Modeling Probability Knowledge and Choice in Decisions from Experience

Guy E. Hawkins
(guy.e.hawkins@gmail.com)
University of New South Wales

Adrian R. Camilleri
(adrian.camilleri@duke.edu)
Duke University

Andrew Heathcote
(andrew.heathcote@newcastle.edu.au)
University of Newcastle

Ben R. Newell
(ben.newell@unsw.edu.au)
University of New South Wales

Scott D. Brown
(scott.brown@newcastle.edu.au)
University of Newcastle

Abstract

In most everyday decisions we learn about the outcomes of alternative courses of action through experience: a sampling process. Current models of these decisions from experience do not explain how the sample outcomes are used to form a representation of the distribution of outcomes. We overcome this limitation by developing a new and simple model, the Exemplar Confusion (ExCon) model. In a novel experiment, the model predicted participants' choices and their knowledge of outcome probabilities, when choosing among multiple-outcome gambles in sampling and feedback versions of the task. The model also performed at least as well as other leading choice models when evaluated against benchmark data from the Technion Prediction Tournament. Our approach advances current understanding by proposing a psychological mechanism for how probability estimates arise rather than using estimates solely as inputs to choice models.

Keywords: Experience-based choice; Probability estimation; Sampling; Feedback; Exemplar model

Utility Maximization Models

Most choice models describe behavior as if people multiply some function of the probability of an outcome by that outcome's value, and maximize. The utility maximization framework is fundamental to many successful models of choice including Prospect Theory (Kahneman & Tversky, 1979) and Cumulative Prospect Theory (Tversky & Kahneman, 1992). Evaluation of these models has relied heavily on "decisions from description", where outcomes and their associated probabilities are explicitly stated (Barron & Erev, 2003; Rakow & Newell, 2010). This focus neglects cognitive processes that must underlie many everyday decisions where probabilities and outcomes are not explicitly provided. In such "decisions from experience" (Rakow & Newell, 2010), decision-makers must explore their environment to establish the range of potential outcomes and the probability with which each occurs. Here, we extend utility maximization models to capture the process of acquiring knowledge of potential outcomes and the representation of their probabilities.

Decisions from Experience

In "decisions from experience" the decision-maker learns about potential outcomes and their respective probabilities by sequentially sampling outcomes from the environment. This situation is modeled in the laboratory using two similar paradigms. In the *feedback paradigm* the outcome of each sample is added to a running total that the decision-maker is tasked with maximizing. The decision-maker must

trade-off the objectives of learning more about the options (to "explore") while trying to maximize the payoff from an unknown number of choices (to "exploit"). The *sampling paradigm* separates these goals, allowing the decision-maker to freely sample during an exploration stage and, at a time of their choosing, proceed to an exploitation stage where a single choice is made.

Preferences in the sampling and feedback paradigms tend to be highly correlated (Erev et al., 2010) and suggest that decision-makers *under*weight low-probability outcomes (Camilleri & Newell, 2011b; Ungemach, Chater, & Stewart, 2009). This contrasts with preferences observed in the conventional description paradigm, where decision-makers appear to *over*weight low-probability outcomes (Kahneman & Tversky, 1979). This "description-experience gap" can be reduced or eliminated in the sampling paradigm when samples are equated (Rakow, Demes, & Newell, 2008; Camilleri & Newell, 2011a). This suggests the description and experience formats may share common core features – the need to combine probability information with outcome values to inform choice. Recent models of experience-based choice have not retained utility maximization as a core feature. This may be due to a failure to adequately capture the process by which outcome distributions are represented.

Modeling Decisions from Experience

Several models of experience-based choice were recently tested in the Technion Prediction Tournament (TPT), where the organizers collected large datasets in the description, sampling, and feedback paradigms (Erev et al., 2010). The winner in each paradigm was the model that best predicted average choice in its own paradigm, with no incentive to develop models that generalized across different tasks, such as the sampling and feedback paradigms. The competition structure therefore did not prioritize capturing how sequentially observed outcomes were integrated into a representation nor how this representation could be used to model knowledge of the outcome probabilities. In fact, some models took probability information as *inputs* to the modeling process.

