



City Research Online

City, University of London Institutional Repository

Citation: Hotaling, J., Donkin, C., Jarvstad, A. & Newell, B. R. (2020). MEM-EX: An exemplar memory model of decisions from experience (10.31234/osf.io/fjhr9). .

This is the preprint version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/25754/>

Link to published version: <https://doi.org/10.31234/osf.io/fjhr9>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

MEM-EX: An exemplar memory model of decisions from experience

Jared M. Hotaling^{1,2}, Chris Donkin¹, Andreas Jarvstad³, Ben R. Newell¹

¹University of New South Wales

²University of Illinois at Urbana-Champaign

³City, University of London

Please send all correspondence to:

Jared Hotaling

Department of Psychology

University of Illinois at Urbana-Champaign

Champaign, IL 61820

jaredhotaling@gmail.com

Abstract

Many real-world decisions must be made on basis of experienced outcomes. However, there is little consensus about the mechanisms by which people make these decisions from experience (DfE). Across five experiments, we identified several factors influencing DfE. We also introduce a novel computational modeling framework, the *memory for exemplars model* (MEM-EX), which posits that decision makers rely on memory for previously experienced outcomes to make choices. Using MEM-EX, we demonstrate how cognitive mechanisms provide intuitive and parsimonious explanations for the effects of value-ignorance, salience, outcome order, and sample size. We also conduct a cross-validation analysis of several models within the MEM-EX framework, as well as a baseline model built on principles of reinforcement-learning. We find that MEM-EX consistently outperforms this baseline, demonstrating its value as a tool for making quantitative predictions without overfitting. We discuss the implications of these findings on our understanding of the interplay between attention, memory, and experience-based choice.

Keywords: decisions from experience; computational models; decision making; cognitive mechanisms; exemplar memory; reinforcement-learning.

1. Introduction

Many important choices are based on past experience. In these situations, decision makers lack precise descriptions of their options, and must instead draw from their knowledge of previous outcomes to guide their behavior. Memory therefore plays a central role in experience-based decisions. In this article, we investigate the psychological processes by which people use memory for previously observed outcomes to make choices.

1.1. The importance of Understanding Decisions from Experience

People use past experience to guide their choices in many situations. These includes mundane issues, like choosing at which restaurant to eat based on past meals or choosing the route on which to drive home based on past traffic patterns. However, many important choices – in domains related to health, consumption, social interactions, and investment – also rely on previously experienced outcomes. In light of the recent COVID-19 pandemic, governments and public health organizations around the world are interested in how individuals judge the risks associated with various behaviors. Some worry that people will underestimate or ignore advice regarding social distancing and isolation because they have not personally experienced rare, but dire, consequences of the outbreak. While others worry that sensationalized reporting of virus-related deaths may distort decision making regarding the balance between tightening restrictions to reduce infections and loosening them to ease the economic and social costs of prolonged shutdowns to businesses, school, and public services.

In recent years there has been an explosion of interest in decisions from experience (DfE). Much of this has focused on the so-called *description-experience gap* (see Wulff, Mergenthaler-Canseco, & Hertwig, 2018 for a recent review), wherein decisions from description typically result in choices that imply overweighting of rare events (Kahneman & Tversky, 1979; Rieskamp, 2008), whereas DfE do not (Camilleri & Newell, 2011b; Hertwig, Barron, Weber, & Erev, 2004; Lejarraga & Gonzalez, 2011; Rakow &

Newell, 2010; Yechiam & Busemeyer, 2006). However, despite these investigations, the causes of the gap remain elusive (e.g. Glöckner, Hilbig, Henninger, & Fiedler, 2016; Wulff et al., 2018).

In this article, we take a different approach. Rather than compare DfE to decisions from description, we use a combination of experimental manipulations and cognitive modeling to probe the psychological processes underlying experience-based choices. For example, through our *value-ignorance* manipulation we vary the juncture at which values and probabilities are integrated in order to investigate the cognitive mechanisms by which people record and update experiences in memory. In so doing, we aim to move beyond simple comparisons of description and experience to further elucidate the decision-making process.

1.2. Types of Decisions from Experience

Experience-based decisions come in several forms, which can be classified into two broad categories. In the *repeated-choice* paradigm, individuals make a series of choices between two or more uncertain alternatives. Each choice is consequential, with the decision maker receiving the resultant outcome. Under this paradigm individuals learn as they choose, and must balance the competing interests of gathering information about outcome distributions (exploration) and maximizing payoffs based on their current knowledge (exploitation).

In the present article, we focus on the alternative *one-shot choice* or *sampling* paradigm. Here, an individual freely samples outcomes from each alternative in an effort to learn the underlying distributions. After completing sampling, the individual makes one consequential choice and receives the resultant outcome. The advantage of using this approach to study experience-based decision making is that learning and choosing are separated, because individuals do not face a tradeoff between exploration and exploitation. In the experiments described below, we impose the additional constraint that individuals observe a representative sample of outcomes from each alternative before making a choice. That is, participants sample a predetermined set of outcomes that perfectly match the

underlying outcome distribution of each alternative. After completing sampling, participants make a single consequential choice. By virtue of using a representative sample, we also eliminate sampling error as a source of individual differences (Camilleri & Newell, 2011a; Rakow, Domes, & Newell, 2008).

1.3. Models of Decisions from Experience

Earlier work has used computational modeling to shed light on the psychology of DfE. A prime example comes from the Technion Prediction Tournament (Erev et al., 2010), where teams of researchers were challenged to submit models that were evaluated with regard to their ability to predict choices. Two datasets were created, with the estimation set being used to fit models and the competition set being used to evaluate them. Behavioral results were aggregated across individuals and compared to each model's predicted choice proportions. Model performance was measured as the mean squared distance between predicted and observed choice proportions across the twelve choice problems in the competition set. For the competition involving one-shot DfE – the condition most relevant to this article – the winning model was an ensemble of four equiprobable choice rules: two variants of the *natural mean heuristic* (see Hertwig & Pleskac, 2008), a version of the *priority heuristic* (Brandstätter, Gigerenzer, & Hertwig, 2006), and a variant of *cumulative prospect theory* (CPT; Tversky & Kahneman, 1992). This result shows the importance of assuming multiple decision strategies, though the model does not specify to what degree this variability occurred within subjects versus between subjects. Although we view the Technion Prediction Tournament as making a valuable contribution to the literature on DfE, we pursue a different goal in our research. Rather than focus on predictive accuracy alone, we aim to develop a deeper understanding of the psychological processes underlying DfE, and therefore use behavioral data to test and compare how well cognitive mechanisms explain choice patterns.

The *instance-based learning model* (IBL; Lejarraaga, Dutt, & Gonzalez, 2012) has been used to account for DfE in both one-shot and repeated-choice paradigms. At its core, the model assumes that

individuals store in memory an *instance* representing each unique outcome type and the choice alternative that produced it. For one-shot decisions, after sampling is completed the decision maker chooses the alternative with the highest *blended value*. This value is the sum of all observed outcomes for a given alternative, weighted by their probability of being retrieved from memory. Retrieval is governed by recency and frequency, with more recent and more frequent outcomes having greater memory activation. In this article, we develop a new model that shares IBL's foundational assumption that individuals draw from memory of past outcomes to make decisions. We refrain from directly comparing it to IBL because our primary aim is to test the explanatory value of cognitive mechanisms within a single modeling framework, rather than to compare theoretically similar frameworks with different auxiliary assumptions (i.e. noise mechanisms). To preview our results, we find evidence supporting IBL's core assumption that event memory drives DfE.

Reinforcement learning (RL; Sutton & Barto, 1998) is also a popular modeling framework for DfE. Two key virtues of these models lie in their relative simplicity and generality. Their success in modeling experience-based choice in various contexts and domains makes RL models a useful baseline for comparing new models. Below we develop the *Value-Updating model*, which uses a simple RL process in which the decision maker is assumed to update the subjective value of each choice alternative after observing each outcome it produces (see also Hertwig, Barron, Weber, & Erev, 2006). By showing that an alternative model outperforms this baseline, we demonstrate the utility of our approach in both explaining and predicting behavior.

1.4. The Memory for Exemplars Model

We propose a novel computational framework inspired by models of exemplar memory (see also Hawkins, Camilleri, Heathcote, Newell, & Brown, 2014; Lin, Donkin, & Newell, 2015). The *memory for exemplars model* (MEM-EX) posits that individuals use memory for previously experienced outcomes to guide their choices. Although the model can be modified for other paradigms, we will focus on the

version for binary one-shot decisions. We also begin with an overview of the simplest version of MEM-EX, followed by descriptions of four additional cognitive mechanisms designed to explain the influence of various factors on behavior.

According to MEM-EX, individuals represent outcomes using two stores: one for each alternative. The model posits that after taking each new sample (e.g. a reward of 10 points), the observed outcome is recorded in the appropriate store, with each new sample producing a new exemplar in memory. When sampling is finished, each memory store will have been populated with a record of the observed outcomes for that alternative. The model now calculates a subjective value, V , for each alternative by taking the average of these outcomes. For example, if sampling from Option A produced the sequence of outcomes [0, 10, 0, 10, 0, 0, 0, 0, 0, 0] and sampling from Option B produced the sequence [0, 0, 0, 6, 0, 6, 0, 0, 0, 6], $V_A = 2.0$ and $V_B = 1.8$.

