



City Research Online

City, University of London Institutional Repository

Citation: Hotaling, J. M., Donkin, C., Jarvstad, A. & Newell, B. R. (2022). MEM-EX: An exemplar memory model of decisions from experience. *Cognitive Psychology*, 138, 101517. doi: 10.1016/j.cogpsych.2022.101517

This is the accepted 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/28859/>

Link to published version: <https://doi.org/10.1016/j.cogpsych.2022.101517>

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.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

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

Jared M. Hotaling^{1,2}, Chris Donkin^{1,4}, Andreas Jarvstad³, Ben R. Newell¹

¹University of New South Wales

²University of Illinois, Urbana-Champaign

³City, University of London

⁴Ludwig Maximilian University of Munich

Please send all correspondence to:

Jared Hotaling

Department of Psychology

University of Illinois, Urbana-Champaign

Champaign, IL 61820

hotaling@illinois.edu

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. We compare these to three alternative models; two baseline models built on the principle of expected value maximization, and another employing a suite of choice methods previously shown to perform well in prediction tournaments. We find that MEM-EX consistently outperforms these competitors, demonstrating its value as a tool for making quantitative predictions without overfitting. We discuss the implications of these findings for our understanding of the interplay between attention, memory, and experience-based choice.

Keywords: decisions from experience; computational models; decision making; cognitive mechanisms; exemplar memory; BEAST.

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 include 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 (Erev, Plonsky, & Roth, 2020).

In recent years there has been a great deal 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).

Here we use a combination of experimental manipulations and cognitive modeling to probe the psychological processes underlying experience-based choices. Our specific goal is to investigate the cognitive mechanisms by which people record and update experiences in memory. In so doing, we aim to further elucidate the memory processes involved in decisions-from-experience.

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. One might think of the sampling paradigm as an experimental analog to researching products before a purchase, or imagining the consequences of a decision before it is made (cf. Lieder, Griffiths, & Hsu, 2018). Furthermore, there are situations in which people literally ‘try-before-they-buy’ (e.g., magazine, online media-subscriptions). Here, a period of costless sampling precedes a consequential commitment, much like in the sampling paradigm. The clear procedural advantage of using the sampling 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; Fox & Hadar, 2006; Rakow, Demes, & 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 the Technion Prediction Tournament makes a valuable contribution to the literature on DfE by discovering models with impressive predictive accuracy, we pursue a different goal in our research. Rather than focus on predictive accuracy, we aim to develop a deeper understanding of the

psychological processes underlying DfE, and therefore use behavioral data to test how well cognitive mechanisms explain choice patterns.

We aim to examine the hypothesis that the theoretical framework of *exemplar memory* can explain how people make decisions from experience. We see this approach as complementing existing models that adopt similar frameworks as well as descriptive theories that use alternative frameworks, such as CPT, to characterize behavior in terms of preference functions. Rather than offering an alternative account, we aim to unpack and further explain behavior in terms of psychological processes. To preview our results, we find that models constructed within this exemplar memory framework (the *memory for exemplars model* or MEM-EX – outlined in detail below) provide accurate and parsimonious explanations for various behavioral patterns. To complement these findings, we compare MEM-EX to two simple baseline strategies. These models represent the hypotheses that individuals choose so as to maximize their expected value (EV), either with (*Logistic 2*) or without (*Logistic 1*) a bias for or against making risky choices. Comparing these models to MEM-EX allows us to assess the additional value provided by memory-based cognitive mechanisms in predicting choices and explaining behavior.

We also consider another alternative model, the *best estimate and simulation tools* (BEAST) because it has previously been found to perform well in choice prediction tournaments involving many kinds of DfE (Erev, Ert, Plonsky, Cohen, & Cohen, 2017). Our purpose for including BEAST is to compare MEM-EX to a well-supported alternative that shares some elements of memory-based sampling. However, our aim is not to definitively answer the question of whether other theoretical frameworks can be adapted to account for all our observations. Instead, our goal is to compare the exemplar memory framework to reasonable benchmark and alternative models that can provide parsimonious accounts of other DfE, but for which we did not make additional theoretically motivated assumptions to accommodate the specific outcomes we observe.

Our framework shares the foundational assumption that individuals draw from memory of past outcomes to make decisions with the *instance-based learning model* (IBL; Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012), which has been used to account for DfE in both one-shot and repeated-choice paradigms. At its core, the IBL 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.

We refrain from directly comparing our MEM-EX framework 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. For example, IBL determines memory activation not only by combining recency and frequency of experienced outcomes, but also by including a noise parameter which is “important for capturing the variability of human behavior” (Gonzalez & Dutt, 2011, p. 527). The framework we develop can be seen as a complement and extension to this approach wherein we attempt to specify *how* such variability is introduced via theoretically-derived cognitive mechanisms, and so hope to clarify and build upon this kind of ‘catch-all’ noise mechanism. To preview our results, we find evidence supporting IBL’s core assumption that event memory drives DfE, but we go beyond IBL by describing mechanisms that capture how memory operates in specific settings, which tend to differ from those for which IBL was originally designed. We contend that a direct quantitative comparison with IBL would offer little in additional explanatory value

given the overall goal of our modelling exercise, since our experiments were not designed to require the cognitive machinery implemented in IBL.¹

1.4. The Memory for Exemplars Modeling Framework

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. (For applications of a model within the same framework of exemplary memory confusion to choices among multiple outcomes see Hawkins et al., 2014). We 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 of the alternatives presented to participants in our experiments (see Figure 1). 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], then $V_A = 2.0$ and $V_B = 1.8$.

¹ We also note that the success of IBL in predicting one-shot DfE choices is attributable to some degree to its ability to model participants' sampling behavior – which option they choose to sample from on each trial and when they stop search. We cannot model this behavior in our experiments – because search is exhaustive rather than self-terminating – thus, again making a direct comparison to IBL inappropriate.

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 require a greater standard of evidence to overcome that default. Risk bias takes the form of a constant added to the subjective value of the riskier alternative. That is, the model predicts that individuals will choose the riskier option – here A, which offers a lower probability higher value reward – if $V_A + \beta > V_B$, and will choose the safer Option B if $V_A + \beta < V_B$ ². β is a free parameter (constrained to be between -10 and 10 in model fits) representing the amount of additional evidence required to overcome an individual's default preference. Values below 0 indicates a bias toward the safer option, while values above 0 correspond to a bias in favor of the riskier alternative. We define the riskier alternative as that which offers the lower probability of a reward. In MEM-EX, this determination is based on the sample of items drawn from memory, rather than the (objectively true) set of events observed during sampling (see Section 1.4.2. more details on MEM-EX's memory sampling error mechanism).

1.4.1. Value-Assignment Error

To this basic framework, MEM-EX adds four cognitive mechanisms that influence how information is processed. Over the course of the article, we will explain why these mechanisms were necessary to add to the model to account for the empirical phenomena observed in our experiments. According to the model, when an outcome is sampled, the event is recorded in the appropriate memory store as an exemplar and is assigned the observed value. However, it is unlikely that any cognitive mechanism is infallible (Simon, 1982). The *value-assignment error* mechanism therefore 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 option³. For instance, if

² A random choice is made if $V_A + \beta = V_B$.

³ Though theoretically possible, we did not encounter any observations that required between-alternative confusion.

Option A produces outcomes of either 10 or 0, MEM-EX posits that at the time of choice, an observed outcome of 10 may be mistakenly remembered as a 0 (and vice versa). The free parameter λ (with values between 0 and 0.5 in model fits) 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 options⁴. Since there were only ever two outcomes in the experiments considered here, we do not need to assign probabilities to the possible incorrect memories, though this is an interesting theoretical issue.

To implement this value-assignment error mechanism in MEM-EX, we need not make any assumptions about when exactly the memory for the value of an outcome becomes faulty, since the final representation at choice is all that matters. As such, we remain largely agnostic as to whether value-assignment errors occur at the time of encoding, retrieval, or somewhere in between. While it is unlikely that at the time of encoding the participant mistakenly believes that the 10 they see is a 0, momentary distraction or confusion while memory traces are established could lead to error-prone encoding. Retrieval errors also seem likely, because the contexts of the study events are rather indiscriminable, and so individual memory traces may be largely interchangeable (Dennis & Humphreys, 2001; Shiffrin & Steyvers, 1997). In the Memory Confusion section below, we describe our attempt to model errors that may occur between encoding and retrieval.

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 and Figure 1), value-assignment errors were more likely. Our speculation is that these additional errors occur because when value information is provided after sampling is finished,

⁴ For simplicity, when doing model simulations, the model effectively ‘knew’ what the possible outcomes were, and so was able to make assignment errors before both outcomes had been experienced. Though unrealistic, the ultimate effect on any claims we make is negligible, especially considered against the required increase in computational complexity to implement the realistic assumption.

participants are forced to retrieve their existing memories and update them, to include this additional value information. As such, the errors that plague typical sampling decisions may be repeated. The effect of these additional assignment errors is to 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. The idea that people rely on small or imperfect samples is common in many theories of decision making (e.g. Hertwig et al., 2004; Stewart, Chater, & Brown, 2006), with some arguing that it is adaptive in realistic choice environments (Anderson & Schooler, 1991; Hertwig & Pleskac, 2010; Plonsky, Teodorescu, & Erev, 2015). In MEM-EX, this is represented with the *memory sampling error* mechanism, according to which the individual randomly samples κ items from each memory store to compute subjective values for each alternative. In a simple version of the model, each item in memory is assumed to have an equal probability of being chosen, with items sampled from memory with replacement. For instance, imagine that $\kappa = 6$ 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 (see Section 1.4), to the benefit of Option B, which now appears more attractive due to memory sampling error. In model fits, κ was constrained to integers from 1 to 8.

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, Glzman, & Hertwig, 2008; Lieder et al., 2018). Formally, each possible outcome is assigned a weight, w , and the probability of sampling outcome i is $w_i / \sum_j w_j$.