An alternative approach is illustrated by exemplar models such as the ACT-R modeling framework (Anderson & Lebiere, 1998) and the GCM (Nosofsky, 1986). Exemplar models assume that a memory trace is recorded each time a stimulus is encountered, which might include information about the stimulus, the corresponding feedback, and the gen-

eral context in which the stimulus was encountered. Later, the properties of an unfamiliar stimulus can be inferred from the properties of related exemplars stored in memory. Exemplar models are not new in economic decision making – e.g., the *k*-sampler model (Erev et al., 2010) and the Instance Based Learning (IBL) model (Gonzalez & Dutt, 2011) – and also have a much longer history in explaining many aspects of higher-level cognition, such as categorization (e.g., Nosofsky, 1986) and probability judgment (e.g., Juslin & Persson, 2002). We argue that the storage of exemplars in memory provides a simple and psychologically plausible approach to explain decision-makers' choices and the estimation of probability information about gambles in risky choice problems.

Modeling Probability Knowledge

Studies of experience-based choice have previously asked decision-makers to estimate outcome probabilities and observed that estimates are either well calibrated (Fox & Hadar, 2006) or that rare events are overestimated (e.g., Camilleri & Newell, 2009; Ungemach et al., 2009). This reveals a subtle paradox where decision-makers overestimate the probability of rare events yet make choices as if they underweight rare events. A successful model of experience-based choice should account for this overestimation-underweighting paradox, but this is only possible if the model accounts for probability knowledge to begin with.

A New Approach

In developing our modeling approach, we highlight two limitations of the current choice modeling enterprise, exemplified by the TPT. First, many current models are highly specific and only successful in the task for which they were originally designed. Second, existing data and procedures make it difficult to discriminate between models – models that make a variety of assumptions can perform nearly equally well. This is due to the limited variability in the problem sets (a binary choice between a safe option and a two-outcome risky option), limited data sources used to constrain the models (solely choice), and a very general prediction goal (prediction of aggregated choice proportions). To the best of our knowledge, no model has attempted to simultaneously model choices and the outcome distributions associated with the alternative options. Therefore, we constructed a general, exemplar-based model and evaluated it with respect to two streams of data in multiple-outcome gambles: decision-makers' choices and their estimates of outcome probabilities, in the sampling and feedback paradigms.

Experiment

Decision-makers were presented with pairs of five-outcome lotteries in the sampling or feedback paradigms, and asked to choose between the lotteries and estimate the probability of each outcome. This allowed us to collect ten outcome probability estimates per gamble (five for each lottery), which were used to constrain the model in subsequent analyses. We expected similar patterns of choice between the sampling and feedback paradigms (cf. Erev et al., 2010). Participants observed more samples in the feedback than sampling condition, so we expected more accurate probability estimates from those participants.

Method

Participants

107 undergraduate students from two Australian universities participated for payment contingent on performance.

Materials and Design

Participants always made choices between two competing lotteries; we denote such a pair of lotteries as a "problem". We used six lotteries in total, adopted from Lopes and Oden (1999). Each lottery had five possible outcomes ranging from $0 to $348. The outcome distribution was unique for each lottery but all had expected values of approximately $100 (Figure 1). The lotteries were unlabeled during the experiment, but for discussion we adopt Lopes and Oden's lottery names: Riskless, Rectangular, Peaked, Bimodal, Shortshot, or Longshot, depending on the specific distribution of outcomes.
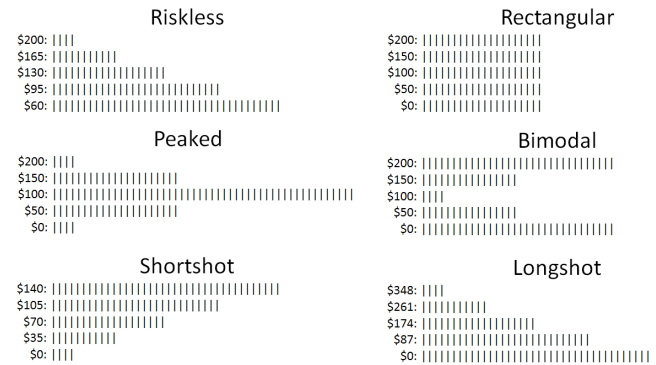

Figure 1: Visual representation of the outcome distribution for the six lotteries, adapted from Lopes and Oden (1999). Participants were not presented with labels or figures.