Rather than always choosing the alternative with the higher subjective value, MEM-EX uses a *risk bias* mechanism to represent that individuals have a default preference and a threshold of evidence – expressed in terms of a ratio of values – required to overcome that default. The riskier option – here A, which offers a lower probability higher value reward – is chosen if $\frac{V_A}{V_B} > 1 - \beta$, and the safer Option B is chosen if $\frac{V_A}{V_B} < 1 - \beta$.¹ β is a free parameter between -3 and 0.75 representing the amount of additional evidence required to overcome an individual's default preference. A value of -1 indicates a bias toward the safer option, with A being chosen only if its value is twice that of B. A value of .5 corresponds to an equally strong bias in favor of the riskier alternative, while a value of 0 indicates no bias.

1.4.1. Value-Assignment Error

To this basic framework MEM-EX adds four cognitive mechanisms that influence how information is processed. According to the model, when an outcome is sampled, the event is recorded in

¹ A random choice is made if $\frac{V_A}{V_B} = 1 - \beta$.

the appropriate memory store as an exemplar and is assigned the observed value. However, the *value-assignment error* mechanism allows for the possibility that people make errors in assigning values to exemplars. Here, we assume that these assignment errors occur within-alternative, such that participants may confuse outcomes sampled from one option with each other, but not with any outcomes sampled from the other option². For instance, if Option A produces outcomes of either 10 or 0, MEM-EX posits that when an outcome of 10 is observed individuals will sometimes mistakenly record this as a 0. The free parameter λ – with values between 0 and 0.4 – determines the probability of making a value-assignment error for each item in memory. These errors are assumed to occur independently – i.e. each exemplar has a probability λ of being assigned an incorrect value – at the same rate for both outcomes of both options.

To preview a result that we find in each of our experiments, when value information was withheld during sampling and participants were forced to delay value assignments until the moment of choice (see Section 2.1.4), value-assignment errors were more likely. These additional assignment errors increase the frequency of rare events in memory because most errors will occur after sampling the common \$0 outcome. Because rewards are less frequent and have a greater value for the riskier option, value-assignment errors increase its subjective value more than for the safer option, which in turn increases the likelihood that the riskier option is chosen. With this mechanism MEM-EX provides a new and deeper explanation – in terms of psychological processes – for behavior that might otherwise be described with the opaque concept of ‘overweighting’ of rare events.

1.4.2. Memory Sampling Error

MEM-EX also posits that individuals do not use all available information to make their choice, but rather estimate the value of each option based on an imperfect sample of items in memory. This is

² Though theoretically possible, we did not encounter any observations that required between-alternative confusion.

represented with the *memory sampling error* mechanism, according to which the individual randomly draws δ items from each memory store (with replacement) to compute subjective values for each alternative. For instance, imagine that $\delta = 6$ and the sequences [0, 10, 0, 0, 0, 0] and [0, 6, 0, 6, 0, 6] are sampled from memory for Options A and B, respectively. The resulting subjective values ($V_A = 1.67$ and $V_B = 3.0$) differ from those calculated earlier, to the benefit of Option B, which now appears more attractive due to memory sampling error.

1.4.3. Memory Priming

MEM-EX's *memory priming* mechanism states that when a decision maker samples outcomes from memory, they may do so in a biased manner, preferentially sampling salient outcomes (Erev, Glozman, & Hertwig, 2008; Lieder, Griffiths, & Hsu, 2018). Formally, each outcome is assigned a weight, w , and the probability of sampling outcome i is $w_i / \sum_j w_j$. Here, we assume a simple two-state salience framework meant to represent the impact of a single type of perceptual highlighting that is either on or off when an outcome is observed (See Experiment 1 for details). If an outcome was not highlighted, its weight was 1. If it was highlighted, its weight was $1 + \zeta$, with ζ being a free parameter having a value between 0 and 1. Thus, perceptually highlighting an outcome increases its likelihood of being retrieved from memory. Here again, MEM-EX provides a process account of behavior that implies overweighting of salient events.

1.4.4. Memory Confusion

The final mechanism introduced in this article is designed to explain the effect of outcome order on DfE. The *memory confusion* mechanism represents a process by which decision makers mistake past outcomes for new ones. Specifically, when an outcome is sampled, there is a probability, ϕ , that each previously remembered outcome in memory is replaced by the currently sampled outcome. For example, imagine that the memory store for Option A is currently [0, 10, 0, 10, 0, 0] when a new sample is drawn with value 10. According to the memory confusion mechanism, each exemplar in memory has a

probability ϕ of having its value changed to 10 (i.e. the value of the new sample). Like with value-assignment errors, these confusions are assumed to occur within each box, with no confusions between boxes (see also Hawkins et al., 2014; Lin et al., 2015).

This mechanism naturally produces retroactive interference because outcomes that appear at the beginning of a sequence of observations are more likely to be replaced in memory later. For example, an observation of 10 on the first sample might have eight subsequent opportunities to be confused with a later 0 observation, i.e. one for each 0 appearing later in the sequence. However, the penultimate observation will only have one such opportunity, and is therefore more likely to survive until choice³.

It is worth noting that MEM-EX is a stochastic model, with variability emerging naturally from its memory sampling error, value-assignment error, and memory confusion mechanisms. Consequently, the model provides a process account of choice variability that might otherwise be represented with algebraic noise (e.g. logistic choice sensitivity) or a separate heuristic process (e.g. trembling hand error). In this sense, MEM-EX constitutes a strong theory of choice variability because it must explain variability using the same mechanisms that explain people's overall behavior. This is preferable to weaker theories that assume noise components that operate independently from valuation. With such models it can be difficult to disentangle the relative contributions of core and auxiliary (noise) components to performance.

1.5. Goals

In this article, we investigate the psychological processes involved in experience-based decisions. We pursue this goal via two complementary approaches. First, in a series of five experiments we examine the impact of several experiential factors on decision making behavior. From these we gain

³ Note that, although unlikely, multiple confusions can occur sequentially to a single exemplar, e.g. with a 10 being switched to a 0, then later back to a 10.

valuable insights into the roles that value information, outcome salience, outcome order, and sample size play in choice. Second, we develop a new computational framework that we use to study the cognitive mechanisms underlying behavior. Our goal here is to provide a unified framework within which we can test hypotheses regarding cognitive processes, while holding constant auxiliary assumptions (e.g. response-error functions).

We structure this article as follows. To begin, we present behavioral results from several laboratory experiments. After each of these, we use MEM-EX to present an account of the cognitive mechanisms that explain the key behavioral patterns we observe. Finally, we use data from all experiments to compare MEM-EX to an alternative model that posits that people do not store outcomes as exemplars but rather track a single summary value for each option, updating these value after each new observation. Our aim here is to test how well each theoretical framework – one using analogical representations vs. one using summary representations – accounts for behavior at an individual level. To preview our results, we find that MEM-EX provides the best account for the vast majority of individuals in each experiment.

2. Experiments 1 & 2

All of the behavioral data reported in this article comes from an experimental paradigm where participants made a series of choices between two risky alternatives framed as boxes containing colored balls to denote outcomes (see Figure 1). The participants' goal was to sample outcomes from each alternative to learn which of two boxes they would prefer to draw from 'for real' (i.e., for a potential payment) at the choice stage. Data from Experiments 1 and 2 were first reported in Hotelling, Jarvstad, Donkin, and Newell (2019)⁴. In this study, we investigated the impact of rare events on DfE using two manipulations.

⁴ The present Experiments 1 and 2 were labeled as Experiments 3 and 4, respectively in Hotelling et al, 2019.

With the first, we examine how people record and update experienced outcomes in memory by varying the juncture at which outcome values were presented. In the Standard condition, when a sample was drawn its magnitude was displayed on screen, allowing participants to learn about both the values and probabilities of outcomes. In contrast, for the Value-Ignorance condition magnitude information was absent during sampling, with participants only able to learn about the likelihood of receiving a reward (i.e. a blue ball) or not (red balls were always worth \$0). Here, outcome values were revealed after sampling, at the time of choice, necessitating further processing of previously stored outcome information in memory. Hotelling et al. (2019) analyzed these results using cumulative prospect theory (CPT; Tversky & Kahneman, 1992) to measure risk preference, and found that Value-Ignorance led to a greater weighting of rare events in choice. In the present study, we seek to further unpack this result to uncover the cognitive mechanisms that produce this preference. Specifically, MEM-EX predicts that the Value-Ignorance condition will require additional mental operations and will therefore produce more value-assignment errors, which will increase the frequency of rare rewards in memory. This will tend to increase the subjective value of the riskier alternative more, leading to greater risk taking.

The second factor manipulated in Experiments 1 and 2 was outcome salience. Here we test whether perceptually highlighting a rare event during sampling increases its impact on choice. Hotelling et al. (2019) used CPT to determine that the salience manipulation led to greater weighting of rare events. Here we extend this work to investigate the mechanisms by which emphasizing a rare event influenced people's choices. According to MEM-EX's memory priming mechanism, highlighted outcomes will be more prominent or available when sampling items from memory to determine an option's value. With the salience manipulation applied to the rare event of riskier alternative, greater sampling of rare rewards will increase risk taking.