Here, we assume two possible states of salience, 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 in model fits. Thus, perceptually highlighting an outcome increases its likelihood of being retrieved from memory.

Here again, MEM-EX provides new insights into cognitive processes producing behavior that implies overweighting of salient events. As is typically found in memory research, MEM-EX posits that items that are attended to more – for example by receiving greater looking time – are encoded more strongly and are easier to retrieve (Ratcliff, Clark, & Shiffrin, 1990; Tulving & Hastie, 1972). Although articulating the specifics of this priming process remains a goal for future research, reframing overweighting in terms of memory priming is an important step toward a unified cognitive theory of experience-based choice.

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 effectively overwrite old memories. Specifically, when an outcome is sampled, there is a probability, ϕ , that any given 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), with ϕ being a free parameter between 0 and 0.5 in model fits. Like with value-assignment errors, these confusions are assumed to occur within each alternative, with no confusions between options (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 have more opportunities to be replaced by later samples. 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 choice⁵.

1.4.5. Stochasticity

It is worth noting that MEM-EX is a stochastic model, with variability emerging naturally from its memory sampling error (κ), value-assignment error (λ), and memory confusion mechanisms (ϕ). Consequently, the model provides a deeper process account of choice variability. Rather than representing variability with algebraic noise (e.g. logistic choice sensitivity) or a separate heuristic process (e.g. trembling hand error) – both of which allow variability to change independent of other factors – MEM-EX embeds several sources of variability within the system, and so forces them to covary with other factors affecting choice. In this sense, MEM-EX constitutes a stronger theory of choice variability because the realm of possible predicted data patterns is constrained.

1.5. Goals

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

⁵ 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.

impact of several experiential factors on decision making behavior. From these we gain valuable insights into the roles that value information, outcome salience, outcome order, and sample size play in choice. Second, we use the new computational framework just outlined 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).

Although we believe that several mechanisms are present in decision makers, the evidence for a given mechanism will be a function of the experimental design. For each experiment we therefore consider versions of the model that only include mechanisms that could be plausibly affected by our manipulations. That is, without the right manipulation a mechanism may not be noticeable, but this should not undermine our belief in the mechanism based on evidence from other experiments that include an appropriate manipulation. For example, though memory confusion may be expected to influence choices in all experiments, its systematic effect on observed behavior in most experimental designs will not be large enough to warrant the additional complexity it adds to the model. Similarly, in experiments without an explicit salience mechanism, different participants may attend to different features of the outcomes (e.g., the largest or smallest values), but such effects, while present, will exert little systematic effect.

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 consider the accuracy of MEM-EX's predictions at an individual level by comparing it to three theoretically interesting alternative models. Our aim here is to test how well MEM-EX – and its underlying framework of analogical representations of experienced outcomes – predicts choices by comparing it with simple baselines (Logistic 1 and Logistic 2) and a previously successful predictive

model of choice (BEAST). 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 Hotaling, Jarvstad, Donkin, and Newell (2019)⁶. In this study, we investigated the impact of rare events on DfE using two manipulations.

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.

Hotaling 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-

⁶ The present Experiments 1 and 2 were labeled as Experiments 3 and 4, respectively in Hotaling et al, 2019.

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. Hotaling 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 the 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 that used identical methods and procedures (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.

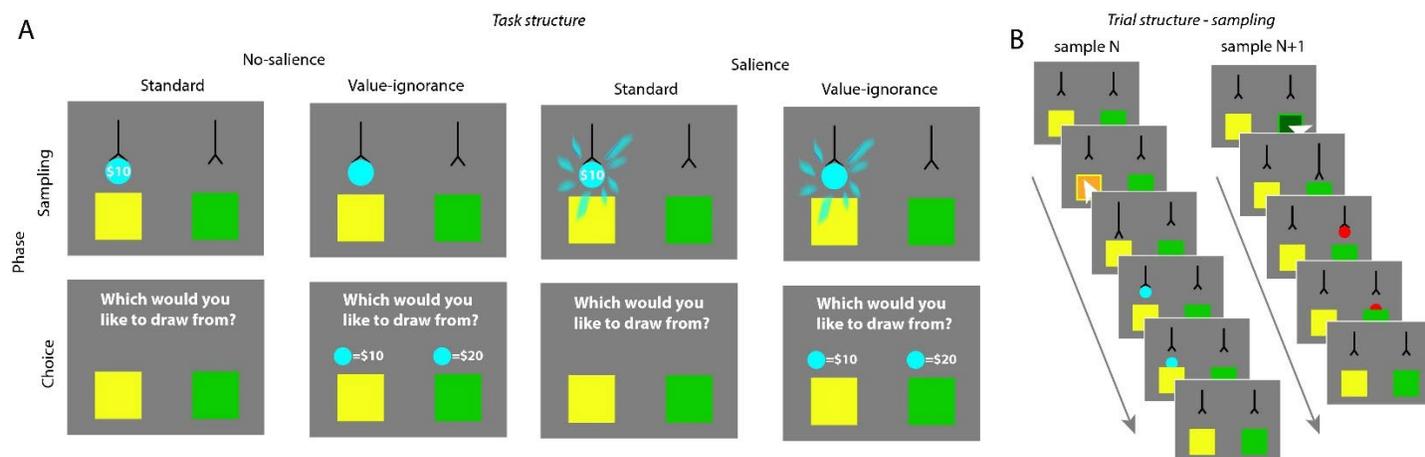


Figure 1. Robot-sampling task A) Main design. Each trial was composed of two phases: a sampling phase (top row), and a choice phase (bottom row). For a given sample, participants could observe either a blue ball, or a red ball. Red balls were worth \$0 and blue balls were worth some reward. In the Standard 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. For the salience manipulation (right two columns of A) highlighting occurred whenever participants sampled the

rare event for the riskier alternative. This was implemented in both the Standard and Value-Ignorance conditions. 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.

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 EVs. To better understand the task, consider an example trial involving a riskier option on the left offering a

1/10 (10%) chance of winning 16 points, and a safer option on the right offering an 8/10 (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 (third and fourth columns of Figure 1A). 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 ($M_{standard} = .40$, $M_{val-ign} = .46$, $SD_{standard} = .19$, $SD_{val-ign} = .26$, $g = .27$)^{7,8}. 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. The reader may be concerned that this result is simply regression to the mean, or increasing levels of indifference/random responding in the Value-Ignorance condition. In the Appendix (Figure A1) we reproduce data from Hotaling et al. (2019) that should allay such concerns. In brief, the pattern of gamble-wise effects shows that when the riskier option was unlikely to produce a reward ($p < .5$), participants in the Value-Ignorance condition made riskier choices than participants in the Standard condition. This result matches the overall effect on risky choice. However, for gambles in which the riskier option was likely to yield a reward ($p > .5$), there were either no appreciable differences between conditions or a trend toward fewer risky choices in the Value-Ignorance condition. This is the pattern one would expect if the Value-Ignorance condition induced additional value-assignment errors, as MEM-EX predicts (see Section 2.2.2. for additional discussion).

⁷ We report Hedge's g as a measure of effect size (Hedges, 1981).

⁸ The reader may worry about the "reliability" of our description of the empirical data. We note that in all experiments that 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.

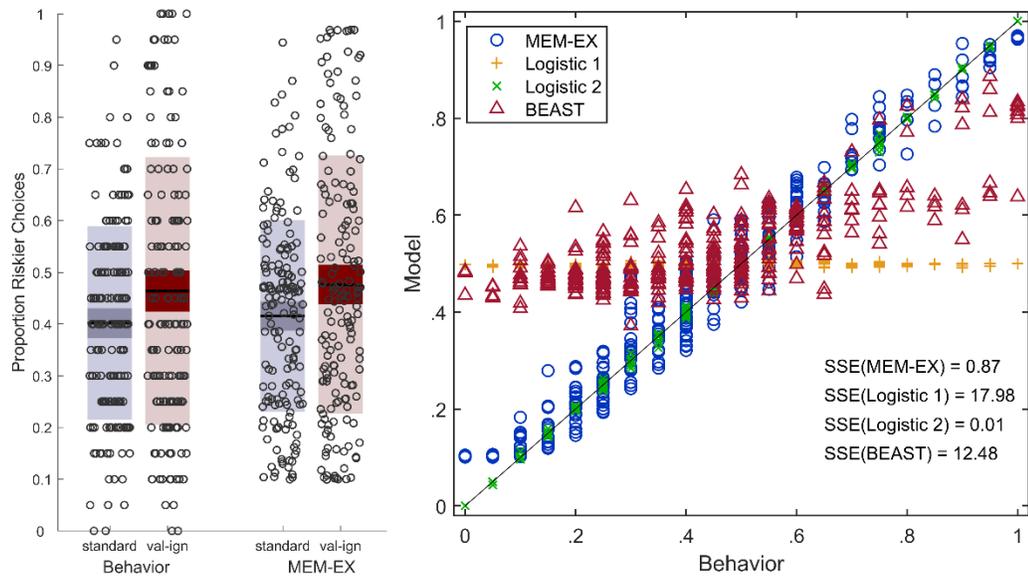


Figure 2. Behavior and model predictions in Experiments 1 and 2. The left panel shows mean proportion of choices in favor of the riskier alternative in each condition. Each dot represents an individual. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations. The right panel compares observed and predicted individual mean risky choice proportions for each model. SSE = sum of squared errors.

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 Hotaling 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 ($M_{salience} = .43$, $M_{no-salience} = .38$, $SD_{salience} = .28$, $SD_{no-salience} = .26$, $g = .19$). That is, perceptually highlighting rare events during sampling increased the likelihood that participants would choose the risky option, particularly in the Standard condition.