Participants were allocated to the sampling or feedback task. We measured participants' preferences in each problem, and their estimates of the probabilities of the outcomes from the lotteries in the problem. In the sampling condition, lottery preference was operationalized as the one-shot choice made after the free sampling phase. In the feedback condition, lottery preference was operationalized as the deck selected most frequently in the final 50 samples.[1]

The six lotteries can be ordered in 15 unique pairings (ignoring order, and without identical choices). Participants in the sampling condition played each pairing once, with order randomized across participants. To equate experiment length, participants in the feedback condition played a random sample of six of the fifteen problems. Data were excluded from

---

[1]Similar results were obtained when using the mode across all 100 samples or just the final sample.

problems where participants failed to sample from one of the lotteries. Some participants did not complete all problems during the one hour experiment so there were unequal sample sizes: 40 participants in the sampling condition experienced a total of 528 problems (32–39 data points per problem), and 67 participants in the feedback condition experienced a total of 386 problems (22–32 data points per problem).

Procedure

Participants made a number of choices between computerized decks of cards and were instructed to use their choices to earn as much money as possible. The screen position (left, right) and order of lotteries was randomized. Participants were presented with two unlabeled images of decks of cards, each associated with a lottery from Figure 1. When a deck was selected, an outcome was randomly sampled (with replacement) from the associated lottery and displayed briefly as if the participant had turned over a playing card. We call the act of selecting a deck and observing an outcome a "sample".

The sampling task began with an exploration phase where the participant could learn about the lotteries, so each sample had no consequence. Participants terminated exploration at any time to make a final choice indicating their lottery preference. The outcome of the final choice was added to the participant's running total for the experiment. In the feedback task, each problem granted 100 samples (participants were not told the number of samples). Each of the 100 samples was consequential: the sample outcomes were added to the participant's running score, which was constantly displayed on screen.

After making a choice (sampling) or taking 100 samples (feedback), participants estimated the probability of different types of cards in each deck – the probability of the different outcomes in the lotteries. Six different outcome values were presented beside adjustable sliders with order randomized across problems and participants. The default starting estimate was 0. Participants moved sliders between 0 and 100 to indicate the estimated percentage. One of the six outcomes was a "foil" that was not an outcome from the lottery in question but was selected from one of the other decks. The foil was used to identify participants who did not attend to the sampling process. Participants proceeded to the next problem when the sum of the sliders for each deck equaled 100% and a confirmation button was clicked. At the end of the experiment participants were reimbursed by converting every 100 experiment dollars to AU$1 (contingent on choices, not probability estimates).

Results

Preferences

Figure 2 displays the percentage of participants preferring each lottery, averaged over all problems in which the lottery was presented. Participants generally favored lotteries that minimized the possibility of obtaining zero (riskless, peaked, shortshot), consistent with a negatively accelerated utility function for gains, as assumed in many theories of choice. Preferences in the feedback and sampling paradigms were positively correlated ($r = .64$, $p = .01$), consistent with previous research (Erev et al., 2010).
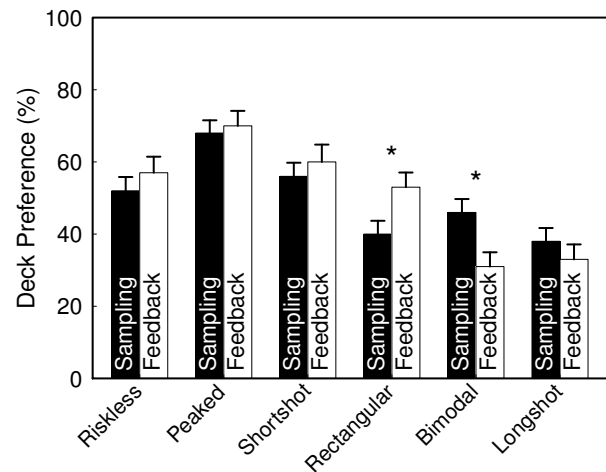

Figure 2: Percentage of participants preferring each lottery, averaged over problems.