2.1. Method

2.1.1. Ethics

Ethical approval for all experiments was obtained through the institutional review boards of the School of Psychology at the University of New South Wales (UNSW).

2.1.2. Participants

All participants were UNSW students and received course credit plus a monetary bonus (\$0 - \$20) based on a randomly selected trial. 149 (99 female, age 18-53, $M = 22.93$, $SD = 4.63$) individuals participated in Experiment 1, and 177 (106 female, age 18-58, $M = 20.49$, $SD = 3.92$) participated in Experiment 2. Experiment 2 was a preregistered replication of Experiment 1 (details can be found at <https://osf.io/bw7ps>).

2.1.3. Procedure

After giving informed consent, participants were placed in a computer booth where they read the following instructions:

“In this task you will draw balls from pairs of virtual boxes. In each box, there are 100 balls, some of which are blue and some of which are red. Blue balls are associated with reward and red balls are not (reward for a red ball = \$0).”

Participants began by completing a practice trial to familiarize them with the task (Figure 1). They were instructed that each trial in the experiment involved a new pair of boxes and that they would have to learn anew the values and proportions of balls within each box. To emphasize that boxes were different across trials, each box was given a unique color. On each trial, participants were required to sample the entire sample set for both alternatives before making their choice. Participants were able to sample freely (e.g., alternating between boxes, sampling exhaustively from one then the other, etc.), with sampling disabled once the entire set from each box had been seen. To provide a financial incentive,

participants were instructed that one of their choices would be used to draw a ball for a bonus payment at the end of the experiment.




Figure 1. Robot-sampling task A) Main design. Each trial was composed of two phases: a sampling phase (top row), and a choice phase (bottom row). For a given sample, participants could observe either a blue ball, or a red ball. Red balls were worth \$0 and blue balls were worth some reward. In the Standard condition, the value of the blue ball was revealed during sampling (left column). In the Value-Ignorance condition, the value of the blue ball was not revealed until the choice stage. Thus, during sampling under value-ignorance, the probability of drawing a blue ball could be learned, but not its value. B) Example of a sampling sequence. Once a box was selected for sampling (having been clicked), an animation showed the box shaking (to 'mix' the balls), then a 'robot arm' reached down and grabbed a ball, lifted it up to reveal it and then dropped it back down again (illustrating sampling with replacement). Participants were required to sample each box a set number of times but were free to sample in any order. Reprinted from "How to Change the Weight of Rare Events in Decisions from Experience", by J. M. Hotaling, A. Jarvstad, C. Donkin, and B. R. Newell, 2019, *Psychological Science*, 30(12), p. 1768. Copyrighted 2019 by Sage.

Importantly, the samples that participants observed matched the true underlying probabilities of each outcome, thus mitigating other factors that may give rise to illusory 'gaps' (e.g., biased sampling and reliance on small samples; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hertwig & Pleskac, 2010; Rakow et al., 2008).

2.1.4. Materials and Design

The values of boxes (monetary gambles) were determined as follows. Each choice alternative was defined by a reward value, v (range \$1 to \$20), and a probability of reward, π (range .083 to .917). With these values we created a sample set for each alternative representing the proportion of red and blue balls. The size of the sample set ranged from 10 to 12 and the frequency of rewards was determined by π .

Red balls were always worth \$0. The value of blue balls was fixed within each box, but varied across boxes and trials. For example, the value of a blue ball may be \$16 in the left box and \$2 in the right box. In the Standard condition (Figure 1A), each sampled ball was labeled with the outcome value. In the Value-Ignorance condition, sampled balls were not labeled with values, though the instruction indicated that red balls were worth \$0 and the blue balls were worth some reward. Participants could therefore learn the relative proportions of balls in each box, but not their values, with values revealed in the choice phase (Value-Ignorance Figure 1A).

Choice pairs were constructed with the goal of exposing participants to a range of problems. For example, problems could involve zero, one, or two risky options (i.e. $\pi < .5$), and equal or unequal expected values. To better understand the task, consider an example trial involving a riskier option on the left offering a 10% chance of winning 16 points, and a safer option on the right offering an 80% chance of winning 2 points. While sampling from the riskier box participants would observe one blue ball and nine red balls. From the safer box they would sample eight blue balls and two red balls. Each participant was assigned to one of four conditions in a 2 (Standard vs. Value-Ignorance) x 2 (No-Salience vs. Salience) factorial design and received the same twenty decision problems in a random order. See the Supplementary Materials for the specific gambles used in all experiments.

2.1.4.1. Salience Manipulation. The No-Salience condition proceeded as described above, while the Salience condition introduced an additional manipulation whereby, during sampling, some balls

were perceptually highlighted. When a highlighted ball was drawn, an auditory tone played and the ball flashed on screen for approximately 700ms before returning to the box as usual.

The highlighting occurred whenever participants sampled the rare event for the riskier alternative (defined as the alternative with the lower π). This resulted in two types of problems. For fourteen Type 1 (*best-outcome salient*) problems, salience highlighted a rare reward, and was expected to increase the likelihood of choosing the risky option. For six Type 2 (*worst-outcome salient*) problems, salience highlighted an outcome of \$0, and could be expected to decrease the likelihood of choosing the risky option.

2.2. Results

2.2.1. Behavioral Analysis

Since Experiments 1 and 2 used identical methods, we analyze them together. On each trial, we define the risky alternative as the one giving the lower probability of reward, and the safe alternative as the one giving the higher probability of reward.

Figure 2 displays the individual and group mean proportions of choices in favor of the risky alternative across conditions. It shows that the Value-Ignorance condition produced a higher proportion of risky choices than the Standard condition ($M_{\text{standard}} = .40$, $M_{\text{val-ign}} = .46$, $SD_{\text{standard}} = .19$, $SD_{\text{val-ign}} = .26$, $g = .27$)⁵⁶. That is, on average, participants who did not know the values associated with each outcome during sampling chose the riskier option more often than participants for whom value information was present during sampling.

⁵ We report Hedge's g as a measure of effect size (Hedges, 1981).

⁶ The reader may worry about the "reliability" of our description of the empirical data. To belay such concerns, we note that in all experiments we report (and have run) we observe the same pattern of increased risky choices in the value-ignorance condition. Also note that, in addition to the statistical analyses accompanying some claims we discuss here (and also analyzed in Hotaling et al. (2019)), our model-based cross-validation analyses are consistent with our statements about the empirical effects.




Figure 2. Behavior and MEM-EX predictions across value conditions in Experiments 1 and 2. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

The salience manipulation carried the risk of introducing a demand characteristic whereby participants would be encouraged to choose the highlighted option, regardless of which outcome was emphasized. Type 2 problems therefore served as a manipulation check because salience highlighted an unattractive outcome of \$0, rather than a rare reward. Since our primary interest, and the majority of the data, involved Type 1 problems, we focus our analyses on these. In the Supplementary Materials of Hotelling et al. (2019), we show that the salience manipulation had no effect on choices for Type 2 problems, therefore ruling out this potential confound.

The effect of salience on choices in Type 1 trials can be seen in Figure 3. In the Salience condition, participants made riskier choices than in the No-Salience condition ($M_{\text{salience}} = .43$, $M_{\text{no-salience}} = .38$, $SD_{\text{salience}} = .28$, $SD_{\text{no-salience}} = .26$, $g = .19$). That is, perceptually highlighting rare events during sampling increased the likelihood that participants would choose the risky option, particularly in the Standard condition.

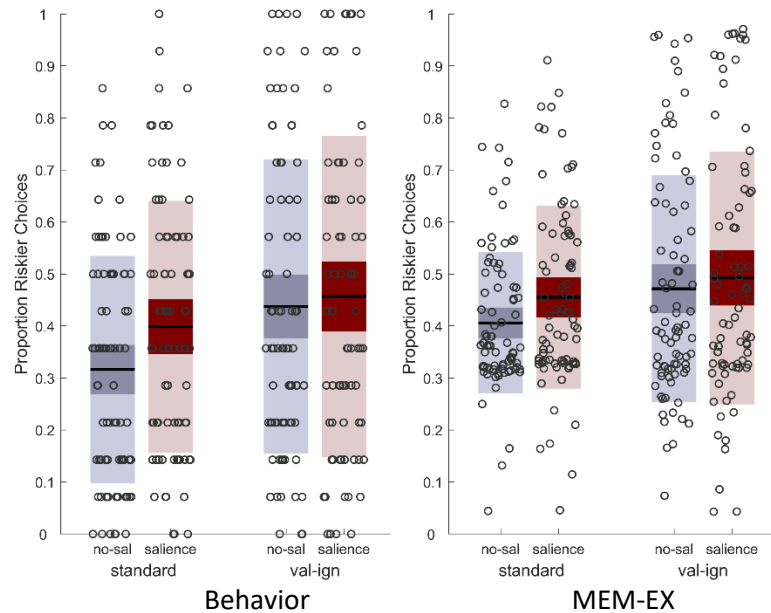


Figure 3. Behavior and MEM-EX predictions across value and salience conditions in Experiments 1 and 2. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

2.2.2. Modeling Analysis

To better understand how our manipulations influenced choices, we now present an account of results in Experiments 1 and 2 using MEM-EX. In our Model Comparison section, we provide details of the modeling procedure, along with results from a comparison of several different models. Here, we focus on the best-performing version of the model, with the goal of demonstrating how MEM-EX uses cognitive mechanisms to explain key behavioral effects⁷.