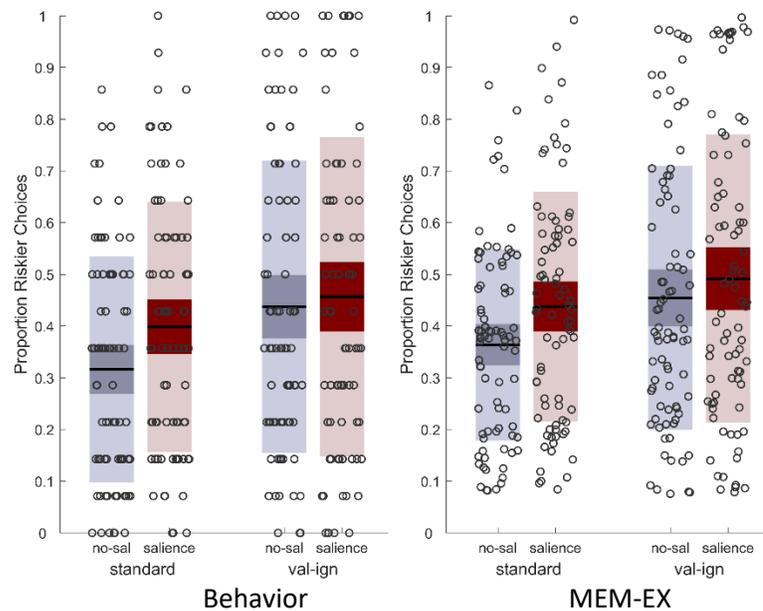


Figure 3. Behavior and MEM-EX predictions across value and salience conditions in Experiments 1 and 2. Each dot represents an individual mean proportion of choices in favor of the riskier alternative. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and 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 more specific details of how the model was fit, 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 effects⁹.

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-

⁹ Qualitative modeling results for Experiments 1-5 are based on fitting each model to all responses 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 and participant. Such an explanation makes intuitive sense, as participants in the Standard condition were allowed to immediately record value information in memory, 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 2.40% of samples ($Mdn_{\lambda_{\text{standard}}} = .024$, $SD_{\lambda_{\text{standard}}} = .130$). 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 assign values to each outcome in memory. 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 3.39% of samples ($Mdn_{\lambda_{\text{val-ign}}} = .034$, $SD_{\lambda_{\text{val-ign}}} = .129$). 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 five outcomes from memory for each alternative ($Mdn_{\kappa} = 5.00$, $SD_{\kappa} = 2.543$). According to MEM-EX's memory-priming

mechanism, in the Saliency condition this sampling was done in a biased manner¹⁰. On average, participants were approximately 19% more likely to sample salient outcomes from memory ($Mdn_{\zeta} = .191$, $SD_{\zeta} = 1.731$).

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 likely to be retrieved from memory during the decision process. 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 error 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 therefore more likely to be sampled from memory when making a choice. Since we highlighted the rare

¹⁰ MEM-EX assumed unbiased sampling in the No-Saliency condition because no outcomes were highlighted.

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 can be explained by a systematic 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 build in the constraint that each individual has an overarching risk bias that is constant across Standard and Value-Ignorance conditions, meaning that only the value-assignment error mechanism can now explain 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 20 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 ($M_{standard} = .40$, $M_{val-ign} = .45$, $SD_{standard} = .18$, $SD_{val-ign} = .24$, $g = .25$). This was also true at the individual level, with 25 participants (60%) making riskier choices under value-ignorance and only 13 participants (31%) making riskier choices in the Standard condition.

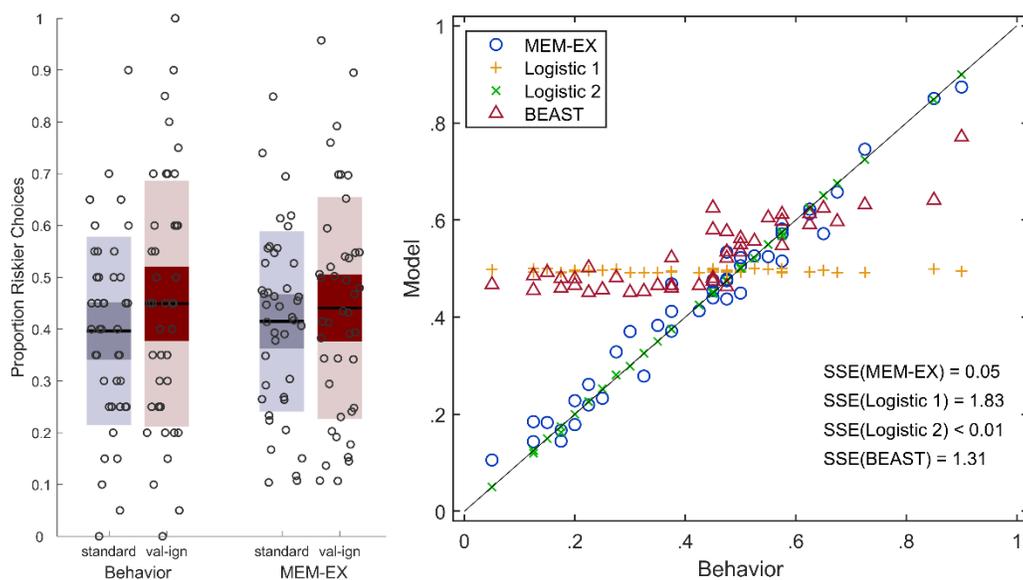


Figure 4. Behavior and model predictions in Experiment 3. The left panel shows mean proportion of choices in favor of the riskier alternative in each condition. Each dot represents an individual. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations. The right panel compares observed and predicted individual mean risky choice proportions for each model. SSE = sum of squared errors.

3.2.2. Modeling Analysis

Our modeling analysis revealed substantial individual differences with respect to the value-ignorance effect. We considered three versions of MEM-EX – each representing a different hypothesis regarding the value-ignorance effects. A large group of participants (45%) were best fit by one of these versions (see Model Comparison for details). The largest group of these (26%) 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 next largest group of individuals (14%), a version of MEM-EX with separate λ parameters (MEM-EX_{base}, see Table 1) gave the best overall account. We also tested an alternative version of MEM-EX (MEM-EX_{alt}, see Table 1), in which we removed the value-assignment error mechanism, and replaced it with the memory confusion mechanism. This model estimated separate ϕ parameters for each condition to represent the possibility that value-ignorance led to additional memory confusion errors. MEM-EX_{alt} performed poorly in our model comparison, giving the best fit to only 2 individuals (5%).

Focusing on MEM-EX_{base} – which assumes separate value-assignment error rates – Figure 4 shows that it produces the same pattern seen in the behavioral data: riskier choices under value ignorance¹¹. 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. This model produced choice proportions that match the pattern seen in participants, predicting that 23 participants

¹¹ These results are based on fitting MEM-EX_{base} to all participants, regardless of which model version produced the best results in cross-validation.

(55%) made riskier choices under value-ignorance and only 19 participants (45%) made riskier choices in the Standard condition.

Figure 5 reveals an interesting pattern relating individual differences in value-assignment error rates ($\lambda_{val-ign} - \lambda_{standard}$) to individual differences in the value-ignorance effect sizes (Value-Ignorance – Standard). As noted above, we can see considerable variability in observed behavior, with 25 participants (60%) showing the typical value-ignorance effect (i.e. riskier choices under value-ignorance), but another 13 (31%) displaying the opposite pattern (i.e. safer choices under value-ignorance), and 4 (10%) showing no effect. Figure 5 also reveals a similar distribution of parameter differences, with 23 individuals (55%) having a higher error rate under value-ignorance, compared to 19 (45%) with a lower error rate under value-ignorance. Most importantly, parameter differences correlated significantly with behavior ($r = .560, p = .001$), demonstrating that MEM-EX provides a new framework for understanding of individual differences in terms of differences in cognitive mechanisms.

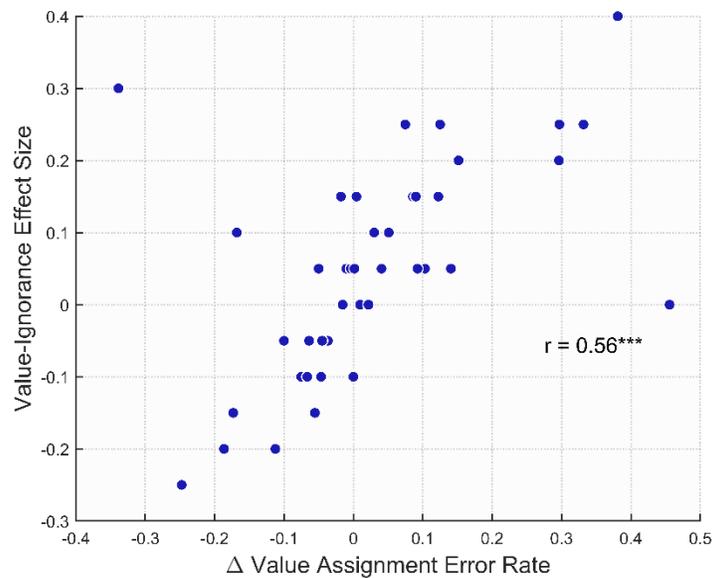


Figure 5. The relationship between value-ignorance effect size – defined as the difference in risky choice proportion – and the difference in value-assignment error rate parameters – defined as $\lambda_{val-ign} - \lambda_{standard}$ – in Experiment 3. Each dot represents an individual. *** $p < .001$.

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 expand our approach and use MEM-EX to examine a variable commonly found to affect DfE: outcome order.

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 because – although a recency bias makes a high-level claim about how events experienced more recently have a disproportionate influence on choice compared to those experienced earlier – such an explanation does not specify how that influence manifests. 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/a264x¹²>).

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 ($M_{standard} = .41$, $M_{val-ign} = .47$, $SD_{val-ign} = .23$, $SD_{standard} = .21$, $g = .30$).

¹² In the OSF preregistration Experiment 4 is labeled as *Experiment 1*. See the General Discussion for more information about the second study, labeled *Experiment 2*, which we do not report here.

