision making tasks to understand how information sampling interacts with cognitive processes, giving rise to dynamic patterns in choices and stopping times (Busemeyer & Townsend, 1993; Edwards, 1965; Nosofsky & Palmeri, 1997; Ratcliff, 1978). The key principle of this class of models is that people sequentially accumulate samples of evidence that favor one or another option. Our model builds on the key assumption of decision field theory (Busemeyer & Townsend, 1993; Diederich, 2003), which is that the rate of evidence accumulation is tied to the properties of the option, defined both objectively (e.g., the variance of outcomes) and in terms of the decision maker's subjective evaluation (e.g., encoding, attentional fluctuation, or differential weighting of gains vs. losses). However, our model differs from decision field theory in at least two respects. First, it uses cumulative prospect theory (CPT; Tversky & Kahneman, 1992) to model the attentional and subjective evaluation process. Second, it predicts the discrete number of samples people draw rather than response times, such that evidence is accumulated until a threshold level of preference is reached, indicating the option to be chosen. Thus, this approach provides a framework within which variations in external choice environments and individual decision makers can be related to how much information they collect and which options they ultimately choose. In this paper, we describe the CHASE model, showing how it can capture the dependence of sample size and choice in the well-studied sampling paradigm. We then apply the model to a number of existing datasets to test hypotheses about the sampling and choice process and their interrelationship.

Choice from Accumulated Samples of Experience (CHASE) model

Sampling paradigm

In the sampling paradigm, each decision problem typically involves a choice between two gambles H and L with higher and lower expected values, respectively. The gambles typically offer a chance of obtaining outcome x with probability q, otherwise y $(x,q;y)$. Participants are not told about the properties of the gambles, but can learn about their payoff distributions using a set of buttons. Pressing a button produces a random draw from the respective gamble's payoff distribution. Participants are instructed to draw as many samples as they wish and only then decide from which distribution to make a single draw resulting in a real payoff.

The CHASE model accounts for both choice and sample size through their relationship to an underlying process of preference formation. In contrast, most if not all existing models of DFE focus on predicting final choices between H and L alone, often by taking into account the statistics of the experienced sample. However, choices (at the aggregate level) can often be predicted with a high degree of accuracy without relying on observed sampling sequences (Erev et al., 2010). Accordingly, we test the CHASE framework with choice and sample size data, and we subsequently evaluate its ability to account for a number of existing findings in the sampling paradigm. Regardless of whether we model aggregate or individual-level data, one distinct advantage of our approach is that the theory is specified analytically, allowing us to fit the model with maximum likelihood estimation.

The model

According to the model, decisions are based on a preference state representing the relative preference for gamble H over gamble L. The preference state evolves over time via samples of valence information about the options, and a decision is made as soon as the preference state exceeds a threshold θ for one of the choice options. This threshold θ is adopted by the respondents at the outset of a decision problem and determines the preference level necessary to terminate sampling and choose an option. The H option is chosen when the preference reaches the threshold θ and the L option is chosen when the preference reaches the threshold $-\theta$. The thresholds define a discrete state space of size m ranging from $-\theta$ to θ representing different levels of preference.

We assume that when a sample is drawn from an option, it is evaluated and compared with the mean valence for the other option, and it is this relative valence that accumulates over multiple sampling trials. This relative evaluation process is characterized by the drift rate d, which is the average trajectory that preference takes over time. In particular,

$$d = \frac{V_H - V_L}{c\sqrt{\sigma^2}}, \tag{1}$$

where V_H and V_L are the mean valences for each option, σ^2 is the variance of the valences, and c is a constant scaling factor (for present purposes it is fixed to 2).

We model the subjective evaluation process with CPT, thereby specifying the mean valence and variance in terms of CPT's utility and weighting function (Tversky & Kahneman, 1992). One advantage of this approach is that the form of CPT's probability weighting function and the associated implications for choice have been of central interest to experience-based decision research. Thus, as we will discuss later, the CHASE model provides a framework to model the weighting function while controlling for different decision thresholds and other factors relevant to the sampling process.