The *value-assignment error* mechanism provides an intuitive account of the value-ignorance effect. According to the model, after observing an outcome an exemplar is placed in memory and a value is assigned to that exemplar based on the observed outcome magnitude. Crucially, in the Value-

⁷ Qualitative modeling results for Experiments 1-5 are based on fitting each model at the individual level. The procedure for estimating optimal parameter values was otherwise identical to that described in the Model Comparison section for the cross-validation analysis.

Ignorance condition, participants could not assign values online during sampling, but were instead required to wait until sampling was completed before values were revealed.

MEM-EX explains that people made riskier choices under value-ignorance because value-assignment errors were more frequent. We represent this difference by estimating separate λ parameters for each condition. Such an explanation makes intuitive sense, as participants in the Standard condition were allowed to immediately record value information, perhaps while the information was still perceptually available. Here we would expect relatively few errors, which fits with MEM-EX's account indicating that participants assigned the wrong value for only 6.76% of samples ($\bar{\lambda} = .068$, $SD = .094$). In contrast, participants in the Value-Ignorance condition were required to store only the event that a particular colored ball was observed, with no information about the value attached to the ball. When values were revealed on the choice screen, they would then need to reactivate these memories and assign values to each outcome. According to MEM-EX, this additional processing introduces more errors, since it must rely on the same mechanisms that gave rise to the original errors, with participants assigning the wrong value for 11.23% of samples ($\bar{\lambda} = .112$, $SD = .132$). Since additional value-assignment errors lead to the rare events being recalled more frequently the model reproduces the observed behavioral effect (Figure 2).

MEM-EX also provides an explanation of the salience effect – whereby perceptually highlighting rare rewards for the riskier alternative during sampling led to riskier choices – via *memory-sampling error* and *memory priming*. The former represents the notion that individuals do not use all available information to make their choice, but rather estimate the frequency of rare events by sampling items from memory. Model parameters indicated that participants sampled an average of four or five outcomes from memory for each alternative ($\bar{\delta} = 4.371$, $SD = 2.760$). According to MEM-EX's memory-

priming mechanism, in the Saliency condition this sampling was done in a biased manner⁸. On average, participants were approximately 25% more likely to sample salient outcomes from memory ($\bar{\zeta} = .252$, $SD = .294$).

Figure 3 shows that memory priming allows MEM-EX to capture the observed saliency effect, with higher mean risky choice proportions in the Saliency condition for both the Standard and Value-Ignorance conditions. Thus, the model explains the saliency effect by positing that perceptually highlighting rare outcomes during sampling led to these outcomes being more salient in the mental representation of a gamble, which led to an increase in the likelihood of sampling these outcomes from memory and using them to make a decision. Since Type 1 trials highlighted rare rewards for the riskier option, this resulted in more risky choices.

2.3. Discussion

In Experiments 1 and 2 we saw that the presence of value information during sampling influenced people's choices. We analyzed these data within MEM-EX and found that value-assignment errors provided an intuitive and parsimonious explanation of this effect. According to the model, participants in the Value-Ignorance condition were forced to perform additional mental operations to assign values to previously sampled events. This resulted in higher value-assignment errors rates under value-ignorance, which served to increase the proportion of rewards in the mental representation. Although this mechanism applied to both gambles, the riskier gamble had the higher reward, and its subjective value was therefore more greatly affected by each error.

We also found that our saliency manipulation led participants to make riskier choices, which MEM-EX explained using its memory priming mechanism. That is, when outcomes were perceptually highlighted during sampling, they became more salient in the mental representation, and were

⁸ MEM-EX assumed unbiased sampling in the No-Saliency condition because no outcomes were highlighted.

therefore more likely to be sampled from memory when making a choice. Since we highlighted the rare rewards for the riskier alternative, this increased the subjective value of the riskier alternative and increased its choice share. This account bears some similarity to the concept of *availability* (Tversky & Kahneman, 1973) because salient outcomes were more available for sampling in memory. Under this view, memory sampling can be seen as a process of using past events to imagine the likely outcome of choosing each alternative. Salience caused some rare events to be more available, and therefore seem subjectively more likely.

It is worth comparing this model analysis with the one reported in Hotaling et al. (2019), as the two complement each other and provide converging evidence. Hotaling et al. used CPT as a measurement model to understand behavior in term of latent preference functions. That analysis showed that the value-ignorance and salience effects resulted from over-weighting of rare events in the Value-Ignorance and Salience condition, respectively. With MEM-EX, we further unpack these results by developing cognitive mechanisms to provide process explanations for CPT's preference functions. Our central insight is that the overweighting that characterized the results of Hotaling et al.'s analysis arose from a systemic over-representation of rare events in memory.

3. Experiment 3

The between-subjects design of Experiments 1 and 2 limited our ability to draw conclusions using MEM-EX because risk bias parameters were estimated separately for each participant, and might therefore contribute to the model's predicted effects. To remedy this, in Experiment 3 we manipulated the presence of value information within-subjects. We can now assume that each individual has an overarching risk bias that is constant across Standard and Value-Ignorance conditions, allowing us to focus on the value-assignment error mechanism as the only explanation of the value-ignorance effect. Again, MEM-EX predicts higher value-assignment error rates under value-ignorance, leading to riskier choices.

3.1. Method

3.1.1. Participants

42 (19 female, age 17-28, $M = 19.17$, $SD = 1.99$) UNSW students participated and received course credit plus a monetary bonus (\$0 - \$20) based on a randomly selected trial.

3.1.2. Procedure, Material, and Design

Experiment 3 was identical to Experiments 1 and 2, with two exceptions. First, we removed the Salience condition from the design. Second, we manipulated value-ignorance within-subjects. Each participant completed 20 trials from the Standard condition, followed by twenty trials from the Value-Ignorance condition.

3.2. Results

3.2.1. Behavioral Analysis

Figure 4 displays the individual and group mean proportions of choices in favor of the risky alternative across conditions. We find a within-subjects effect of value-ignorance mirroring that seen in Experiments 1 and 2, with participants making riskier choices in the Value-Ignorance condition ($M_{standard} = .40$, $M_{val-ign} = .45$, $SD_{standard} = .18$, $SD_{val-ign} = .24$, $g = .25$).




Figure 4. Behavior and MEM-EX predictions across in Experiment 3. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

3.2.2. Modeling Analysis

Our modeling analysis revealed substantial individual differences with respect to the value-ignorance effect (see Model Comparison for details). The majority of participants (64%) were best fit by a version of MEM-EX (MEM-EX_{null}, see Table 1) that assumed no difference between conditions. That is, this model used the same value-assignment error rate in the Standard and Value-Ignorance conditions. For the remaining individuals, a version of MEM-EX with separate λ parameters (MEM-EX_{base}, see Table 1) again gave the best overall account.

Focusing on the model that assumes separate value-assignment rates, Figure 4 shows that the model produces the same pattern seen in the behavioral data: riskier choices under value ignorance⁹. Importantly, because the presentation of value information was manipulated within subjects, the only parameter that varied across conditions was λ . This means that the value-assignment error mechanism was solely responsible for producing differences across conditions. According to this analysis, participants assigned the wrong value for 8.29% of samples in the Standard condition ($\bar{\lambda} = .083$, SD = .079), but made assignment errors for 11.90% of samples in the Value-Ignorance condition ($\bar{\lambda} = .119$, SD = .120). Figure 5 plots the relationship between individual value-ignorance effect sizes (Value-Ignorance – Standard) and the difference in value-assignment error rates across conditions ($\lambda_{val-ign} - \lambda_{standard}$). Here we can see that this mechanism allows us to understand individual differences in terms of differences in cognitive mechanisms, with parameter differences closely tracking behavioral differences.

⁹ These results are based on fitting MEM-EX_{base} to all participants, regardless of which model version produced the best results in cross-validation.




Figure 5. The relationship between value-ignorance effect size – defined as the difference in risky choice proportion – and the difference in value-assignment error rate parameters – defined as $\lambda_{val-ign} - \lambda_{standard}$ – in Experiment 3. Each dot represents an individual.

3.3. Discussion

In Experiment 3 we replicated the value-ignorance effect, this time at the within-subjects level. This change in design posed a new challenge to MEM-EX because the model was now tasked with explaining the observed effect without appealing to individual differences in risk bias. Its success under these conditions lends more support to the idea that in the Value-Ignorance condition, participants performed additional mental operations at the time of choice, resulting in more value-assignment errors. MEM-EX also explained individual differences in the value-ignorance effect as the result of differences in value-assignment error rates. That is, though the model-selection exercise reveals a number of individuals not clearly affected by the value-ignorance manipulation, a model assuming no effect for all participants would fail to explain those participants who showed a larger effect of the manipulation. In Experiment 4, we build on this insight and use MEM-EX to understand the effect of outcome order on DfE.