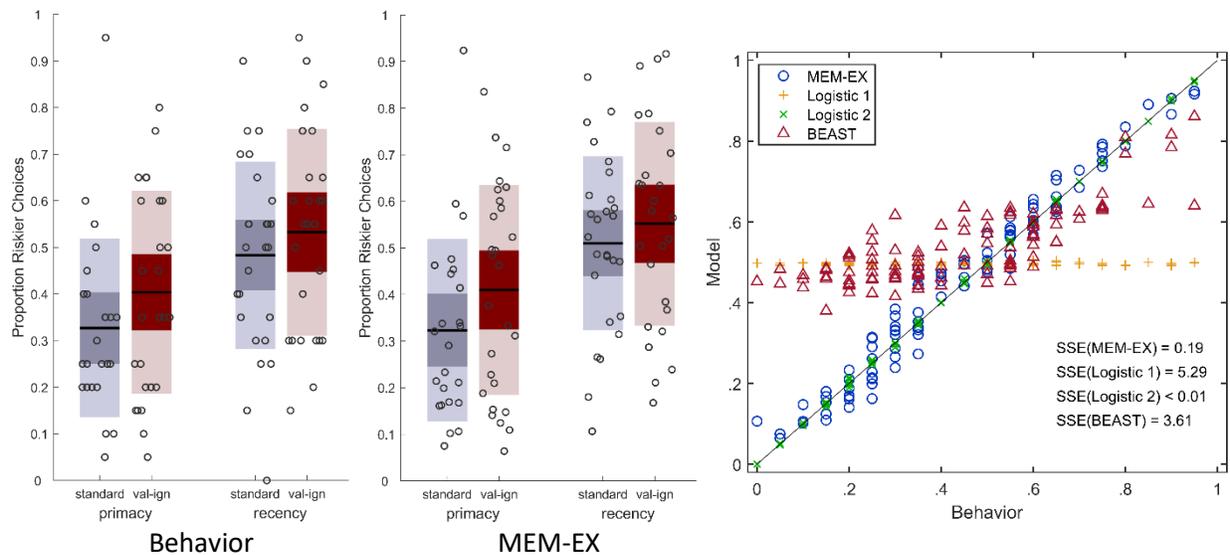


Figure 6. Behavior and model predictions in Experiment 4. The left panel shows mean proportion of choices in favor of the riskier alternative in each condition. Each dot represents an individual. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations. The right panel compares observed and predicted individual mean risky choice proportions for each model. SSE = sum of squared errors.

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. Here, 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 approximately 9.88% of exemplars after each sample ($Mdn_{\psi} = .099$, $SD_{\psi} = .156$).

Once again, MEM-EX captures the observed value-ignorance effect, with higher value-assignment error rates under value-ignorance ($Mdn_{\lambda \text{ val-ign}} = .073$, $SD_{\lambda \text{ val-ign}} = .110$) than in the Standard condition ($Mdn_{\lambda \text{ standard}} = .033$, $SD_{\lambda \text{ standard}} = .103$).

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.

Despite this, there remain alternative explanations for the observed order effect. For example, if individuals believed that the likely set of outcomes changed over time – which conflicts with our cover story, but is conceivable given the order manipulation – they might base their decision on only the most recent outcomes (Navarro, Newell, & Schulze, 2016). A strong version of this might resemble MEM-EX's memory-sampling error mechanism, with the modification that individuals base their decisions on only the final κ items from each alternative, ignoring all earlier information. A more flexible version could add to this something like the memory priming mechanism, with items' retrieval probability being an increasing function of recency (rather than salience). Each of these hypotheses could be tested within the MEM-EX framework, thereby providing a common set of auxiliary assumptions and a level playing field for future comparisons. We therefore acknowledge the possibility of alternative memory-based accounts of order effects, and we think that MEM-EX can play an important role in future investigations of these.

5. Experiment 5

In Experiment 5, we investigate how the number of experienced outcomes affects DfE. Experiment 4 found clear evidence that front/back loading samples produced behavior that could be

explained via memory-confusion errors. Here we explored the intuition that such errors and the retroactive interference that accompany them might be more likely as the number of items in memory increases. MEM-EX makes this basic prediction; 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. We note that although this basic prediction that larger samples lead to more opportunities for memory confusion errors might appear to contradict the law of large numbers, our exhaustive-sampling design ensures that samples are always representative, thus there is no convergence to a more accurate representation with additional sampling.

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

Figure 7 shows that there was no substantial difference in choices across size conditions ($M_{small} = .42$, $M_{large} = .41$, $SD_{small} = .23$, $SD_{large} = .23$, $g = .05$). In contrast, the value-ignorance effect replicated in Experiment 5 with value-ignorance again producing riskier choices ($M_{standard} = .37$, $M_{val-ign} = .46$, $SD_{standard} = .18$, $SD_{val-ign} = .25$, $g = .40$). To analyze the impact of sample size more directly 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.

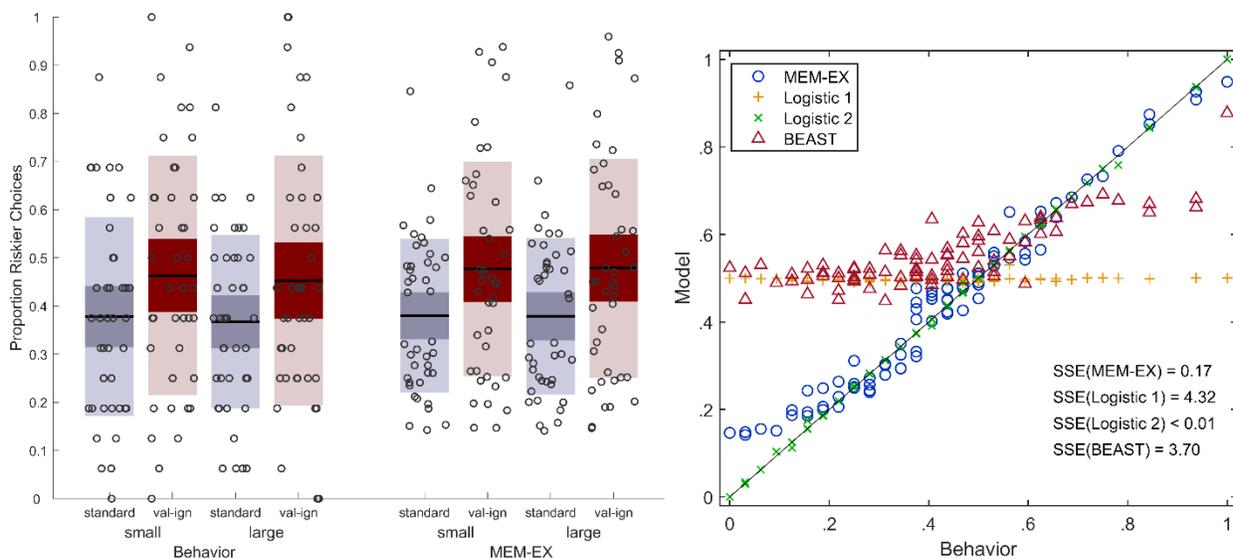


Figure 7. Behavior and model predictions in Experiment 5. The left panel shows mean proportion of choices in favor of the riskier alternative in each condition. Each dot represents an individual. Group mean values are indicated by the solid lines. Dark bands indicate 95% confidence intervals, and light bands indicate standard deviations. The right panel compares observed and predicted individual mean risky choice proportions for each model. SSE = sum of squared errors.

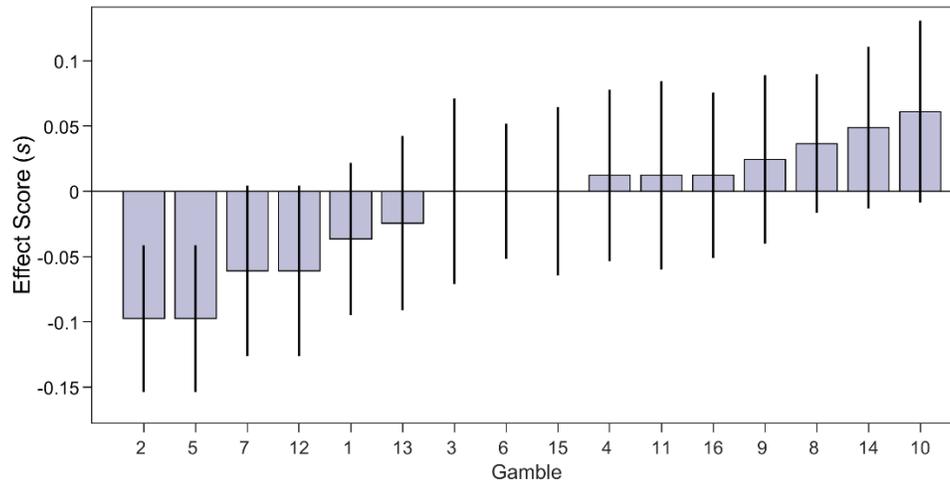


Figure 8. Mean sample size effect scores in Experiment 5. Error bars indicate standard errors.

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, it is value-assignment errors that account for differences across value conditions, with participants making errors an average of 10% of time under value-ignorance, compared with 7% of the time in the Standard condition.

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/floor effects. Designing trials where one alternative is significantly more attractive also runs the risk of encouraging new decision strategies.

The lack of evidence for memory confusion further suggests the need for more research on order effects. Experiment 4 used a strong manipulation and found large effects and strong support for memory confusion, whereas Experiment 5 used no direct order manipulation and consequently yielded no clear evidence for the effects or mechanism. Although this seems a sensible pattern of results, and is consistent with the notion that the mechanisms used should depend on the demands of the experiment (i.e., its design), it nonetheless emphasizes the importance of careful investigation into the factors governing order effects and the mechanisms that produces them. This result also points to potential alternative explanations for the behavior observed in Experiment 4. Perhaps order effects did not result from an interference mechanism that overwrites older items during encoding, but rather a mechanism affecting memory retrieval. For example, a recency effect may indicate 'memory decay', whereby older items have a lower probability of being sampled from memory when making a choice. Similarly, individuals may employ a 'moving window' of attention that prevents old items from being retrieved at all. Such possibilities provide ample fodder for future model-based investigations into the cognitive mechanisms of DfE.