Due to the rank-dependent nature of CPT, the mean valence of a gamble $(x,q;y)$ for a set of gains $x > y > 0$ is

$$V = w(q)u(x) + [w(1.0) - w(q)]u(y). \tag{2}$$

For the corresponding functions for losses and mixed gambles, refer to Tversky and Kahneman (1992). The utility function $u(x)$ follows the form of CPT with

$$u(x) = \begin{cases} x^{\alpha}, & \text{if } x \geq 0, \\ -\lambda(-x)^{\alpha}, & \text{if } x < 0, \end{cases} \tag{3}$$

where the parameter α defines the curvature of the utility function for both gains and losses; λ determines the degree of loss aversion.

The utilities of each outcome are multiplied by the probability weights, which are obtained for positive outcomes (i.e., gains) by transforming the decumulative probabilities with a weighting function w. Following Prelec (1998), we use the two-parameter function:

$$w(q) = exp(-\delta(-ln\, q)^{\gamma})). \tag{4}$$

The γ controls the curvature of the weighting function, which determines the sensitivity to changes in probabilities. The δ controls the elevation of the weighting function to produce pessimistic or optimistic weights.

Finally, the variance of the valences σ^2 is

$$\sigma^2 = s_H^2 + s_L^2, \tag{5}$$

where s_H^2 and s_L^2 are the perceived variance of each option, which is equal to

$$s^2 = w(q)[u(x) - V]^2 + [w(1.0) - w(q)][u(y) - V]^2. \tag{6}$$

We model the change in preference resulting from sequential evaluations of samples as a *birth–death chain* in which each observed outcome results in one of three effects: a transition up one state (preference in favor of the H option), a transition down one state (preference in favor the L option), or staying in the same state (no change) (Diederich & Busemeyer, 2003). These probabilities are given by

$$p_{i,j} = \begin{cases} \frac{1-p_{stay}}{2}(1-d), & \text{for } j-i = -1, \\ \frac{1-p_{stay}}{2}(1+d), & \text{for } j-i = +1, \\ p_{stay}, & \text{for } j = i, \\ 0, & \text{otherwise,} \end{cases} \tag{7}$$

where the transition probability $p_{i,j}$ represents the probability of moving from state i to state j.[1] p_{stay} is a free parameter that dictates the probability of remaining in the same state. The transition probabilities are represented in a transition probability matrix $\mathbf{P}$. An example with $\theta = 2$ and therefore a state space of $S = \{-2, -1, 0, +1, +2\}$ is

$$\mathbf{P} = \begin{array}{ccc} \textit{index} & & \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ & \textit{state} & \begin{array}{ccccc} -2 & -1 & 0 & +1 & +2 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} & \begin{array}{c} -2 \\ -1 \\ 0 \\ +1 \\ +2 \end{array} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ p_{21} & p_{22} & p_{23} & 0 & 0 \\ 0 & p_{32} & p_{33} & p_{34} & 0 \\ 0 & 0 & p_{43} & p_{44} & p_{45} \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{array}. \tag{8}$$

[1] Equation 7 requires that the drift rate d (Equation 1) is bounded by -1 and 1.

To calculate the predicted choice probability and sample size distribution, we rearranged the transition matrix into the following form:

$$\mathbf{P} = \begin{bmatrix} \mathbf{P_I} & 0 \\ \mathbf{R} & \mathbf{Q} \end{bmatrix}, \tag{9}$$

where $\mathbf{P_I}$ is a 2×2 identity matrix (corresponding to absorbing states θ and $-\theta$), $\mathbf{R}$ is an $(m-2) \times 2$ matrix containing the one-step probabilities of transitioning into an absorbing state, and $\mathbf{Q}$ contains transition probabilities between transient (i.e., non-absorbing) states.