4. Experiment 4

The order in which people experience outcomes has been shown to have a significant effect on choice (Rakow et al., 2008; Wulff et al., 2018). Although some have explained these effects in terms of a recency (or primacy) bias, these accounts typically fall short of articulating a specific cognitive process. In Experiment 4 we test the effects of outcome order on DfE, and use MEM-EX to elucidate the mechanism through which they manifest. The model predicts that memory confusions will produce retroactive interference, with early observations partially replaced by later ones. As a result, options whose rewards appeared at the end of the sample sequence will appear to have a greater value and will be increasingly chosen. Methods and hypotheses for Experiment 4 were preregistered (details can be found at <https://osf.io/a264x>).

4.1. Method

4.1.1. Participants

104 (42 female, age 17-35, $M = 19.41$, $SD = 2.55$) UNSW students participated in Experiment 4. Each received course credit plus a monetary bonus (\$0 - \$20) based on a randomly selected trial.

4.1.2. Procedure, Material, and Design

Experiment 4 used a 2 (Standard vs. Value-Ignorance) x 2 (Primacy vs. Recency) between-subjects factorial design. The presentation of value information was manipulated between subjects, as in Experiments 1 and 2. Additionally, half of participants were randomly allocated to each order condition. In the Primacy condition, outcomes from both alternatives were ordered such that all of the rewards (blue balls) appeared at the beginning of the sequence of sampled outcomes. In the Recency condition, all of the non-rewards (red balls) appeared at the beginning of the sequence. Participants completed the same 20 gambles from Experiments 1-3.

4.2. Results

4.2.1. Behavioral Analysis

Figure 6 indicates two interesting findings. First, participants made riskier choices in the Recency condition compared to the Primacy condition ($M_{recency} = .51$, $M_{primacy} = .37$, $SD_{recency} = .21$, $SD_{primacy} = .21$, $g = .66$). Second, the value-ignorance effect replicated, with riskier choices under value-ignorance ($M_{standard} = .41$, $M_{val-ign} = .47$, $SD_{val-ign} = .23$, $SD_{standard} = .21$, $g = .30$).

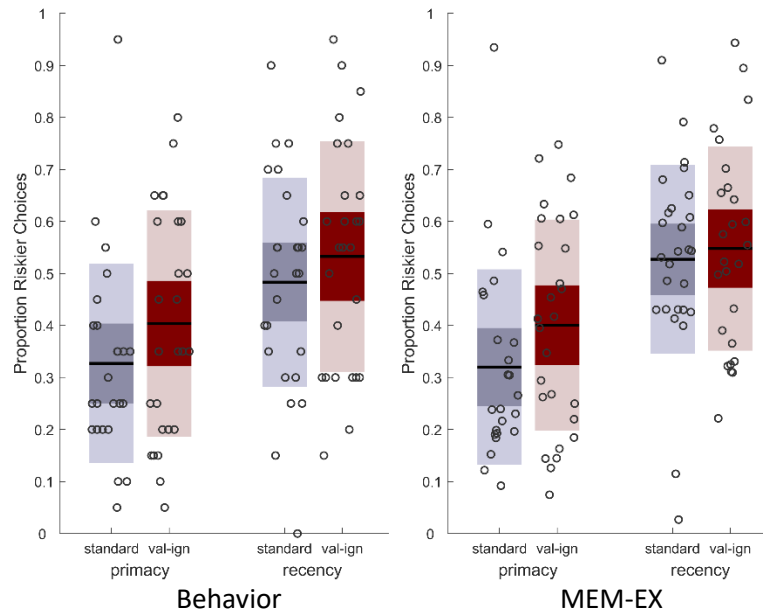


Figure 6. Behavior and MEM-EX predictions in Experiment 4. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

4.2.2. Modeling Analysis

To explain the observed effect of outcome order MEM-EX uses its *memory confusion* mechanism. Figure 6 shows that this indeed produces the kind of order effects we observe in behavior. When rewards appeared at the beginning of the sequence (Primacy), they were more likely to be replaced in memory later with \$0 outcomes. This will tend to decrease the subjective value of the riskier option more than the safe option because the former involves a smaller number of higher magnitude rewards, so each replacement has a larger impact. In contrast, when rewards appeared at the end of the sequence (Recency), the effect was in the opposite direction. For the riskier alternative, early 0s would be replaced with larger magnitude rewards than for the safer option, causing its subjective value to rise

more. According to MEM-EX, memory confusions were made for an average of 10.07% of exemplars after each sample ($\bar{\phi} = .101$, $SD = .097$).

Once again, MEM-EX captures the observed value-ignorance effect, with higher value-assignment error rates under value-ignorance ($\bar{\lambda} = .116$, $SD = .109$) than in the Standard condition ($\bar{\lambda} = .077$, $SD = .100$).

4.3. Discussion

In Experiment 4 we see that the order of outcomes can influence DfE, and that MEM-EX provides new insights into this behavior. The memory confusion mechanism provides an intuitive explanation for the observed order effects. A natural interpretation of this mechanism is that, as a decision maker performs the operation to add a new sample to memory, this operation might also be mistakenly applied at the location of a previously remembered item.

5. Experiment 5

In Experiment 5, we investigate how the number of experienced outcomes affects DfE. MEM-EX predicts that sampling a greater number of items provides more opportunities for memory confusion errors and retroactive interference. However, the direction and magnitude of the predicted effect depends on several factors, including gamble variables, outcome orders, and model parameters. We therefore use Experiment 5 to explore the impact of sample size across of range of new gambles.

5.1. Method

5.1.1. Participants

82 (23 female, age 18-27, $M = 19.21$, $SD = 1.65$) UNSW students participated and received course credit plus a monetary bonus (\$0 - \$22) based on a randomly selected trial.

5.1.2. Procedure, Material, and Design

Experiment 5 was similar to previous experiments, with a few exceptions. As in all experiments except Experiment 3, the presence of value information was manipulated between subjects. Sample size

was manipulated within subjects, with each participant receiving each pair of gambles twice. In the Small Sample condition, sample sizes ranged from 8 to 12 outcomes. In the Large Sample conditions, sample sizes were three times larger, and ranged from 24 to 36. Small and Large trials were randomly intermixed, with the restriction that the same gamble pair could not repeat on consecutive trials.

Sixteen gamble pairs were created with the aim of presenting participants with a new and diverse set of problems. This new set of gambles had a wider range of expected values and of expected value differences between alternatives (see Supplementary Materials for details). Each gamble pair was presented twice, for a total of 32 trials.

5.2. Results

Unsurprisingly, the value-ignorance effect replicated in Experiment 5. Figure 7 shows that value-ignorance again produced riskier choices ($M_{standard} = .37$, $M_{val-ign} = .46$, $SD_{standard} = .18$, $SD_{val-ign} = .25$, $g = .40$). To analyze the effect of sample size we began by computing an effect score, s , for each trial and individual. For a given trial, if a participant chose the same option in both Sample Size conditions, $s = 0$. If they chose the safe option in the Small Size condition and the risky option in the Large Size condition, $s = 1$. If they chose the risky option in the Small Size condition and the safe option in the Large Size condition, $s = -1$. Figure 8 shows that mean effect scores were roughly centered on 0, with substantial individual differences. There was no substantial difference in choices across size conditions ($M_{small} = .42$, $M_{large} = .41$, $SD_{small} = .23$, $SD_{large} = .23$, $g = .05$).




Figure 7. Behavior and MEM-EX predictions in Experiment 5. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.




Figure 8. Mean sample size effect scores in Experiment 5. Error bars indicate standard errors.

The absence of a robust group-level sample size effect fits with insights from MEM-EX. We found that the majority of individuals in Experiment 5 did not show evidence of the model's memory confusion mechanism. They were instead best accounted for with the same version of MEM-EX used in Experiment 3 (MEM-EX_{base}, see Table 1). Again, its value-assignment errors (6.85% for Standard vs. 10.01% for Value-Ignorance) accounting for differences across value conditions.

5.3. Discussion

These results highlight the limitations of using behavioral data alone to test hypotheses about cognitive mechanisms. Although we failed to detect a sample size effect in choices, we must take care to interpret this properly. If sample size effects varied across individuals, gambles, and outcome sequences – as MEM-EX predicts – collapsing across any of these factors would obscure individual-level effects. This poses problems for future investigations because these sources of variability are largely unavoidable. For example, one alternative approach would be to design trials so that predicted sample-size effects are extreme enough to assure that all participants' choices move in the same directions. However, this amounts to finding gamble pairs for which most participants have very strong preferences, and would likely result in ceiling or floor effects. Designing trials where one alternative is significantly more attractive also runs the risk of encouraging new decision strategies.

Given that our theory does not predict a simple difference between experimental conditions, we now turn to an alternative approach focused on testing psychological theories through model comparison. This method has the advantage of incorporating many sources of variability because they naturally interact with model mechanisms. That is, models like MEM-EX provide a formal theory for how experimental manipulations should affect behavior in complex and hard-to-intuit ways. As we show below, using these models, rather than experimental variables, to organize analyses can offer a more comprehensive and psychologically principled means of understanding DfE.