Preferences in the sampling and feedback conditions differed in two respects: under feedback there was a stronger preference for the rectangular lottery ($\chi^2(1) = 4.89$, $p = .027$) but a weaker preference for the bimodal lottery ($\chi^2(1) = 7.07$, $p = .008$). Table 1 shows preferences between pairs of lotteries. Cell entries indicate the average preference for the column-named lottery over the row-named lottery. Asterisks denote preferences where a lottery was significantly preferred over indifference (i.e., 50–50) by z-test. For example, in the first row of Table 1 the 79 value is asterisked, indicating that a significant proportion of participants preferred the peaked lottery over the longshot lottery in the sampling condition.

Probability Knowledge

Figure 3 plots median probability estimates assigned to outcomes against the actual sampled frequency of those outcomes, in the samples observed by participants.[2] Participants had good knowledge of the outcome probabilities: in both conditions the median estimate assigned to outcomes increased almost always with increasing sample probability. Nevertheless, there was a tendency to overestimate the probability of rare outcomes and underestimate the probability of frequent outcomes, indicated by the inverted-S shapes in Figure 3. This pattern appeared in the sampling and feedback paradigms, and was unchanged when probability estimates were graphed against the *population* probability of the outcome (i.e., the proportion of times it would appear in the long run, Figure 1) rather than the sampled frequency of the out-

---

[2]Participant estimates of the probability of foil outcomes were accurate. Foil cards were assigned zero probability in 62% of problems. On the remaining problems, this estimate was still small – 11% on average – and on 37% of these trials the foil was assigned the smallest of the estimated probabilities. Since the foils were mostly well identified by participants we do not analyze those data further.

Table 1: Percentage of participants preferring the column-named lottery over the row-named lottery, separately for the sampling and feedback conditions. ExCon model predictions are shown in parentheses. Lottery pairings for which data and model had the same modal preference are shown in bold face. Abbreviations refer to deck type: RL=riskless, PK=peaked, SS=shortshot, RC=rectangular, BM=bimodal, LS=longshot. * $p < .05$ by z-test.

Sampling

	RL	PK	SS	RC	BM
LS	60 (**73**)	79*(**62**)	75* (48)	46 (**39**)	49 (**23**)
BM	55 (**88**)	69 (**77**)	59 (**59**)	39 (**48**)	
RC	59 (**89**)	64 (**73**)	58 (**55**)		
SS	47 (73)	66 (**82**)			
PK	39 (78)				

Feedback

	RL	PK	SS	RC	BM
LS	62 (**84**)	72 (**73**)	88*(**61**)	63 (46)	55 (**52**)
BM	82*(**99**)	70 (**76**)	88*(**79**)	62 (**80**)	
RC	59 (**95**)	79*(**73**)	56 (**56**)		
SS	44 (95)	70 (**63**)			
PK	39 (86)				

Figure 3: Probability estimates from participants (y-axes) against the sampled probability for the corresponding outcome (x-axes) in the sampling (left panel) and feedback (right panel) conditions. Circles show medians, whiskers show 5^{th} and 95^{th} percentiles across participants. Gray identity lines indicate probability estimates that perfectly reflect the sampled probabilities. Black lines show median probability estimates of the ExCon model.

come (i.e., the proportion of times it really did appear, in the samples observed by participants).