Finally, an initial distribution defined by the vector $\mathbf{Z}$ with individual elements z_i gives the probability of starting in each transient state S_i and can be used to model bias towards one option or the other. However, when choice options are coded as H and L (whose identities are not known to the learner), an unbiased initial distribution is more appropriate. We assumed that the initial state was distributed according to a normalized exponential function of the distance from the state 0,

$$z_i = \frac{exp(-|S_i|/\tau)}{\sum_{j=2}^{m-1} exp(-|S_j|/\tau)} \tag{10}$$

for $i = 2, 3, ..., m-1$, where τ is a temperature parameter controlling how peaked the distribution is at state 0.

The probability of choosing the H or L option is

$$[P(H), P(L)] = \mathbf{Z} \cdot (\mathbf{I} - \mathbf{Q})^{-1} \cdot \mathbf{R} \tag{11}$$

and the conditional probability of stopping at a particular sample size n is defined by the first passage distribution

$$[P(N = n|H), P(N = n|L)] = \mathbf{Z} \cdot \mathbf{Q}^{n-1} \cdot \mathbf{R}./[P(H), P(L)], \tag{12}$$

where the ./ denotes element-wise division. For derivations of the choice probabilities and the first passage distribution, see Busemeyer and Diederich (2010).

Model predictions

The CHASE model makes a number of predictions with respect to the sampling paradigm and can capture qualitative effects on sample size that have previously been reported. For instance, people tend to sample longer when faced with gambles with high variance (Lejarraga et al., 2012; Pachur & Scheibehenne, 2012). This is predicted by the model because the drift rate decreases with increasing variance (see Equation 1), causing the process to take longer to reach a fixed threshold relative to low variance options. Similarly, increased sample sizes in the face of potential losses (Lejarraga et al., 2012) could be accounted for by a decrease in the drift rate with increased loss aversion (Equation 3).

By virtue of the weighting function in Equation 4, the model may also be used to examine a long-standing question in DFE research: To what extent is a distortion of objective