6. Model Comparison

To complement and extend our behavioral analyses, we conducted a quantitative comparison of models for each experimental dataset. The majority of these models are constructed within the MEM-EX framework, with versions sharing the same foundation, but differing in their specific cognitive mechanisms. In cases where we compare multiple versions of MEM-EX (Experiments 3 and 5), we can examine the importance of these cognitive mechanisms while holding constant less-interesting auxiliary

assumptions. That is, we vary hypotheses about specific mechanisms, but maintain the fundamental assumption that decision makers rely on an analogical representation of sampled outcomes in memory. To assess the overall merits of the MEM-EX framework, we also compare it to a simple alternative. The *Value-Updating* model is built of the basic principles of reinforcement learning, and provides a handy baseline characterization of performance for any novel model to achieve.

6.1. MEM-EX Models

We focus our analyses on models that offer psychologically plausible explanations for the behavioral patterns we observe in each experiment. Because the experiments involved different manipulations, for each dataset we test the versions of MEM-EX that include only the relevant cognitive mechanisms. Table 1 summarizes the four versions of MEM-EX we consider. MEM-EX_{base} is the standard version of the model, with value-assignment error and memory-sampling error mechanisms. MEM-EX_{null} is a simpler model, nested within MEM-EX_{base}, which only uses one value-assignment error-rate parameter across Standard and Value-Ignorance conditions. In Experiment 3, we compare these two versions of MEM-EX to test if error rates differed across conditions. MEM-EX_{prime} includes the memory-priming mechanism to explain the salience manipulation in Experiments 1 and 2. MEM-EX_{conf} includes the memory-confusion mechanism used to explain order effects in Experiments 4 and 5.

Table 1. *Cognitive mechanisms present in each model.*

<i>Cognitive Mechanisms</i>				
Model	Mechanism			
	Value-assignment error	Memory-sampling error	Memory Priming	Memory Confusion
MEM-EX _{base}	x	x		
MEM-EX _{null}	x	x		
MEM-EX _{prime}	x	x	x	
MEM-EX _{conf}	x	x		x

6.2. Value-Updating Models

The *Value-Updating* model (ValUp) stands in contrast to the central claim of MEM-EX. Rather than storing a record of each outcome, ValUp proposes that decision makers simply track the subjective value of each gamble, updating this impression after each new outcome. This modeling framework is built on principles of simple reinforcement learning, which have proved powerful tools for characterizing and predicting experience-based choices (Sutton & Barto, 1998). Under this view, ValUp serves as a theoretically relevant baseline for comparison. If MEM-EX can outperform ValUp in predicting choices and can also provide valuable insights into underlying cognitive mechanisms, this would support its theoretical claims. If instead, ValUp is the more accurate predictive model, this may lead us to question the utility of MEM-EX's core memory assumptions.

In the Standard condition, ValUp posits that decision makers use outcomes to learn the subjective value of each alternative. After sampling a ball from Option A, the value for Option A is updated according to $V_A^{t+1} = V_A^t + \delta(o_A^t - V_A^t)$, where V_A^t is the estimated value after sample t, o_A^t is the value of the sampled ball, and δ is a learning rate parameter between 0 and 1. An identical process governs the updating of V_B^t . After completing sampling, the decision maker compares the estimated value of each option and chooses accordingly.

In the Value-Ignorance condition, outcome values were absent during sampling, so ValUp focuses on learning the reward probabilities for each alternative. After each sample, the decision maker updates the probability of receiving a reward from Box A, according to $p_A^{t+1} = p_A^t + \delta(c_A^t - p_A^t)$, where p_A^t is the estimated probability of a reward at sample t and c_A^t is the outcome of the sample (1 if a reward, otherwise 0). On the choice screen, the values of the rewards from each box are revealed, and the decision maker integrates these by multiplying each value with its learned probability. Again, the decision maker compares the estimated value of each option to make a choice. Importantly, these different learning methods may give rise to different behavior across conditions, potentially accounting for the value-ignorance effects we observe in behavior. The model uses the same simple risk bias

mechanism as MEM-EX to represent individual differences in overall risk attitude. To give ValUp the best chance of fitting behavioral data, we initialized the values of V and p to the mean expected value and reward probability, respectively, in each experiment.

This simple version of ValUp performed poorly in our model comparison, so we focus on a version with one additional mechanism. This model incorporates the possibility that decision makers sometimes ignore information. After sampling a new outcome, updating occurs with probability ψ , otherwise the sample is ignored and no learning occurs. This mechanism plays a similar role to MEM-EX's memory-sampling error mechanism by controlling the number of outcomes a decision maker uses to make their decision. This additional flexibility improved the model, though performance was still relatively poor. It is likely that the ValUp framework could be improved by introducing of additional components to account for specific behavioral effects, but this simple version is sufficient as a benchmark standard of model performance. We leave the development of a theoretical explanation for our results in terms of reinforcement learning for future work.

6.3. Method

Our central aim in this article is to find cognitive mechanisms that give good explanations of behavioral phenomena in DfE. However, to test each model's ability to predict behavior while avoiding the problem of overfitting, we used a cross-validation analysis (Busemeyer & Wang, 2000). For each participant in a given experiment, we randomly selected half of their choices to be in the training set and half to be in the validation set. All models assumed a binomial error process to connect their predicted choice probabilities to observed choices, and the parameters for each model were fit to the responses from each individual's training set. We used Matlab's *ga* genetic algorithm (Chipperfield & Fleming, 1995) to maximize the likelihood of responses according to each model. The best-fitting parameters value for each individual were then used to predict responses for trials in the validation set. The accuracy of these predictions indicates how well each model explains behavior. Comparing this

measure across models tests how well each accounts for behavioral patterns, while also implicitly taking model complexity into account. An overly flexible model might overfit the training set, causing it to perform badly when making out-of-sample predictions to the validation set. The entire cross-validation process was repeated ten times, with each replicate making a new random allocation of responses to training and validation sets.

6.4. Results

We calculated the log-likelihood of the validation set for each model, and took the average across the ten replicates to summarize the results of cross-validation. Table 2 summarizes our findings in terms of the number and proportion of individuals for whom each model gave the largest log-likelihood in the validation set. Although results varied across experiments, the MEM-EX framework always yielded the best performing model, while ValUp struggled with cross-validation, and was the worst performing model in each experiment.

Table 2. *Number (and proportion) of participants for whom each model gives the best out-of-sample predictions in our cross-validation analysis.*

<i>Cross-Validation Results</i>				
Model	Experiment			
	1 & 2	3	4	5
MEM-EX _{null}	—	27 (64.29%)	—	—
MEM-EX _{base}	—	13 (30.95%)	—	62 (75.61%)
MEM-EX _{prime}	282 (86.50%)	—	—	—
MEM-EX _{conf}	—	—	91 (87.50%)	11 (13.41%)
ValUp	44 (13.50%)	2 (4.76%)	13 (12.50%)	9 (10.98%)

In the following sections, we examine the modeling results for each experiment in more detail.

To quantify and visualize the relative performance of these models, we converted individual mean

(across cross-validation replicates) log-likelihoods to probabilities (labeled as *model weight* in the figures below)¹⁰.

6.4.1. Experiments 1 & 2

For the dataset combining Experiments 1 and 2 we tested two models, with the goal of explaining both the value-ignorance effect and the salience effect. Figure 9 shows individual, mean, and median model weights. MEM-EX_{prime} was the better performing model, and gave the best account of choices for the majority of individuals (87%). As noted earlier, this model is successful in explaining group-level effects; using value-assignment errors for the value-ignorance effect, and memory priming for the salience effect. The cross-validation results indicate that the model also best explained choices on an individual-level. ValUp performed poorly, and was the preferred model for only 14% of individuals. This result shows that MEM-EX can predict choices more accurately than a reasonable baseline model, while also providing new insights into the cognitive mechanisms underlying behavior.




Figure 9. Cross-validation model weights for Experiments 1 and 2. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

¹⁰ The weight for Model i in the set of j models was calculated as $e^{LL_i} / \sum_j e^{LL_j}$.

6.4.2. Experiment 3

In Experiment 3 we found individual differences in the within-subjects value-ignorance effect. 64% of participants were best characterized by MEM-EX_{null}, which kept value-assignment error rates constant across conditions. Most of the remaining participants (31%) preferred MEM-EX_{base}, which tended to estimate higher value-assignment error rates in the Value-Ignorance condition. ValUp struggled again, with fewer than 5% of participants attributed. These results fit nicely with those from Experiments 1 & 2, reinforcing the utility of the MEM-EX framework in general, and the value-assignment error mechanism specifically.




Figure 10. Cross-validation model weights for Experiment 3. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

6.4.3. Experiment 4

MEM-EX_{conf} was the best performing model for 88% of individuals in Experiment 4. The model's good performance here further indicates the importance of its memory confusion mechanism for explaining the impact of outcome order on choice. ValUp performed poorly, with 13% of participants best fit.




