Modeling adaptive exploration in decisions from experience: A sequential sampling approach

Markant, D.[a], Pleskac, T. J.[a], Diederich, A.[b], Pachur, T.[a], and Hertwig, R.[a]

a. Center for Adaptive Rationality, Max Planck Institute for Human Development
b. Department of Psychology, Jacobs University Bremen

Max-Planck-Institut für Bildungsforschung
Max Planck Institute for Human Development

Email: markant@mpib-berlin.mpg.de
Web: http://www.dougmarkant.com

How are people's choices shaped by the way they explore the environment? A key question in decisions from experience (DFE) is how people gather experiences with a set of available options prior to making a choice between them (e.g., how many times do you test-drive different cars before deciding which one to purchase?).

The *sampling paradigm* is a common experimental tool for studying DFE. Final choices are preceded by a period of exploration during which people sample experiences with individual options (i.e., outcomes that are generated according to underlying probability distributions for each option). Previous research on this task has shown that *sample size* (the number of draws prior to a choice) varies depending on a number of factors, including outcome variability (Lejarraga et al., 2012) and payoff magnitude (Hau et al., 2008). Although existing models offer potential explanations for how final choices are related to sampled experiences, they do not account for such examples of adaptive exploration.

We present a model—Choice from Accumulated Samples of Experience (CHASE)—which formalizes DFE as a sequential sampling process by which external exploration drives the accumulation of relative preference between two choice options. In the present study, we demonstrate that CHASE can account for both choice and sample size in an experiment based on the sampling paradigm. In addition, we test three key predictions of the model with respect to how people adapt their exploration in response to 1) sampling costs, 2) option uncertainty, and 3) contextual variability.

The sampling paradigm

- On each trial, participants must choose between two options *L* and *H* with lower and higher expected values, respectively. Each option is defined by a Normal distribution that is used to randomly generate outcomes.
- Each trial begins with an exploration stage, during which a participant draws individual outcomes from each option. Participants are free to sample options any number of times and in any order.
- When participants are ready to choose an option, they stop sampling and make a final choice. A final outcome is then generated from the chosen option to determine their score for the trial.

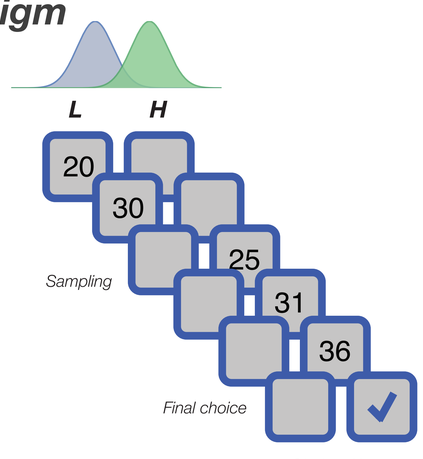

Choice: **H**
Sample size: **5**

CHASE: Choice from Accumulated Samples of Experience

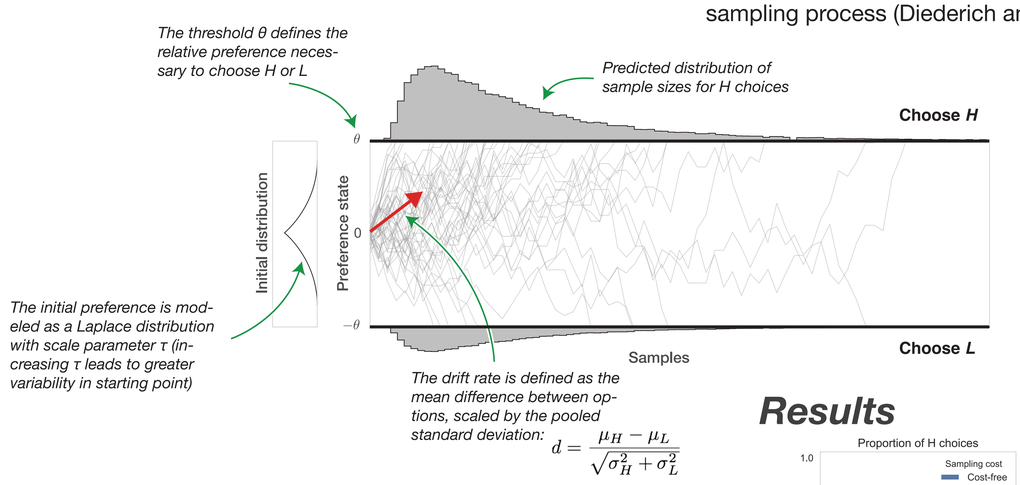

$$d = \frac{\mu_H - \mu_L}{\sqrt{\sigma_H^2 + \sigma_L^2}}$$

Under CHASE, exploration drives the accumulation of relative preference for option *H* over *L*. As outcomes are sampled from choice options, the preference state evolves according to a drift rate *d* until reaching one of two decision thresholds corresponding to each option (θ or -θ). Predicted choice probabilities and sample size distributions are derived from a matrix approximation to this sequential sampling process (Diederich and Busemeyer, 2003; Markant et al., 2015).

