Accumulation of Evidence and Information Search in Experiential Decisions

Katja Mehlhorn (katjam@cmu.edu)
Varun Dutt (varundutt@cmu.edu)
Noam Ben-Asher (noamba@andrew.cmu.edu)
Cleotilde Gonzalez (coty@cmu.edu)
Dynamic Decision Making Laboratory, Carnegie Mellon University
4609 Winthrop Street, Pittsburgh, PA 15213 USA

Keywords: information search, decisions from experience, sampling paradigm, cognitive model, evidence accumulation.

Imagine you are about to write a paper for an upcoming conference. Before you decide how to write it, you are likely to read a number of related articles. How many articles do you read? When do you stop searching for more articles and start writing? While we know much about how people make decisions, little is known about the search process that precedes a consequential decision. Here, we analyze participants search behavior in a binary choice task and present a cognitive model that is able to explain both, the process of information search, as well as the subsequent consequential choice.

Information search in decisions from experience

In the classical decision-making literature, decision-making has often been studied by explicitly presenting the decision maker with the information relevant for the decision. Thus, this literature largely ignores the process of information search. Recently, experimental paradigms have studied experiential decisions and allowed the investigation of the search process before decisions are made. One prominent example is the sampling paradigm (e.g., Hertwig, Barron, Weber, & Erev, 2004). In this paradigm, participants are presented with two options (visualized as two buttons on the screen) that are associated with monetary payoff distributions. For example, in the decision problem shown in Figure 1, the left button (risky button) yields a high payoff (17.1) with 10% probability and a low payoff (6.9) with 90% probability. The right button (safe button) yields a medium payoff (8) with 100% probability. Participants are not explicitly told about the payoffs or their probabilities. Instead, they are asked to sample from both options until they feel confident to make a consequential choice for the option they prefer.

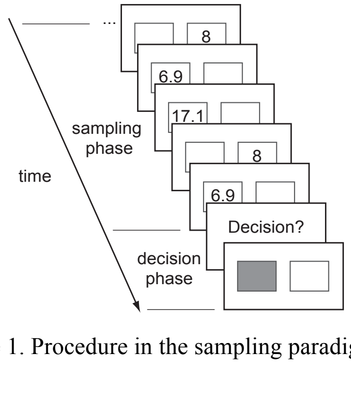

Figure 1. Procedure in the sampling paradigm.

Research has shown that, overall, people tend to take only a small number of samples (medians range from 9-19 in most studies; Erev et al., 2010; Hau, Pleskac, & Hertwig, 2010). This small sample size is surprising, given that the chance to correctly estimate the payoff distribution is likely to increase with the sample size (Gonzalez & Dutt, 2012). While the sample size seems to be affected by several characteristics of the decision maker, not much is known about how it is influenced by different properties of the task.

To learn more about when people stop searching for information, we analyzed data from the Technion Prediction Tournament (Erev, et al., 2010), which is the largest data set on the sampling paradigm currently available (79 participants, each solving 30 out of 120 problems). In agreement with the small sample size effect, Figure 2 shows that in the TPT data set the distribution of the sample size is heavily right-tailed. To take a closer look at how the sample size is affected by properties of the decision problems, we investigated two factors: experienced variability and payoff domain. As shown in Figure 3, the sample size increases when variability was experienced in a problem (i.e., when the risky button displayed two possible outcomes, as in Figure 1), compared to problems where no variability was experienced (the risky button displayed only one outcome); average medians: 11.1 vs. 15.3. Furthermore, the number of samples is lower if the observed payoffs constituted gains (as in Figure 1), rather than losses (where outcomes on both buttons were negative); average medians: 12.2 vs. 14.0. A linear mixed effects model, showed the effects of both factors to be significant (variability: $\beta=4.65$, $p<.001$; domain: $\beta=1.64$, $p=.001$).

Modeling information search

Recently, Gonzalez and Dutt (2011) have shown that a computational model based on instance-based learning theory (IBLT; Gonzalez, Lerch, & Lebiere, 2003) can not only explain how people make experiential decisions, but also explain how information is sampled before a consequential decision is made. The basic idea in the IBL model is that, during sampling, instances of the observed payoffs in both options are stored in memory. Behavior during sampling and at final choice depends on the experienced utility (or bended value) of the options. The experienced utility of an option is a function of its associated payoffs and the probability of retrieving these payoffs (i.e., instances) from memory, using a simplified ACT-R activation mechanism (Anderson & Lebiere, 1998).