Figure 11. Cross-validation model weights for Experiment 4. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

6.4.4. Experiment 5

Experiment 5 manipulated sample size within-subjects, but failed to produce clear behavioral results. Due to various sources of between-subjects variability, we could not perform a satisfactory test of our hypotheses using independent variables alone. However, through our model comparison we gained insight into the psychological processes at work. The majority of participants (76%) were best characterized by MEM-EX_{base}'s value-assignment error and memory-sampling error mechanisms. A smaller proportion (13%) supported the addition of MEM-EX_{conf}'s memory confusion mechanism, which we expected to interact with sample size. ValUp again struggled, and was the preferred model for 11% of individuals.




Figure 12. Cross-validation model weights for Experiment 5. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

6.5. Discussion

These results were remarkably consistent across all experiments, and – together with the qualitative results presented in earlier sections – provide converging evidence for the MEM-EX framework. In presenting the behavioral results from Experiments 1-5 we showed how MEM-EX offers simple, intuitive explanations for behavior using cognitive mechanisms. In our model comparison we used cross-validation to show that these explanations also provide parsimonious accounts of behavior, without overfitting. To further emphasize the utility of cognitive modeling, consider the analysis of Experiment 5 depicted in Figure 13. Here we see how difficult it can be to understand behavior in terms experimental variables alone. For each panel, gambles are sorted according to a different experimental variable. One might expect to see clear patterns emerge, however – with the possible exception of expected value differences (top row) – there are no obvious relationships linking behavior to any measure. This contrasts sharply with the depth of insight we can gain through MEM-EX. What might first

appear to be mysteriously random choice can be succinctly characterized in terms of cognitive mechanisms.

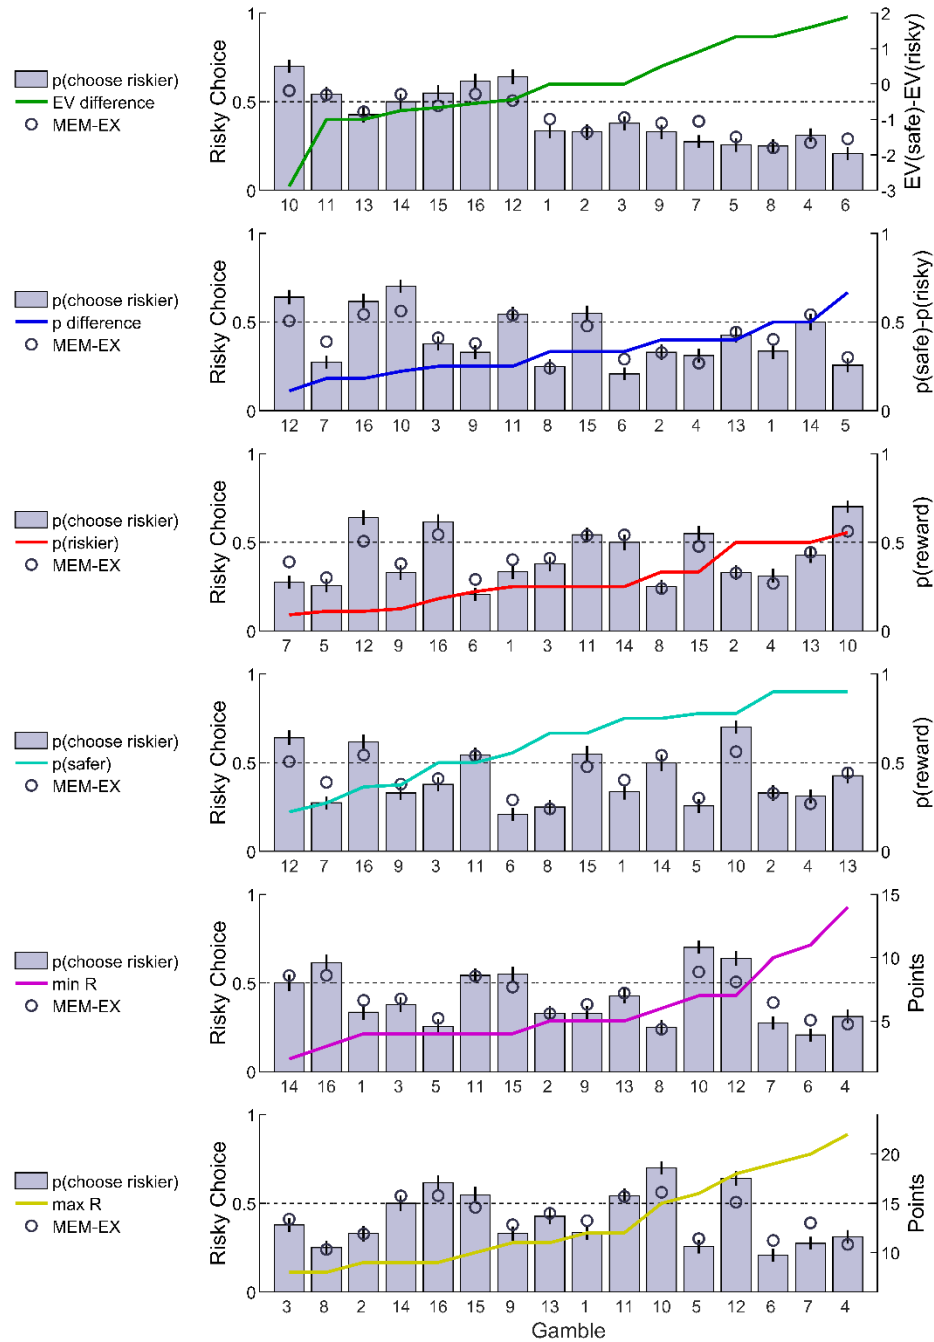


Figure 13. Shaded bars indicate the mean proportion of choices in favor of the riskier alternative in Experiment 5 (left axis). Each row is sorted by the A) difference in expected value across gambles, B) difference in reward probabilities across gambles, C) probability of receiving a reward from safer alternative, D) probability of receiving a reward from riskier alternative, E) reward value for the safer

alternative, and the F) reward value for the riskier alternative (left axis). Error bars indicate standard errors. Circles indicate predictions from MEM-EX.

We contrasted MEM-EX with ValUp, which employs a fundamentally different approach to DfE. This model had consistently poor performance and never best fit more than 20% of individuals. To understand why, we reanalyzed Experiment 5. Rather than cross-validating the models, we fit them to each individual, using 100% of responses to optimize parameter values and maximize log-likelihoods. The results in Figure 14 show a starkly different picture to that observed in cross validation, with ValUp performing better than both MEM-EX models. This reversal (compared to Figure 12) indicates that ValUp is overfitting. In cross-validation, overfitting hurts the model's performance because, in optimizing its parameter values to the training set the model tunes its predictions to random noise. The model then struggles to generalize to the validation set, where those noisy patterns do not hold.




Figure 14. Model weights based on maximum-likelihood fit for Experiment 5. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations.

In contrast, Figures 15 and 16 show how MEM-EX deals more appropriately with different patterns of behavior. Since Experiments 1-4 used the same gambles, we can combine their results into

one figure. Figure 15 plots mean choice proportions for each experiment (symbols) and mean predictions across all four experiments for MEM-EX_{base}¹¹ (solid line) and ValUp (dotted line). In Figure 15, sorting gambles by risky-choice proportion shows a reliable behavioral trend across experiments, which both models capture quite well. However, Figure 16 shows that the result is different when we look at the value-ignorance effect. This figure plots mean effect sizes, and shows much more variability, with values fluctuating a great deal across gambles and experiments. MEM-EX largely ignores these fluctuations, but nonetheless captures the overall positive effect. By contrast, ValUp make more erratic and variable predictions that do not match behavior. These results suggest that MEM-EX's foundation in cognitive principles constrains its ability to fit cross-experiment variability. This gives the model an advantage over ValUp, whose general-purpose learning mechanisms lead to more variable predictions.




Figure 15. Mean proportion of choices in favor of the riskier alternative in Experiments 1-4. Symbols denote behavior, the solid line indicates predictions from MEM-EX, and the dotted-line indicates predictions from ValUp. Error bars indicate standard errors across all 632 participants in Experiments 1-4.

¹¹ We use MEM-EX_{base}, rather than the best-fitting version of MEM-EX from each experiment, because this allows us to examine how a single model deals with cross-experiment behavioral variability.




Figure 16. Mean value-ignorance effect size in Experiments 1-4. Symbols denote behavior, the solid line indicates predictions from MEM-EX, and the dotted-line indicates predictions from ValUp. Error bars indicate standard errors across experiments.

7. General Discussion

7.1. Summary of Behavioral Results

Across five experiments, we explored several factors influencing experience-based decision making. Time and again, we found the presentation of outcome values to reliably affect people's choices. When value information was available during sampling, participants were less likely to choose the risky option compared to when this information was withheld until after sampling. Although the magnitude of this effect was variable and subject to individual differences, it reliably replicated in each experiment, both within and between subjects.