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 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 two alternatives. The *Logistic 1* and *Logistic 2* models are built on the basic principle of expected value maximization and risk bias, and provide useful baselines of performance for any new model to achieve. The *BEAST* shares the baselines' sensitivity to expected value, but adds the assumption that individuals also sample information from memory to make decisions.

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. MEM-EX_{alt} is an alternative model that omits value-assignment errors, and replaces them with memory confusions whose rates are free to vary across Standard and Value-Ignorance conditions. In Experiment 3, we compare these three versions of MEM-EX to test if error rates differed across conditions and whether our results can distinguish between the effects of value-assignment errors and memory confusions. 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.*

Model	Cognitive Mechanisms			
	Value-assignment error (λ)	Memory-sampling error (κ)	Memory Priming (ζ)	Memory Confusion (ϕ)
MEM-EX _{base}	x	x		
MEM-EX _{null}	x	x		
MEM-EX _{alt}		x		x
MEM-EX _{prime}	x	x	x	
MEM-EX _{conf}	x	x		x

6.2. Baseline Models

We built two baseline models for gauging the performance of MEM-EX. Each uses a logistic choice rule to predict behavior according to a simple and intuitive strategy. The first model, *Logistic 1*, represents the intuition that participants choose so as to maximize their expected earnings. According to the model, the likelihood of choosing the riskier alternative was:

$$\frac{1}{1 + e^{-\psi\delta}} \quad (1)$$

where δ is the difference in expected value favoring the riskier alternative (i.e. $EV_{\text{riskier}} - EV_{\text{safer}}$) and ψ is a free parameter controlling each individual's sensitivity to expected value differences. Note that in our designs participants always experienced representative sequences of outcomes that matched the objective probabilities of rewards. Consequently, the expected value for alternative X, was simply:

$$EV_X = v \times \pi. \quad (2)$$

The second baseline model, *Logistic 2*, was identical to Logistic 1, but replaces the choice rule in Equation 1 with:

$$\frac{1}{1 + e^{-\psi(\delta-b)}} \quad (3)$$

where b is a free parameter representing each individual's bias in favor of the riskier alternative. Using these models we establish the degree to which participants' behavior can be accounted for with EV maximization alone (Logistic 1) or with the additional assumption of individual differences in risk attitude.

6.3. BEAST

The *best estimate and simulation tools*, or *BEAST*, is a model developed as part of a choice prediction competition reported in Erev et al. (2017). We chose to include it in our model comparisons because Erev et al. (2017) considered dozens of candidate models and found BEAST to be among the very best, performing well across a range of DfE domains and contexts. In fact, all 12 of the highest ranked submissions to the prediction competition were variants of BEAST. This success in predicting behavior makes it an appropriate benchmark for establishing the additional explanatory value of specifying memory mechanisms for particular effects within the MEM-EX framework. Including BEAST in our comparison therefore provides a means for evaluating the novel contribution of MEM-EX relative to an accepted and well-regarded alternative model.

BEAST is built on the assumption that choice is sensitive to expected value, but also relies on four behavioral tendencies: (a) bias toward equal weighting of possible outcomes, (b) minimization of the probability of regret, (c) sensitivity to the payoff sign, and (d) pessimism. These tendencies are thought to reflect people's use of simple heuristic strategies that might be well-adapted to particular scenarios (see Gigerenzer & Todd, 1999). However, these tools are not thought to determine choice – as is commonly assumed with heuristic models of choice – but instead provide additional and complementary estimates of the likely outcomes of each choice alternative.

For our implementation of BEAST, we began with code provided on a website reporting preliminary results from a follow-up choice prediction tournament (<https://cpc-18.com>). BEAST

represents the intuition that individuals are sensitive to expected value, but may also rely on additional processes that involving sampling from memory. In the case of a binary choice between alternatives A and B, the model predicts that individuals choose A if:

$$\delta + (ST_A - ST_B) + e > 0, \quad (4)$$

where δ is the same as in Equation 1 for Logistic 1, ST_x is the average of k outcomes, where each outcome is generated by using one of three¹³ sampling tools. For the *unbiased* sampling tool, the individual first draws a *luck level* from a uniform distribution between 0 and 1 to be used for both alternatives. This value is used to sample an outcome from the cumulative payoff distribution for each option. In our design, a luck level above $p(\text{reward})$ would result in sampling the reward for that alternative, otherwise the sample would be a 0. The *uniform* sampling tool also draws a random number from $U(0,1)$, but uses it to sample uniformly from each alternative's payoff distribution. Finally, the *sign* sampling tool operates similar to the unbiased tool, except that positive outcomes are replaced with R and negative outcomes are replaced with $-R$, where R is the payoff range across both alternatives, or:

$$R = \max(\text{MAX}_A, \text{MAX}_B) - \min(\text{MIN}_A, \text{MIN}_B). \quad (5)$$

Here, MAX_x and MIN_x are the maximum and minimum payoffs for X, respectively.

With a probability γ , BEAST uses one of the biased sampling tools, with uniform sampling occurring 2/3 of the time and sign sampling occurring 1/3 of the time.¹⁴ The error term, e , is drawn from a normal distribution with a mean of 0 and standard deviation σ . Erev, et al. (2017) focused their analyses on predicting mean choice proportions aggregated across individuals, and therefore treated BEAST as a model of group behavior. For our purposes, we adapt BEAST to model choices at an individual level, and for each participant we estimate three free parameters to account for individual

¹³ Our implementation of BEAST did not include the *contingent pessimism* sampling tool because all alternatives shared the same minimum value 0.

¹⁴ The imbalance between *uniform* and *sign* sampling results from our design always presenting alternatives with the same minimum value, 0. Consequently, *contingent pessimism* sampling – which occurs 1/3 of time – defaults to uniform sampling. See Erev, et al. (2017) for additional details.

differences. The sample size k was constrained to integers between 1 and 8, inclusive. The probability of biased sampling, γ , had a value between 0 and 1. And the standard deviation of the error distribution, σ , had a value between 0 and 10.

6.4. 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. We simulated the models 500 times and averaged the results to compute predicted choice probabilities for each trial. 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 parameter values 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. Matlab code for MEM-EX, BEAST, and baseline models can be found on OSF (<https://osf.io/wjgqd>).

6.5. 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 typically yielded the best performing model, with the exception of Experiment 4 where Logistic 1 was best. BEAST 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.

Model	Experiment			
	1 & 2	3	4	5
MEM-EX _{base}	—	6 (14.29%)	—	41 (50.00%)
MEM-EX _{null}	—	11 (26.19%)	—	—
MEM-EX _{alt}	—	2 (4.76%)	—	—
MEM-EX _{prime}	149 (45.71%)	—	—	—
MEM-EX _{conf}	—	—	37 (35.58%)	7 (8.54%)
Logistic 1	121 (37.12%)	12 (28.57%)	51 (49.04%)	19 (23.17%)
Logistic 2	52 (15.95%)	9 (21.43%)	12 (11.54%)	11 (13.41%)
BEAST	4 (1.23%)	2 (7.14%)	4 (3.85%)	4 (4.88%)

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 Figure 9)¹⁵.

¹⁵ The weight for Model i in the set of j models was calculated as $e^{LL_i} / \sum_j e^{LL_j}$.

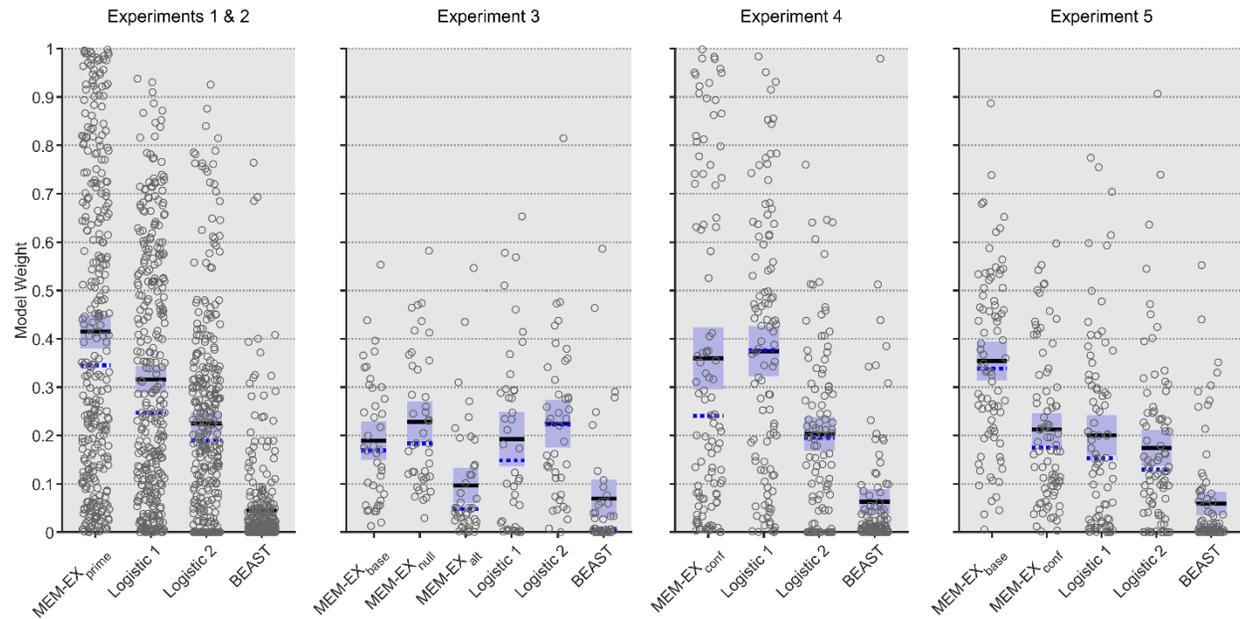


Figure 9. Cross-validation model weights for Experiments 1-5. Each dot represents an individual. Group mean and median values are indicated by the solid and dotted lines, respectively. Shaded bands indicate 95% confidence intervals.