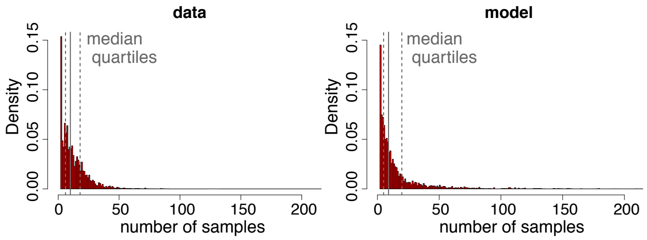

Figure 2. Distribution of the sample size. The x-axes are truncated at the maximum number of samples found in the human data (resulting in cutting off 2% of the tail produced by the model).

The IBL model explains, for example, how participants alternate between the choice options during sampling (for a detailed explanation of the model see Gonzalez & Dutt, 2011).

Although the IBL model is able to account for choices during sampling, the stopping point (i.e., the number of samples drawn during the sampling phase) was previously determined by a random draw from a geometric distribution function that was fitted to behavioral data in TPT. Here, we introduce a stopping mechanism in the model that is grounded in psychological literature related directly to information foraging. This mechanism is motivated from evidence accumulation models (for an overview see e.g., Ratcliff & Smith, 2004). It assumes that people accumulate experienced utilities until a decision criterion is reached. Once this criterion is reached, sampling is stopped and a decision is made based on the experienced utilities as in the original formulation of the IBL model. The revised model is thus identical to the model reported in Gonzalez and Dutt (2011), with the exception that the stopping point is now determined by an evidence accumulation process, rather than by a mathematical distribution function.

To fit the revised model, we calibrated the decision criterion (which consists of an upper bound for positive and a lower bound for negative experienced utilities), to the median number of samples from the human data. All other parameters were left at the original values reported in Gonzalez and Dutt (2011). More specifically, we calibrated the bounds by using a Genetic Algorithm (Holland, 1975) to fit the model's median sample size to the median sample size for half of the human data. The resulting bounds were randomly drawn from U(0.0001, 14.18) for positive, and from U(-24.18, -0.0001) for negative values. Then, we generalized the model to the other half of the data.

With the model merely calibrated to fit the median sample size, we then investigated whether it would reproduce the distribution of the number of samples, as well as the effects of the two task-relevant factors (identified above). As shown in Figure 2, the model captures human sampling behavior well and it reproduces a similar and heavily right-tailed distribution, as it was found in the human data. Figure 3 shows the effects of experienced variability and payoff-domain. The model correctly shows the increase in sample size due to the experienced variability. The effect of domain is less clear in the model. Whereas it correctly predicts a higher number of samples for losses than for gains if

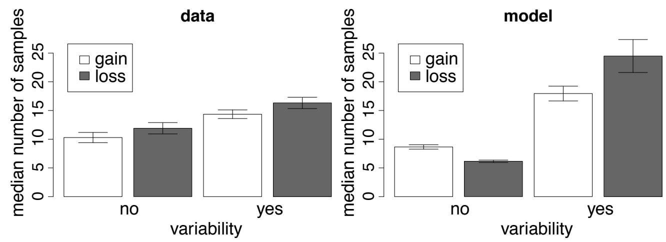

Figure 3. Sample size as a function of experienced variability and domain (shown are the average medians ±1 SE.)

variability is experienced, it does not produce this pattern if no variability is experienced.

Discussion

How much information do people search before making a consequential choice? Our results suggest that at least two task-related factors likely affect the amount of sampled information: Sample size is likely to increase with the experienced variability of outcomes, and it will be higher if losses, rather than gains are experienced. Furthermore, our results suggest that it is possible to extend the IBL model of experiential decisions by incorporating an evidence-accumulation mechanism that predicts when people stop sampling. While the revised model presented here presents a first step into this direction, there are several potent approaches to model the accumulation process. We are currently evaluating these approaches in more detail by exploring their ability to more accurately predict human information search.

References

Anderson, J. R., & Lebiere, C. (1998). The atomic components of thought. Mahway: NJ: Erlbaum.

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*, 15-47.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review, 118*, 523-551.

Gonzalez, C. & Dutt, V. (2012). Transition from Exploration to Exploitation in Decisions from Sampling. Unpublished paper under review.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*, 591-635.

Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making, 23*, 48-68.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15*, 534-539.

Holland, J. H. (1975). Adaptation in natural and artificial systems. Ann Arbor, MI: University of Michigan Press.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111*, 333-367.