We confirmed the statistical reliability of differences between the pattern of estimated versus sampled probabilities and the $y = x$ line (which indicates probability estimates that perfectly reflect the sampled probabilities) using the Wald-Wolfowitz (or "runs") test. This analysis examines the sign of successive residuals under the null hypothesis that residuals should be randomly distributed either side of zero. There was non-random scatter around the $y = x$ line in both paradigms, both p's $< .001$, reflecting the run of positive residuals for low sample probabilities and negative residuals for high sample probabilities.[3]

Exemplar Confusion: An Account of Choices and Probability Knowledge

Exemplar models assume that observers record a memory trace each time they encounter a stimulus, and later use these traces to make inferences about their experiences. The exemplar confusion (ExCon) model employs this idea as an extension of the primed sampler models described by Erev et al. (2010), which we refer to as a k-sampler. The standard k-

---

[3]A mundane explanation for the inverse-S shape of the probability estimates is a mixture of participants, some who were perfectly calibrated ($y = x$ line) and others who were not engaged in the task (probability estimates unrelated to the sampled probabilities – horizontal lines at $y \approx .2$), which could lead to the inverted-S shapes even if no individual participant displayed such a pattern. We ruled out this explanation by removing many participants to leave us with only the very best (those most likely to be engaged in the task), and the inverted-S shape remained (graphs not shown).
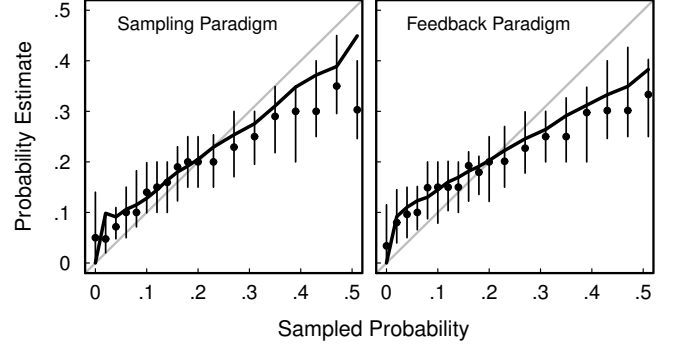
sampler draws k samples from each lottery and prefers the lottery with the greater sample mean. The k-sampler instantiates a limited memory capacity (only k exemplars per lottery are maintained). Despite its simplicity, the k-sampler provides a reasonable account of choice data compared to more complex models (e.g., Erev et al., 2010). However, the k-sampler fails to capture biased probability knowledge because it necessarily predicts accurate estimates, on average.

The ExCon model makes two modifications to the k-sampler. First, ExCon replaces the k-sampler's limit on memory *capacity* (k) with a limit on memory *accuracy*. We instantiate memory imperfection by degrading the information content of exemplars: a sample-by-sample confusion process where new exemplars can interfere with previously stored exemplars. This form of confusion occurs with the passing of events rather than the passing of time alone, which has precedence in the memory literature (e.g., Lewandowsky, Geiger, & Oberauer, 2008).

The confusion process occurs as follows. Each lottery begins with an empty, limitless memory store. When a sample outcome is observed a memory trace of the outcome value is added to the corresponding lottery store, and the confusion process then operates which leads to a small chance of "confusing" the exemplars within that store. Each exemplar in the store has a fixed probability (p) of having its outcome value confused – substituted with the outcome value from another exemplar already in the store. If an exemplar's outcome value is confused, the new outcome value assigned to that exemplar is chosen uniformly from the *list* of all exemplar labels in the store (rather than uniformly from the *distribution* of all exemplars in the store). The p parameter governing the exemplar confusion process is the sole free parameter of the model to be estimated from data.

The ExCon also differs to the *k*-sampler in the adoption of a utility maximization rule: ExCon prefers the lottery whose exemplar store has the highest average utility. We implemented the diminishing marginal utility function for gains specified by Lopes and Oden (1999): $u(x) = x^{0.551}$. This duplication reflects our belief that the choice rule at the core of both description and experience-based choice is the expected utility theory assumption of multiplying some function of probability with an outcome value, and maximizing.