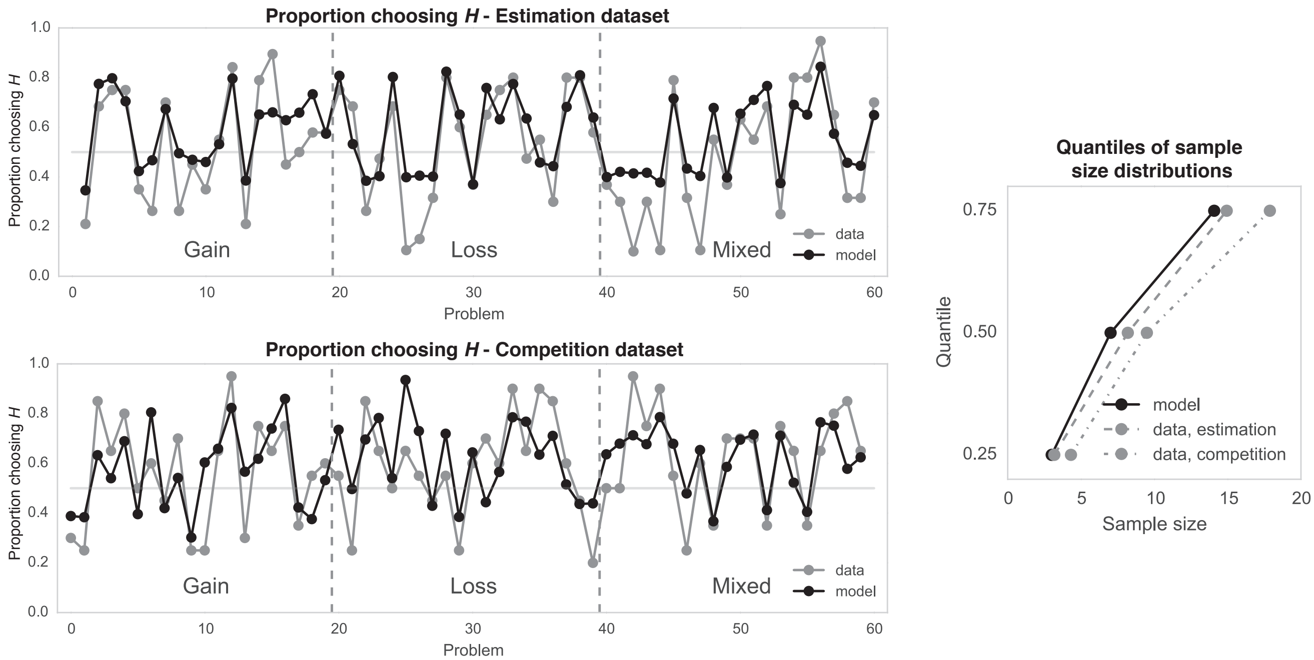

Figure 1: *Left:* Proportion choosing the *H* option and model prediction across 60 problems in the TPT estimation (top) and competition (bottom) datasets. Problems are divided by domain (gains, losses, and mixed gambles). *Right:* Quantiles of observed (gray) and predicted (black) sample size distributions.

probabilities (particularly with respect to rare outcomes) related to not only choice behavior (Fox & Hadar, 2006; Hertwig et al., 2004) but also decreased sample sizes (as would be predicted if such distortion decreases the perceived variance of the options)? In addition, like other sequential sampling models of decision making, the model can account for speed–accuracy tradeoffs by assuming that the decision threshold can vary according to the decision maker's goals or other features of the task. For example, a person may adopt a higher threshold when greater rewards are at stake, leading to larger sample sizes and potentially a higher likelihood of choosing the better option.

As a first step toward evaluating these predictions, we fitted the model at an aggregate level to several existing datasets and tested whether it can account for both the observed choice proportions and distributions of sample sizes. First, using data from the Technion Prediction Tournament (Erev et al., 2010), we tested the model on a relatively heterogeneous set of choice problems, including a large proportion of problems involving rare outcomes. Second, we examined whether changes in behavior resulting from a manipulation of the magnitude of rewards (Hau et al., 2008), an effect that existing models are unable to explain, can be accounted for by differences in decision thresholds under the CHASE model.

Fitting datasets from the sampling paradigm

Technion Prediction Tournament (TPT)

Erev et al. (2010) conducted a modeling tournament with two datasets that were collected using the sampling paradigm: (1) an estimation set used to fit models, and (2) a competition set that was then used to evaluate the models' ability to predict choices in a different group of problems. Each dataset contained 60 problems, with 20 problems each in the loss, gain, and mixed domains. Each individual problem involved a choice between a safe option with a single outcome and a risky option with two possible outcomes. Problems in both datasets were created using the same algorithm by randomly generating outcomes and probabilities, with the constraint that two-thirds of the risky options involved a rare outcome ($q < .1$ or $q > .9$).

Given that the TPT data has been used to evaluate a wide variety of models (Erev et al., 2010; Gonzalez & Dutt, 2011), it serves as a useful benchmark to validate the CHASE model on a representative dataset in the sampling paradigm. We followed the same approach as in Erev et al. (2010) to fit the model using the estimation set and then evaluate its predictions for the competition dataset. Importantly, however, the CHASE model was fitted to maximize the joint likelihood of

Table 1: BIC scores from model fitting

		TPT	Hau et al. (2008)	
Utility	Weighting	Estimation set	Exp. 1	Exp. 2
–	–	9282	1954	2504
–	+	**9183**	**1906**	**2496**
+	–	9280	1907	2498
+	+	9200	1915	2507

Note: A '+' indicates that the corresponding function's parameters were fit in the model, otherwise they were fixed at 1.