6.5.1. Experiments 1 & 2

For the dataset combining Experiments 1 and 2 we tested three 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 best performing model for nearly half of individuals (46%). 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. Logistic 1 was the next most popular, providing the best account for 37% of individuals. This result implies that – although the model almost completely fails to reproduce the observed distribution of mean choice proportions across individuals (see Figure 2) – its strategy of focusing on expected value struck a good balance between flexibility and fit in cross-validation. The addition of a risk bias led to 16% of individuals being best fit by Logistic 2. BEAST performed poorly, and was the preferred model for only 1% of individuals. Together,

these findings indicate that MEM-EX can predict choices more accurately than a reasonable alternative model (BEAST), while MEM-EX also provides new insights into the cognitive mechanisms underlying behavior.

6.5.2. Experiment 3

In Experiment 3 we found individual differences in the within-subjects value-ignorance effect. To better understand the influence of value-ignorance on the decision process we tested three models within the MEM-EX framework. MEM-EX_{null} and MEM-EX_{base} hypothesize that participants made value-assignment errors, while MEM-EX_{alt} tests an alternative hypothesis that memory confusions explain behavior. Together, these three models gave the best account for a total of 19 individuals (45%) in Experiment 3. 26% of participants were best characterized by MEM-EX_{null}, which used a single value-assignment error rate across conditions. Six of the remaining participants (14%) were best captured by MEM-EX_{base}, which tended to estimate higher value-assignment error rates in the Value-Ignorance condition. The poor performance of MEM-EX_{alt} – which gave the best account of only 2 individuals (5%) – further strengthens the case for value-assignment errors. This result also provides some assurance that value-assignment and memory confusion are indeed distinct mechanisms, differentially affected by experimental manipulations, that cannot easily mimic each other. For additional information about mimicry within the MEM-EX framework see Appendix B, where we present results from model recovery analyses of MEM-EX in Experiments 1 and 2 and Experiment 3.

Fifty per-cent of participants were best fit by baseline models, again indicating the predictive utility of such a simple model. BEAST struggled again, with 7% 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.

6.5.3. Experiment 4

Logistic 1 faired best under cross-validation in Experiment 4, and was the preferred model for 49% of individuals. MEM-EX_{conf} was the best performing model for 36% of individuals in Experiment 4, giving credence to the importance of its memory confusion mechanism for explaining the impact of outcome order on choice. BEAST performed poorly, with 4% of participants best fit.

6.5.4. 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. Half of participants (50%) were best characterized by MEM-EX_{base}'s value-assignment error and memory-sampling error mechanisms. A further 7 individuals (9%) supported the addition of MEM-EX_{conf}'s memory confusion mechanism, which we expected to interact with sample size. Logistic 1 and Logistic 2 gave the best account of choices for 23% and 13% of participants, respectively. BEAST again struggled, and was the preferred model for only 5% of individuals.

6.6. Discussion

These results were remarkably consistent across all experiments, and – together with the qualitative results presented in earlier sections – provide converging evidence for the usefulness of 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. Further evidence of the insights we can draw from the MEM-EX framework are explored in Figure A2 of the Appendix. Here we show how difficult it is to find clear relationships between risky choice rates across gambles and variables such as the probability of rewards or the magnitude of points.

A rather surprising aspect of our model comparison exercise is the failure of BEAST to capture behavior. Given the model's excellent performance in previous competitions (e.g., Erev et al., 2017), we had hoped it would provide a reasonable benchmark for comparison. We suspect that its failure is to do with the level of explanation/prediction that we are attempting to achieve in our modeling exercise compared to that pursued in previous prediction tournaments. In those tournaments, the key goal is to model group means; as such, whether variability in parameter estimates comes from within or between-subjects is of no concern. For our comparison we created an individual version of the model where we estimate a set of parameters for *each individual*. In doing so we may have reduced the power of BEAST to flexibly mix different mechanisms in the ways that helps it explain group-level variability in choice. Indeed, the risks and potential problems associated with modeling aggregate data are well-documented (see Birnbaum, 2011; Regenwetter & Robinson, 2017; Spektor & Wulff, 2021; Wulff & van den Bos, 2018). It is also worth noting that BEAST was developed to capture a wide range of benchmark phenomena within a single model fit. The data we used here did not contain such patterns, leaving the model relatively underconstrained, and so unable to capture some of the data patterns we observed. For example, *pessimism sampling* – where decision makers are assumed to sample the worst outcome for each option – is an important tool for producing risk aversion in BEAST because riskier alternatives tend to have worse minimum outcomes compared to safer alternatives (at least in laboratory experiments). However, in our experiments BEAST never used this tool because all alternatives had the same minimum outcome (see Erev, et al., 2017 for additional details). As a result, the model could not account for risk averse individuals in Experiments 1-5 (see Figures 2, 4, 6, and 7). To sum up, our goal here is to shed light on plausible cognitive mechanisms that can account for individual choices within a relatively narrow range of experimental designs – given that criterion, it is perhaps not so surprising that BEAST fails to live up to its apparent potential.

A final noteworthy finding from the comparison was the good performance of Logistic 1 in cross-validation for Experiment 4. That nearly half of individuals were best fit by this model suggests that EV was a good predictor of overall behavior. That said, the poor performance of Logistic 2 and BEAST – which also include an EV component – may indicate that Logistic 1's simplicity is the main driver of its success. That is, although MEM-EX was uniquely capable of explaining the effects of value-ignorance and outcome order – while also successfully modeling the observed individual differences in choice proportions – its additional complexity relative to the baseline models was not merited for many participants.

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.

As noted in the introduction we remain agnostic as to whether value-assignment errors occur during encoding or as part of memory retrieval. We were able to rule out one possible means by which value-assignment errors occur; our model-based analysis of Experiment 3 indicated that the value-assignment-error process cannot be replaced with the memory confusion mechanism. Regarding what explanations for value-assignment errors remain, one possibility is that while retrieving exemplars from memory people may unintentionally edit the values of some items. It is an interesting and open question whether these edits occur in working memory and are limited to the current decision or are made permanent in the long-term memory store. MEM-EX provides a useful framework for investigating these possibilities in future studies. For current purposes it is how these outcomes are represented at the time of choice that is of key interest.

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* – without specifying a particular mechanism for producing order effects – we can now articulate and test a formal psychological theory. There are, however, alternative ways to conceptualize the observed order effects. Perhaps our manipulation led participants to believe that the probability of receiving a reward changed over time; increasing in the recency condition and decreasing in primacy. This is consistent with the memory confusion mechanism, and would even provide it a rational justification. That is, rather than view these confusions as unintended errors, we can think of them as adaptations to dynamic environments where older observations become outdated. Future research will investigate alternative mechanisms for producing similar effects within the MEM-EX framework, such as recency-weighted retrieval using a mechanism similar to memory priming.

7.3. 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 et al., 2015)? What cognitive mechanisms best explain learning and adaptation in dynamic environments (Hotaling, Fakhari, & Busemeyer, 2015; Hotaling & Kellen, in press), where payoffs change over time or as a consequence of the decision maker's actions (Hotaling, Navarro, & Newell, 2018, 2020; Navarro et al., 2016)? MEM-EX represents an important 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 several 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.¹⁶ 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

¹⁶ 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.

A related direction in which to expand the applications of MEM-EX would be to explore the impact of contextual variations on the recruitment of different memory mechanisms. Research into value-based decision making demonstrates that the familiarity of current and prior decision-context influences choices (e.g. Duncan & Shohamy, 2016). Specifically, participants are more likely to choose options that they have selected before when those options are subsequently presented in the same context (e.g., the background image on the choice screen is the same), than when the context changes. This effect holds when the immediately preceding trial (rather than the current one) has the prior familiar context. These results suggest a reliance on episodic memory, and perhaps even a type of ‘retrieval state’ that facilitates recollection of similar episodic memories. In our experiments, we excluded any contextual cues by virtue of presenting different colored pairs of urns for each gamble problem and keeping the background image constant. Nonetheless, it would be straightforward in future experiments to manipulate context and examine its impact on patterns of choice. One might expect that if participants could recruit additional memory cues (e.g., a black background could indicate urns with low probability, high rewards and a white background the opposite), then we would see fewer memory-sampling and memory-confusion errors. (See also Murty, FeldmanHall, Hunter, Phelps, & Davachi, 2016 for related ideas about the impact of context on value-based decision making.)

Our work with MEM-EX also suggests that a memory-based perspective on decision-making might also provide key insights into how memory itself works. For example, while encouraging, the studies herein do not provide a clear adjudication between the specific explanations for observed recency effects. On the one hand, MEM-EX appeals to the effect of retroactive interference on stored memories, while IBL would explain the same result using the idea of memory decay. Deciding between those accounts in memory studies has been a source of ongoing controversy (e.g. Berman, Jonides, &

Lewis, 2009; Oberauer & Lewandowsky, 2008; Portrat, Barrouillet, & Camos, 2008), and it is possible that DfE tasks may provide further avenue for investigation. Similarly, debate lingers over why primacy has such large effects on memory (Kahana, 2012). While the experiments here did not warrant the explicit modeling of primacy, future tests of MEM-EX could examine whether the manipulation of the initial samples could affect behavior in systematic ways.

Another potentially fruitful avenue for future investigations involves comparing the performance and predictions of MEM-EX to those of models built on principles of reinforcement learning (RL; Sutton & Barto, 1998). The RL framework is a popular method for studying DfE because it offers general and often simple mechanisms for modeling experience-based choice across many contexts (see Hertwig, Barron, Weber, & Erev, 2006). For our purposes, RL models – which excel at producing complex behavior by incrementally updating relatively limited knowledge representations – offer an interesting alternative to MEM-EX’s analogical representations of experienced outcomes.

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_{base} vs. MEM-EX_{conf?}). In sum, we find that MEM-EX offers a framework for understanding experience-based choice in terms of cognitive mechanisms, while also helping to motivate future empirical investigations and theoretical developments.