ExCon Implementation

The ExCon model was given the same sequences of outcomes experienced by the participants on a problem-by-problem basis. The sequences of outcomes were shown to the model 100 times (i.e., 100 Monte-Carlo replicates for every real participant). After the sampling process had finished, model preference was inferred differently for the sampling and feedback paradigms – consistent with the human data. In the sampling paradigm, the model preferred the lottery with the highest average utility. In the feedback paradigm, after each sample outcome the ExCon's decision rule determined which lottery would be chosen on the following sample, for all 100 samples in each problem, and the preference of the Monte-Carlo replicate was taken as its modal choice over the last 50 trials.

We simulated model predictions for 20 values of the *p* parameter on $[0, .1]$. The upper end of this interval implies a high degree of confusion: with a .1 chance of confusion at each sample, the probability of accurately maintaining a single memory trace (exemplar) following 10 samples is $(1 - .1)^{10} = .349$, and following 100 samples is only $(1 - .1)^{100} < .001$. The ExCon instantiates that a *single* exemplar is increasingly likely to be confused as new exemplars enter the store. Nonetheless, the *set* of exemplars will approximate the true distribution with increasing accuracy up until a certain sample size, after which the set of exemplars become noisier.

We assessed the goodness-of-fit on the ExCon's ability to predict two outcome measures in data: choice and probability knowledge. Choice prediction accuracy was calculated as the proportion of times the model successfully predicted the choice made by the participant when exposed to the same sequence of samples for each problem, which we refer to as "trial-by-trial agreement". Probability estimate prediction accuracy was calculated as the sum of the squared deviations between participant and model probability estimates across the sampled probability bins. This measures an average distance between the data and the model when represented in a plot such as Figure 3. Probability estimates were derived from the frequency of each outcome in the ExCon memory stores. The best-fitting parameter estimates were those that maximized goodness-of-fit to choices and minimized prediction error in probability estimate data. Parameters were estimated separately for the sampling and feedback tasks.

ExCon Evaluation

Choice Behavior

With $p = .033$ (sampling) and $p = .027$ (feedback), the ExCon model had 65.3% and 70.7% trial-by-trial agreement with preference data from the sampling and feedback tasks, respectively. In isolation it is difficult to determine whether this reflects good performance, so we provide two benchmark comparisons. The first benchmark was the "natural mean" heuristic, which prefers the lottery with the highest observed sample mean (Hertwig & Pleskac, 2008). The ExCon reduces to the natural mean heuristic when $p = 0$ and assuming a linear utility function. The ExCon model outperforms the natural mean heuristic in both tasks (sampling – 62.9%, feedback – 64.5%). The second benchmark maximized choice in light of the inferred expected value of each lottery as indicated by participants' probability estimates, and was also outperformed by the ExCon model (sampling – 62.9%, feedback – 56.7%). The success of the ExCon relative to the two benchmarks highlights the benefit of a choice rule that maximizes utility rather than value.

The average ExCon choice between pairs of lotteries is displayed in Table 1; bold face shows lottery pairings for which ExCon produced the same modal preference as the participants (i.e., both data and model scored above or below 50%). On this lottery-by-lottery basis, the model correctly predicted 24 out of 30 lottery preferences observed in data even though the model was not explicitly fit to this table.

Probability Estimates

The black lines in Figure 3 illustrate the ExCon probability estimate predictions. The model captures the qualitative patterns in the probability estimates – overestimation of rare outcomes and underestimation of common outcomes.

Technion Prediction Tournament

We now demonstrate that the ExCon model performs as well as leading competitor models when given a benchmark set of problems from the TPT (Erev et al., 2010). The TPT evaluated models separately for the sampling and feedback paradigms using a common problem set (for details, see Erev et al., 2010). The TPT used one data set for model calibration, and another data set for the evaluation (from an identical problem distribution as the original problems). We simulated the process of entering the ExCon model in the TPT, following TPT guidelines: we estimated model parameters using the estimation data set, and used those parameter values to evaluate predictive performance in the competition data set. In the TPT, model performance was evaluated using the average proportion of agreement (across problems) between the modal preference of the model and of the participants (e.g., Erev et al., 2010), which we refer to as "PAgree". Our measure of trial-by-trial agreement reported above maintains more information than PAgree, but for ease of comparison we use the common metric of PAgree values to compare our results to those from the TPT (cf. Table 3, Erev et al.).