Predictions

1. People adjust their decision threshold (θ) based on the costs of information search (e.g., decrease threshold when sampling is costly, leading to lower *p(H)* and smaller sample sizes)
2. For a fixed EV difference, increased outcome variance leads to lower drift rates, causing lower *p(H)* and larger sample sizes (see right).
3. Increased starting point variability (τ) leads to smaller sample sizes when option *L* is chosen, particularly when *p(H)* is high (Ratcliff & Rouder, 1998)

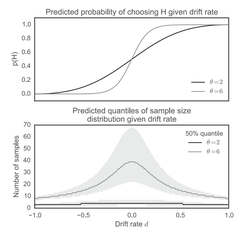

Experiment

- N=100 participants recruited on Amazon MTurk. Each participant played 24 games; bonused based on their final choices in 3 randomly chosen games.v

1 Sampling cost

- *within-subjects*
- *12 games per condition*

CHASE Prediction: *Decrease in decision threshold when sampling is costly*

2 Outcome variance

- *within-subjects*
- *12 games per condition*

CHASE Prediction: *Decreased p(H), increased sample size for high-variance problems*

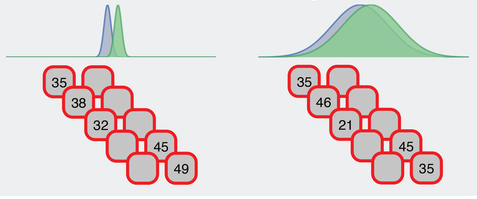

3 Contextual variability

- *between-subjects*
- *N=50 per condition*

CHASE Prediction: *Higher starting point variability when contextual cue (color) recurs across trials → faster L choices*

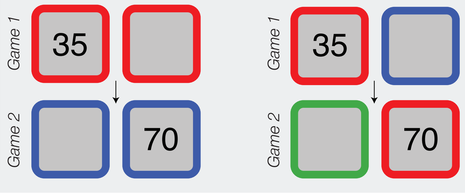

Results

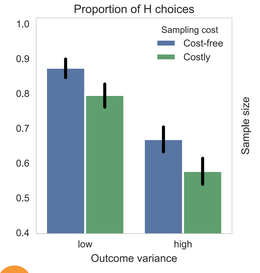

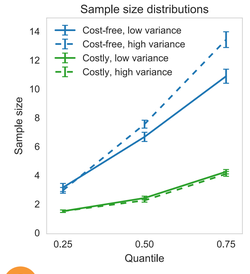

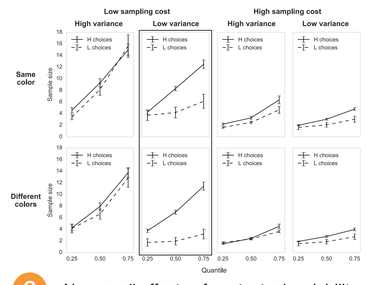

1 - Increased sampling cost resulted in lower *p(H)* [mixed effects logistic regression, p < .001] and smaller sample sizes [mixed effects negative binomial reg., p < .001]

2 - Increased outcome variance resulted in lower *p(H)* [mixed effects logistic regression, p < .001].
- Variance *X* cost interaction on sample size [neg. binomial reg., p < .001], with larger sample sizes in cost-free trials

3 - No overall effects of contextual variability on *p(H)* or sample size.
- Within low-cost, low-variance trials, smaller sample sizes among different-color group for *L* choices [neg. binomial reg., p = .003] but no difference for *H* choices [p > .05].

Model comparison

- CHASE was fit with max likelihood to choice and sample size data for each participant, with separate threhold (θ) parameters for each cost condition [4 parameters: $\theta_{cost\text{-}free}$, θ_{costly}, τ, and p_{stay}].
- Optional stopping version (depicted above) compared with fixed stopping variant that assumes sample sizes follow geometric distribution (with shape parameter *q)* and are independent of preference accumulation [4 parameters: $q_{cost\text{-}free}$, q_{costly}, τ, and p_{stay}]
- 81/100 participants were better described by optional stopping model (according to lower BIC score). Distributions of best-fit parameters show decrease in threshold (θ) from cost-free to costly trials, as well as increase in starting point variability (τ) among different-color group.

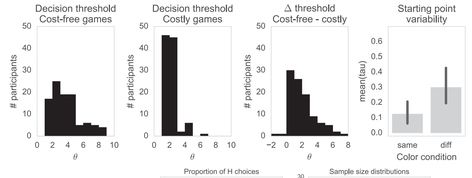

Best-fit model predictions:

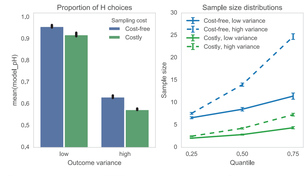

Summary

- Whereas previous research on sequential sampling models has predominantly focused on RT and choice (i.e., over short time-scales and with fixed external stimuli), the present work demonstrates that similar mechanism can account for complex relationships between external, discrete information sampling and final choices (see also Brown et al., 2009).
- Sequential sampling process enables adaptive exploration, such that the amount of information search is adjusted based on the costs of exploration and uncertainty about options' values.
- In contrast to existing models of the sampling paradigm, CHASE provides a framework for understanding variation in exploration and choice under different conditions. Moreover, it generates novel predictions for DFE based on hallmarks of sequential sampling processes in other domains (e.g., effect of starting point variability).

References

Brown, S., Steyvers, M., and Wagenmakers, E. J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, 53(6):453–462.

Diederich, A. and 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.

Hau, R., Pleskac, T. J., Kiefer, J., and 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.

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

Markant, D., Pleskac, T. J., Diederich, A., Pachur, T., and Hertwig, R. (2015). Model- ing choice and search in decisions from experience: A sequential sampling approach. In Dale., R., Jennings, C., Maglio, P., Matlock, T., Noelle, D., Warlaumont, A., and Yoshimi, J., editors, *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, Austin, TX. Cognitive Science Society.

Ratcliff, R. and Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5):347–356.