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

- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6), 396-408. doi:10.1111/j.1467-9280.1991.tb00174.x
- Berman, M. G., Jonides, J., & Lewis, R. L. (2009). In search of decay in verbal short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(2), 317-333. doi:10.1037/a0014873
- Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review*, 118(4), 675-683. doi:10.1037/a0023852
- 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.
- Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, 108(2), 452-478. doi:10.1037/0033-295x.108.2.452
- Duncan, K. D., & Shohamy, D. (2016). Memory states influence value-based decisions. *Journal of Experimental Psychology: General*, 145(11), 1420-1426. doi:10.1037/xge0000231
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, 124(4), 369-409. doi:10.1037/rev0000062
- 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., Glöckner, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, 36(2), 153-177.
- Erev, I., Plonsky, O., & Roth, Y. (2020). Complacency, panic, and the value of gentle rule enforcement in addressing pandemics. *Nature Human Behaviour*, 4(11), 1095-1097. doi:10.1038/s41562-020-00939-z
- Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, 1(2), 159-161.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- 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.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118(4), 523-551. doi:10.1037/a0024558
- 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. E., Camilleri, A. R., Heathcote, A., Newell, B. R., & Brown, S. D. (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. doi:10.1016/j.cognition.2009.12.009
- Hotaling, J. M., Fakhari, P., & Busemeyer, J. R. (2015). Dynamic decision making. In J. D. Wright (Ed.), *International Encyclopedia of the Social and Behavioral Sciences* (2nd ed., pp. 709-714). Oxford: Elsevier.
- 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., & Kellen, D. (in press). Dynamic decision making: Empirical and theoretical directions. In K. D. Federmeier (Ed.), *Psychology of Learning and Motivation* (Vol. 76).
- 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. (2020). Skilled bandits: Learning to choose in a reactive world. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 47*(6), 879-905. doi:10.1037/xlm0000981
- Kahana, M. J. (2012). *Foundations of human memory*. New York, NY, US: Oxford University Press.
- 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.
- Murty, V. P., FeldmanHall, O., Hunter, L. E., Phelps, E. A., & Davachi, L. (2016). Episodic memories predict adaptive value-based decision-making. *Journal of Experimental Psychology: General, 145*(5), 548-558. doi:10.1037/xge0000158
- 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.

- Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall: Decay, temporal distinctiveness, or interference? *Psychological Review*, *115*(3), 544-576. doi:10.1037/0033-295X.115.3.544
- 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
- Portrat, S., Barrouillet, P., & Camos, V. (2008). Time-related decay or interference-based forgetting in working memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1561-1564. doi:10.1037/a0013356
- 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.
- Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). List-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(2), 163-178. doi:10.1037/0278-7393.16.2.163
- Regenwetter, M., & Robinson, M. M. (2017). The construct-behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*(5), 533-550. doi:10.1037/rev0000067
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1446-1465.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, *4*(2), 145-166. doi:10.3758/BF03209391
- Simon, H. A. (1982). *Models of bounded rationality*. Cambridge, MA: The MIT Press.
- Spektor, M. S., & Wulff, D. U. (2021). Myopia drives reckless behavior in response to over-taxation. *Judgment and Decision Making*, *16*(1), 114-130.
- Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive psychology*, *53*, 1-26.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). Cambridge: MIT press.
- Tulving, E., & Hastie, R. (1972). Inhibition effects of intralist repetition in free recall. *Journal of Experimental Psychology*, *92*(3), 297-304.
- 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.
- Wulff, D. U., & van den Bos, W. (2018). Modeling choices in delay discounting. *Psychological Science*, *29*(11), 1890-1894. doi:10.1177/0956797616664342
- 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.

Appendix A

1. Gable-Wise Effects of Value-Ignorance

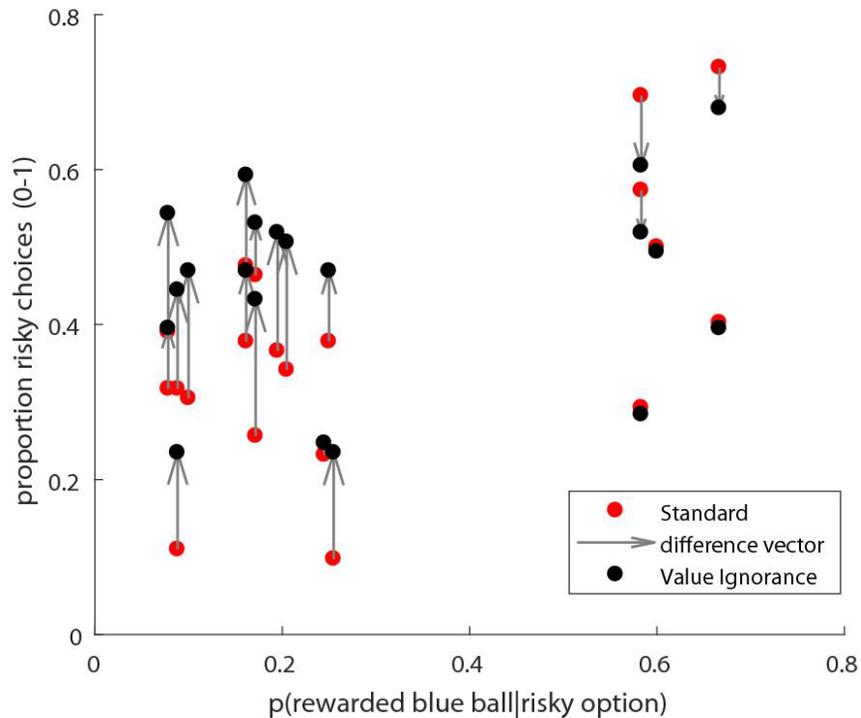


Figure A1. Gamble-wise proportion of risky choices as a function of the probability of drawing a blue ball from the riskier option in Experiments 1 and 2. Each data point comes from a unique gamble in one of the two conditions. In both experiments, there was a wide range of reward probabilities and approximately 80 observations per gamble. The arrows indicate the direction of change in the Value-Ignorance condition relative to the Standard condition. Values on the x-axis have been jittered to improve readability.

2. Gable-Wise Effects of Experimental Variables

To further emphasize the novel contribution of MEM-EX to our understanding of DfE, we analyzed the relationship between several experimental variables and risky choice proportions in Experiment 5. Figure A2 displays our findings and demonstrates how little can be learned by considering each variable in isolation. For each panel, gambles are sorted according to a different experimental variable. The first panel shows a predictable pattern, whereby individuals made riskier choices as the expected value of the riskier alternative increased relative to that of the safer alternative. However, the

remaining five panels display 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 behavior – that appears to be unrelated to gamble variables such as the magnitude or likelihood of rewards – can be succinctly characterized in terms of cognitive mechanisms like value-assignment and memory sampling.

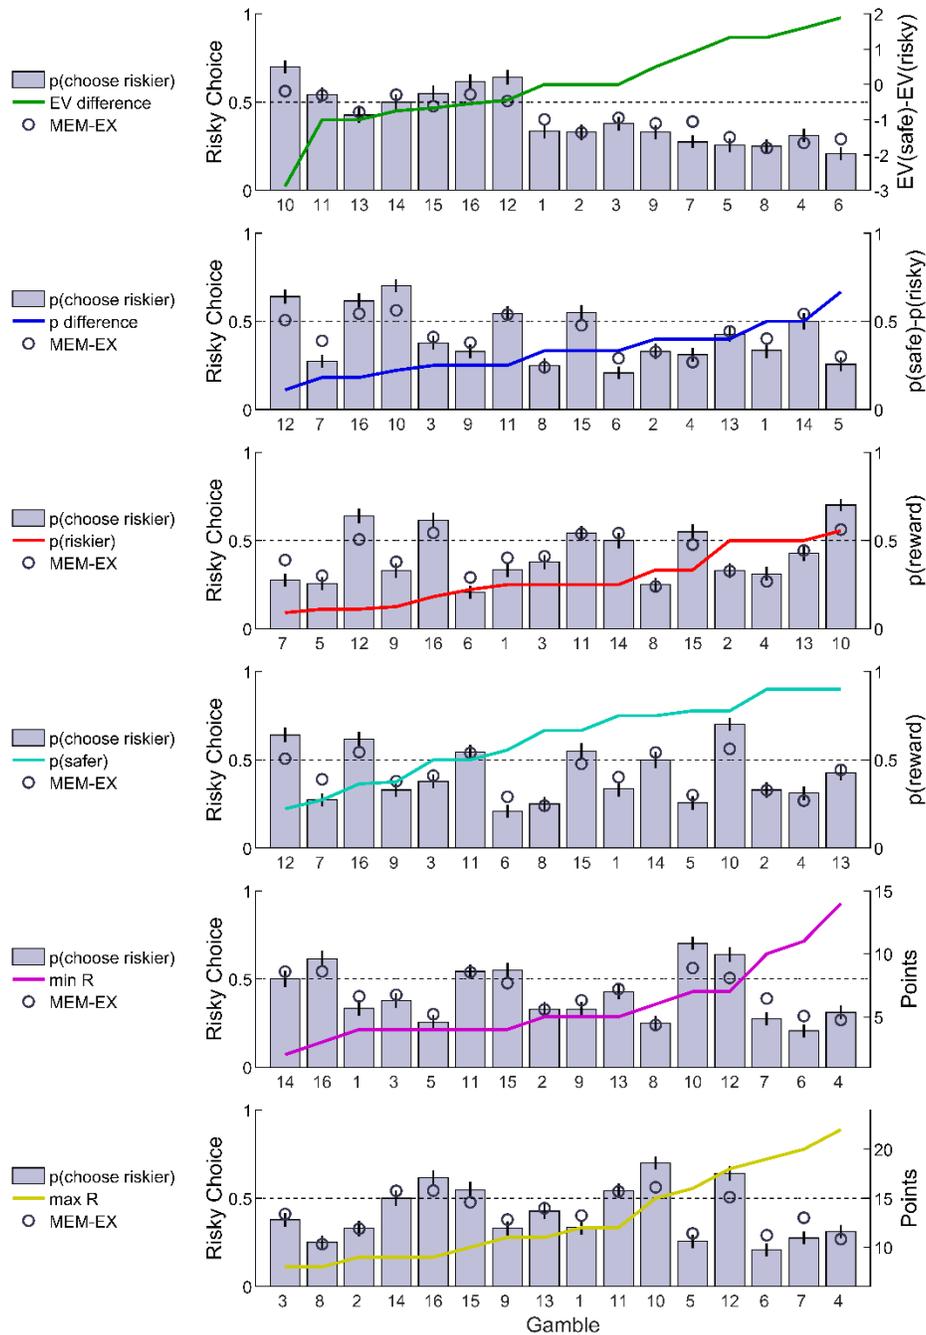


Figure A2. Shaded bars indicate the mean proportion of choices in favor of the riskier alternative in Experiment 5 (left axis). Each row is sorted by the A) difference in expected value across gambles, B) difference in reward probabilities across gambles, C) probability of receiving a reward from safer 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.