The ExCon model performed approximately as well as the winners and runners-up in the sampling *and* feedback paradigms, in contrast to the TPT where each model was successful in the sampling *or* feedback paradigm – see Table 2. The best-fitting parameter values tended toward extremes of the parameter space ($p \approx 0$). However, the choice data from the TPT did not constrain the confusion parameter of the Ex-Con; similar PAgree predictions were observed across the examined parameter space of the model. This result was not restricted to the ExCon: we performed the same model fits with the k-sampler and the three-parameter instance based learning (IBL; Gonzalez & Dutt, 2011) models, and found that the parameters of those models were also under-constrained when the models were examined solely against choice data. This under-constraint reinforces the need for multiple streams of data, and richer problem sets, in evaluating modern choice models.

Table 2: PAgree for the ExCon model and TPT winner in the sampling and feedback paradigms, for the estimation and competition data sets.

	Estimation		Competition	
	ExCon	TPT	ExCon	TPT
Sampling	91.4	95	92.7	83
Feedback	83.3	82	83.3	87

Conclusions

We addressed two issues regarding experience-based choice. Firstly, models of experience-based choice should account not only for participants' choices in multiple paradigms and datasets but also the process by which participants construct and represent the probability distributions upon which those choices are based. To that end we designed a new model – the exemplar confusion model (ExCon) – that provides a process explanation of how people might acquire knowledge of outcome distributions. We assumed that sampled outcomes are stored as memory traces – exemplars – and represent the outcome distribution for alternative options. We also assumed an imperfect, noisy memory system implemented as a confusion process. The confusion process operated each time a new exemplar was stored, similar to retroactive memory interference. The ExCon successfully predicted decision-makers' behavior across two experience-based choice paradigms in a novel data set and an important existing dataset (the TPT).

Our second point relates to model evaluation. Examination of solely choice data, as in the TPT, may not sufficiently constrain models, and therefore permits models with different assumptions to perform equally well. Under-constraint is due to limited variability in the problem sets (binary choices between a safe option and a two-outcome risky option), limited data sources to constrain models (only choice), and a very general prediction goal (aggregated choice proportion). In collecting probability estimates we departed from previous models of risky choice (e.g., Prospect Theory) which in-

fer probability weighting from choices or measure probability estimates to serve as model inputs. One of our novel contributions is to instead use this stream of data to constrain model development. With this constraint we developed a model of how discovered information is used to construct a representation, and how that representation is used to form a preference. A challenge for future research is to develop a comprehensive model that also incorporates information search.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Barron, G., & Erev, I. (2003). Small feedback–based decisions and their limited correspondence to description–based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.

Camilleri, A. R., & Newell, B. R. (2009). The role of representation in experience–based choice. *Judgment and Decision Making*, *4*, 518–529.

Camilleri, A. R., & Newell, B. R. (2011a). Description– and experience–based choice: Does equivalent information equal equivalent choice? *Acta Psychologica*, *136*, 276–284.

Camilleri, A. R., & Newell, B. R. (2011b). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*, 377–384.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 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*, 159–161.

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

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 209–236). Oxford, England: Oxford University Press.

Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, *26*, 563–607.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Lewandowsky, S., Geiger, S. M., & Oberauer, K. (2008). Interference–based forgetting in verbal short–term memory. *Journal of Memory and Language*, *59*, 200–222.

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology*, *43*, 286–313.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Rakow, T., Demes, K., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience–based choice. *Organizational Behavior and Human Decision Processes*, *106*, 168–179.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience–based choice. *Journal of Behavioral Decision Making*, *23*, 1–14.

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

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted, when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473–479.