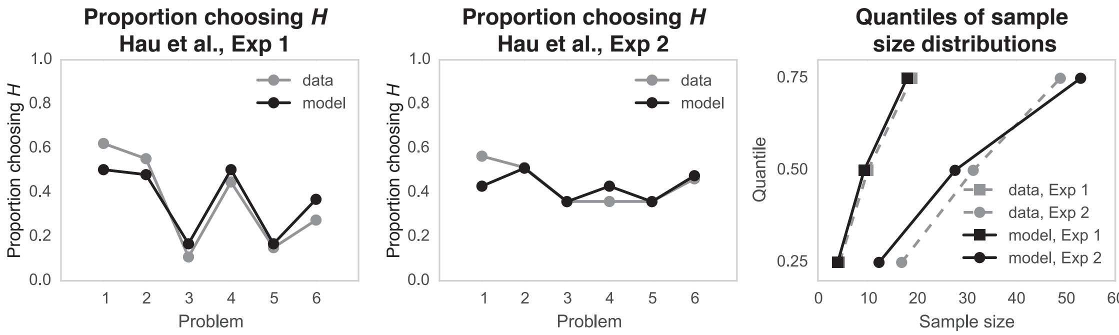

Figure 2: Choice proportions and sample size distributions (averaged across six problems) for best-fitting models compared with data from Hau et al. (2008).

both choices and sample sizes across all problems. This distinguishes the current approach from other models that use a fixed sample size distribution as the basis for predictions and that are fit to choice data alone.

A combination of grid search (for the discrete θ parameter) and numerical optimization (Nelder–Mead, for all other parameters) was used to find the parameter values that maximized the log-likelihood of the data under a given model (summed across all observations). We compared four CHASE models derived from a factorial combination of the utility and weighting functions described above, and used the Bayesian information criterion (BIC) to assess model fit (see Table 1). The parameter values of the best-fitting model are shown in Table 2. The best-fitting model included nonlinear probability weighting ($\gamma = 1.4$, $\delta = 1$) but a linear utility function, and high variability in the starting position ($\tau = 40$, the upper limit used in fitting, indicating a flat initial distribution), for a total of five free parameters. Note that the best-fitting value of $\gamma = 1.4$ corresponds to a probability weighting function with underweighting of low probabilities and overweighting of high probabilities.

The choice proportions and sample size distributions for the best-fitting model are shown in Figure 1. For the purposes of comparison against existing models, for choice proportions we computed the mean-squared deviation *MSD*, Pearson correlation r, and proportion of agreement P_{agree} (proportion of problems in which the model assigns the more frequently observed choice a probability higher than .5). The model was highly accurate for the estimation dataset ($P_{agree} = .94$, $r = .88$, $MSD = .019$) and was comparable to other top models in the tournament. For the competition dataset, the results were $P_{agree} = .9$, $r = .68$, $MSD = .022$. Relative to the top models from the tournament (some of which had many more free parameters), the CHASE model resulted in slightly lower r and *MSD* values, but a higher P_{agree} than any of the reported models ($P_{agree} = .83$ for the winner of the tournament).

Thus, the CHASE model can account for choices in a large, representative set of DFE problems, is competitive with other models when predicting performance on an independent test set, and can also capture the sample size distributions in both datasets. The best model demonstrates the influence of sampling error on choices, which is sensible given the high proportion of problems in these datasets that involved rare outcomes unlikely to be experienced with small numbers of samples. Together with this underweighting of low-probability outcomes, a simple preference accumulation process can account for both the observed distribution of sample sizes as well as variability in choice proportions across problems.

Accounting for longer search under higher stakes

Although a number of existing models perform well on the TPT datasets, they all lack a mechanism that can account for changes in sample size distributions. As such, they cannot explain variations in sample size that occur due to manipulations of instructions, the structure of the option environment, or properties of the individual decision maker.