In Experiments 1 and 2 we also examined the impact of perceptually highlighting outcomes during sampling. In the Salience condition, we found that highlighting rare rewards led people to make riskier choices, as if these events were more prominent or available in memory. In Experiment 4, we also found choices to be affected by outcome order. When rewards appeared early in the sampling

sequence, participants preferred the safer alternative. However, when rewards appeared at the end of the sequence, people made riskier choices.

Our effort to study the role of sample size in Experiment 5 was less successful. We failed to find reliable group-level effects, which may indicate that individual differences produced noisy effects that varied across gambles. This limitation served to emphasize the importance of using cognitive models to more directly interrogate the psychological processes underlying choice behavior.

7.2. Mechanisms of Experience-Based Choice

We used computational models to better understand how the above factors affected decision making. After comparing the performance of competing models, we found strong support for the MEM-EX framework, with four cognitive mechanisms important for explaining behavioral patterns. *Value-assignment errors* accounted for differences in risky choice between Standard and Value-Ignorance conditions by positing that individuals misremembered outcome values with greater frequency under value-ignorance. Functionally, this mechanism mimics ‘overweighting’ of rare events, in that errors effectively reduce the difference in frequency between rare and common events. In this sense, value-assignment errors provide a mechanistic interpretation of the results that Hotaling et al. (2019) found using CPT.

Memory-sampling error was also an important mechanism for capturing behavior in each experiment. It provides a psychologically plausible mechanism for explaining why decision makers fail to maximize. Unlike many alternative error mechanisms – such as ‘trembling-hand’ or softmax (Luce, 1959) – memory-sampling error is couched in terms of well-known psychological constructs. Its virtue can also be seen in its interaction with another cognitive mechanism, *memory priming*. These combine to explain the observed salience effects in Experiments 1 and 2. Memory-sampling error posits that decision makers have limited attention, and therefore sample a subset of information from memory to make a choice. Memory priming adds the intuition that perceptually highlighted events are more salient, and

are therefore more likely to come to mind when sampling from memory. These simple mechanisms formally instantiate many of the ideas contained within extant verbal theories, such as availability. Also, by connecting to concepts like attention and working memory, they yield clear predictions that pave the way for future tests and further theory development.

Finally, we found evidence that *memory confusion* played an important role in many people's decisions. This mechanism describes a process whereby new experiences replace older ones in memory. Its effects are most obvious in Experiment 4, where memory confusion produced retroactive interference effects implying greater weighting of recent outcomes. Once again, this mechanism allows us to recast non-mechanistic theory in terms of cognitive mechanism. Rather than appeal to the abstract notion of *recency bias*, we can now articulate a psychological process that produces order effects.

7.3. Reinforcement-Learning Models of Experience-Based Choice

In addition to the explanatory value of MEM-EX's cognitive mechanisms, we were interested in the accuracy of its quantitative predictions. We contrasted the framework's performance with that of an alternative baseline model, which posits that decision makers learn the subjective value of each alternative through repeated updating of values. Similar reinforcement-learning theories have been used to model choice behavior in various contexts (e.g. Hertwig et al., 2006). However, our findings suggest that the theoretical assumptions in MEM-EX allow the model to surpass a simple version of ValUp. The lack of theoretical constraint placed on the general-purpose learning algorithm used in ValUp grants it the flexibility to apply to various DfE phenomena. Unfortunately, this universality is a detriment when making predictions under cross-validation, where ValUp succeeds in fitting choice data, but fails to predict choices when generalized. To achieve better performance, the model requires a psychological theory to constrain its behavior in accordance with the experimental and psychological context under consideration. The introduction of additional, cognitively inspired mechanisms to the ValUp framework may provide the key to avoiding overfitting.

7.4. Future Directions

Our findings motivate several paths for future study. Applying our model-based analysis to other DfE paradigms may provide new insights into their unique behavioral patterns. For example, in repeated choice, where every action is consequential, how does memory for past outcomes support the balance between exploration and exploitation (Plonsky, Teodorescu, & Erev, 2015)? What cognitive mechanisms best explain learning and adaptation in dynamic environments, where payoffs change over time or as a consequence of the decision makers actions (Hotelling, Navarro, & Newell, 2018, accepted; Navarro, Newell, & Schulze, 2016)? MEM-EX represents a valuable new tool for investigating these questions.

We also hope to deepen our understanding of the cognitive mechanisms described by MEM-EX. For instance, there are presently multiple interpretations of the model's different value-assignment error rates. Might the increased errors under value-ignorance result from a greater cognitive load imposed by the temporal separation of frequency and value information? Or is it the act of 'reopening' one's memory to assign values during the choice phase that produces these errors? New experimental manipulations can shed light on these issues.

Future studies will also examine the role of uncertainty in DfE. To this end, we have begun investigations into new choice scenarios with reduced memory demands. Using a design similar to that of Experiment 4, we tested the effects of value-ignorance and outcome order when participants were certain of the observed outcome sequence. During sampling, the history of sampled outcomes from each box was displayed as a series of balls at the top of the screen, obviating the need to remember the sequence. We find a somewhat puzzling pattern of results.¹² We replicate the value-ignorance effect from Experiments 1-5, but only in the Primacy condition, and we replicate the outcome-order effect from Experiment 4, but only in the Standard condition. Curiously, these effects disappear (and slightly

¹² Results from this pilot experiment can be found at <https://osf.io/x7uqw>.

reverse) in the Recency and Value-Ignorance conditions, respectively, suggesting that participants performed the task differently when they were certain of the outcomes they had sampled.

Future work can also elucidate individual differences by relating constructs like working memory capacity to cognitive mechanisms like memory sampling error (Olschewski, Rieskamp, & Scheibehenne, 2018). Such studies may also allow greater insight into more significant differences in decision strategies (i.e. what factors predict whether an individual uses MEM-EX vs. ValUp?).

Acknowledgements

JMH, CD, and BRN were supported by the Australian Research Council (DP160101186). AJ was supported by a British Academy Postdoctoral Fellowship (PF150005). We thank Jake Embrey, Garston Liang, and Dominic Tran for help with data collection.

References

- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: making choices without trade-offs. *Psychological Review*, 113(2), 409.
- Camilleri, A. R., & Newell, B. R. (2011a). Description-and experience-based choice: Does equivalent information equal equivalent choice? *Acta Psychologica*, 136(3), 276-284. doi:10.1016/j.actpsy.2010.11.007
- 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(2), 377-384.
- Chipperfield, A., & Fleming, P. (1995). The MATLAB genetic algorithm toolbox.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., . . . Lebiere, C. J. J. o. B. D. M. (2010). A choice prediction competition: Choices from experience and from description. 23(1), 15-47.
- Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, 36(2), 153-177.
- Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General*, 145(4), 486-508.
- 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.
- Hawkins, G., Camilleri, A., Heathcote, A., Newell, B., & Brown, S. (2014). Modeling probability knowledge and choice in decisions from experience. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the annual meeting of the 36th annual cognitive science society* (pp. 595-600). Austin, TX: Cognitive Science Society.
- Hedges, L. V. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics*, 6(2), 107-128. doi:10.2307/1164588
- 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(8), 534-539.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006). The role of information sampling in risky choice. In *Information sampling and adaptive cognition*. (pp. 72-91). New York, NY, US: Cambridge University Press.
- 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. 209-235): Oxford University Press.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225-237.
- Hotaling, J. M., Jarvstad, A., Donkin, C., & Newell, B. R. (2019). How to Change the Weight of Rare Events in Decisions From Experience. 30(12), 1767-1779. doi:10.1177/0956797619884324
- Hotaling, J. M., Navarro, D. J., & Newell, B. R. (2018). Skilled bandits: Learning to choose in a reactive world. In C. Kalish, M. Rau, J. Zhu, & T. T. Rogers (Eds.), *Proceedings of the 40th annual meeting of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Hotaling, J. M., Navarro, D. J., & Newell, B. R. (accepted). Skilled bandits: Learning to choose in a reactive world. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. doi:10.31234/osf.io/ghmdx
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision making under risk. *Econometrica*, 47, 263-292.

- Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, 25(2), 143-153.
- Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, 116(2), 286-295.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, 125(1), 1-32.
- Lin, D., Donkin, C., & Newell, B. R. (2015). The exemplar-confusion model: An account of biased probability estimates in decisions from description. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual meeting of the cognitive science society* (pp. 1374-1379). Austin, TX: Cognitive Science Society.
- Luce, R. D. (1959). *Individual choice behavior; a theoretical analysis*. New York: Wiley.
- Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore-exploit dilemma in static and dynamic environments. *Cognitive psychology*, 85, 43-77.
- Olschewski, S., Rieskamp, J., & Scheibehenne, B. (2018). Taxing cognitive capacities reduces choice consistency rather than preference: A model-based test. *Journal of Experimental Psychology: General*, 147(4), 462-484. doi:10.1037/xge0000403
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, 122(4), 621-647. doi:10.1037/a0039413
- Rakow, T., Demes, K. A., & 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(2), 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), 1-14.
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(6), 1446-1465.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). Cambridge: MIT press.
- Tversky, A., & Kahneman, D. (1973). Availability: a heuristic for judging frequency and probability. *Cognitive psychology*, 5(2), 207-232.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297-323.
- Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, 144(2), 140-176.
- Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, 19(1), 1-16.