One such effect was reported in a study by Hau et al. (2008), whereby a manipulation of reward magnitude (by a factor of 10) from Experiment 1 to Experiment 2 was associated with an increase in sample size. A single set of problems was used in both experiments (see Table 3), where each gamble involved a nonzero outcome that occurred with probability q (and zero otherwise). Three problems (3, 5, and 6)

Table 2: Best-fitting parameter estimates

	TPT	Hau, Exp 1.	Hau, Exp 2.
Choice threshold θ	2	3	5
Start point variability τ	40	40	2.46
Probability of staying (p_{stay})	.68	.49	.46
Weighting function γ	1.41	1.15	.92
Weighting function δ	1	1.61	1.30

Table 3: Six problems used in Hau et al. (2008)

Decision problem	x_H	q_H	x_L	q_L
1	4	0.8	3	1
2	4	0.2	3	0.25
3	−3	1	−32	0.1
4	−3	1	−4	0.8
5	32	0.1	3	1
6	32	0.025	3	0.25

involved a large outcome that was relatively rare ($q \leq .1$). Whereas people's choices in Experiment 1 were consistent with underweighting of these rare outcomes, in Experiment 2 people were more likely to choose the H option in those same problems.

According to the CHASE model, these changes are predicted to occur if people adopt a higher decision threshold when larger rewards are at stake. To evaluate whether a difference in thresholds can account for this pattern, we fitted the model separately to each experiment using the same procedure as above. As in the TPT datasets, the best-fitting model according to BIC included parameters for the probability weighting function only (Table 1). Most importantly, the best-fitting values for the decision threshold θ differed between the two experiments as expected, with participants in Experiment 2 best described by a higher threshold. The models successfully capture the overall shift in the distribution of sample sizes as well as a higher proportion of H choices in those problems involving rare outcomes (Figure 2).

Discussion

The sampling paradigm has become a useful tool for studying how people learn about uncertain options prior to making a choice, and how the dynamics of sequential experience can lead to different choices than those made in description-based settings. However, there are no existing theories that account for these dynamics, particularly with respect to how long people decide to explore. A principal advantage of the CHASE model is that it offers a framework for predicting both sample size and choice across a wide range of conditions in this paradigm, as well as a way to interpret such variation in terms of psychological processes related to the evaluation of observed outcomes and comparison of those outcomes against decision thresholds.

Future work will compare the predictions of the model against alternative theories of how people make choices in this task. Although existing models do not account for sample size in general, they are capable of achieving high accuracy in terms of choice data alone, suggesting that the TPT datasets may not be diagnostic as to the underlying processes involved in making experience-based choice. One important goal is to expand the modeling framework presented here to compare the role of alternative stopping and preference processes in predicting behavior, both at the aggregate level and for individual participants.

References

Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Thousand Oaks, CA: SAGE Publications.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, *100*(3), 432.

Diederich, A. (2003). Decision making under conflict: Decision time as a measure of conflict strength. *Psychonomic bulletin & review*, *10*(1), 167–176.

Diederich, A., & Busemeyer, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *Journal of Mathematical Psychology*, *47*(3), 304–322.

Edwards, W. (1965). Optimal strategies for seeking information - models for statistics, choice reaction-times, and human information-processing. *Journal of Mathematical Psychology*, *2*(2), 312-329.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., ... Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*(1), 15–47.

Fox, C. R., & Hadar, L. (2006). Decisions from experience = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, *1*(2), 159–161.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*(4), 523.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518.

Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in cognitive sciences*, *13*(12), 517–523.

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, *115*(2), 225–237.

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, *124*(3), 334–342.

Nosofsky, R., & Palmeri, T. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*(2), 266–300.

Pachur, T., & Scheibehenne, B. (2012). Constructing preference from experience: The endowment effect reflected in external information search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(4), 1108.

Prelec, D. (1998). The probability weighting function. *Econometrica*, 497–527.

Rapoport, A. (1969). Effects of observation cost on sequential search behavior. *Perception & Psychophysics*, *6*(4), 234–240.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, *85*(2), 59.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, *5*(4), 297–323.