Appendix B

1. Model Recovery Analysis

1.1. Experiments 1 and 2

1.1.1. Method

For Experiments 1 and 2, we used MEM-EX_{prime} to investigate model recovery. We began with the best-fitting parameter values for each of the 326 participants, and used these to simulate behavior in our experimental design. We repeated this process 10 times for each participant – i.e. each set a parameters was run through the experiment 10 times – to produce a total of 3,260 simulated agents. We then treated each agent as a participant and fit MEM-EX_{prime} to its simulated choices. The analysis below compares the original parameter values from the 326 participants to the recovered values from the 3,260 simulated agents.

1.1.2. Results & Discussion

Our results indicate a partial success in recovering the key insights provided by MEM-EX. Table B1 and Figure B1 summarize the results. Correlations between original and recovered parameters were moderate for λ , ζ , and κ , while the correlation for β was strong. All correlations were highly significant ($p < .001$). We were particularly interested in recovery of the difference in value-assignment error rates because these were central to MEM-EX's explanation of the value-ignorance effect. Originally, error rates were higher in the Value-Ignorance condition ($Mdn_{\lambda_{\text{standard}}} = .024$, $SD_{\lambda_{\text{standard}}} = .130$, $Mdn_{\lambda_{\text{val-ign}}} = .034$, $SD_{\lambda_{\text{val-ign}}} = .129$). However, recovered parameter values showed a much smaller difference ($Mdn_{\lambda_{\text{standard}}} = .045$, $SD_{\lambda_{\text{standard}}} = .120$, $Mdn_{\lambda_{\text{val-ign}}} = .047$, $SD_{\lambda} = .133$).

Table B1. *Median original and recovered parameters for Experiments 1 and 2. Standard deviations are shown in parentheses. r = correlation coefficient. *** $p < .001$.*

Model Recovery for Experiments 1 and 2

Parameter	Original	Recovered	r
λ	.028 (.129)	.045 (.127)	.584***
β	-1.036 (4.351)	-1.662 (4.539)	.763***
κ	5 (2.543)	5 (2.536)	.485***
ζ	0.191 (1.731)	0.421 (2.049)	.577***

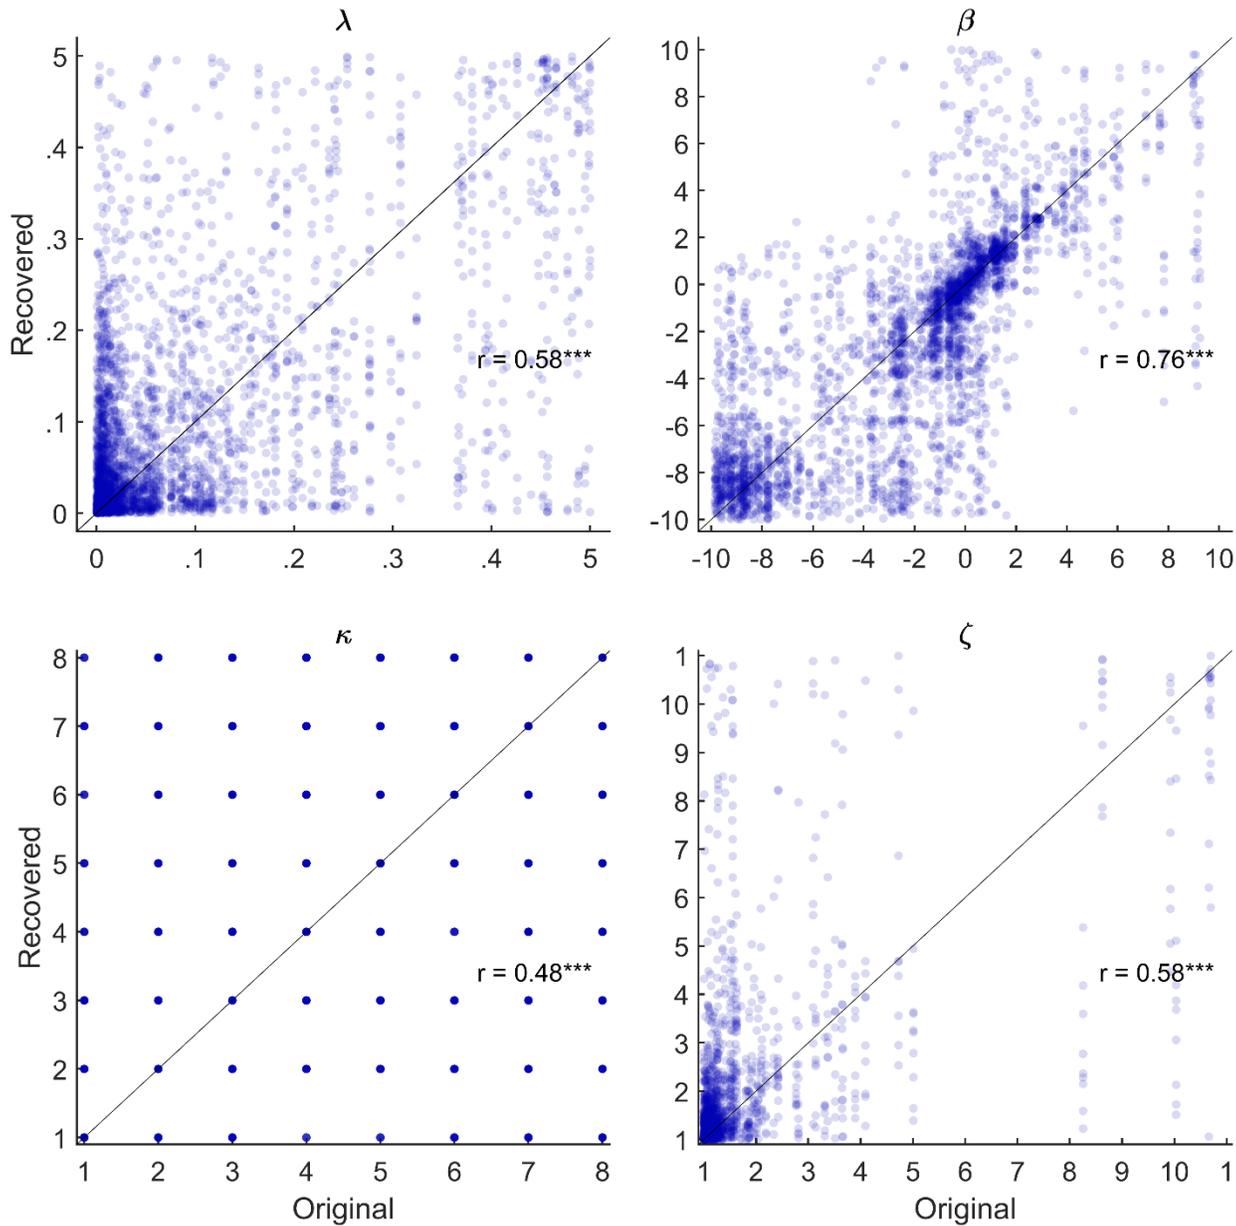


Figure B1. Comparison of original and recovered parameter values for MEM-EX in Experiments 1 and 2.
*** $p < .001$

1.2. Experiment 3

1.2.1. Method

For Experiment 3, we investigated model recovery for MEM-EX_{base} using a method identical to that for Experiments 1 and 2. Best-fitting parameter values for the 42 participants were used to simulate choices for 420 agents.

1.2.2. Results & Discussion

Our model recovery results in Experiment 3 show considerable success. Table B2 and Figure B2 summarize the results. We again found that the correlation between original and recovered parameters was moderate for λ , ζ , and κ , and was strong for β . All correlations were highly significant ($p < .001$). Experiment 3 was focused on the within-subjects value-ignorance effect, so we were especially interested in recovery of within-subject differences in value-assignment error rates. Our original analysis found that 23 individuals (55%) had a higher error rate under value-ignorance, compared to 19 (45%) who had a lower error rate under value-ignorance. With our model recovery when found a very similar pattern, with 240 simulated agents (57%) having a higher error rate under value-ignorance, compared to 179 (43%) who had a lower error rate under value-ignorance (1 agent had equal rates). In sum, we successfully recovered the key insight of MEM-EX in Experiments 3, i.e. that value-ignorance effects result from within-subject differences in value-assignment error rates across conditions.

Table B2. Median original and recovered parameters for Experiment 3. Standard deviations are shown in parentheses. r = correlation coefficient. *** $p < .001$.

Model Recovery for Experiment 3

Parameter	Original	Recovered	r
$\lambda_{standard}$.091 (.108)	.080 (.119)	.535***
$\lambda_{val-ign}$.076 (.146)	.108 (.145)	.662***
β	-2.047 (3.836)	-2.270 (3.957)	.775***
κ	5 (2.709)	5 (2.565)	.693***

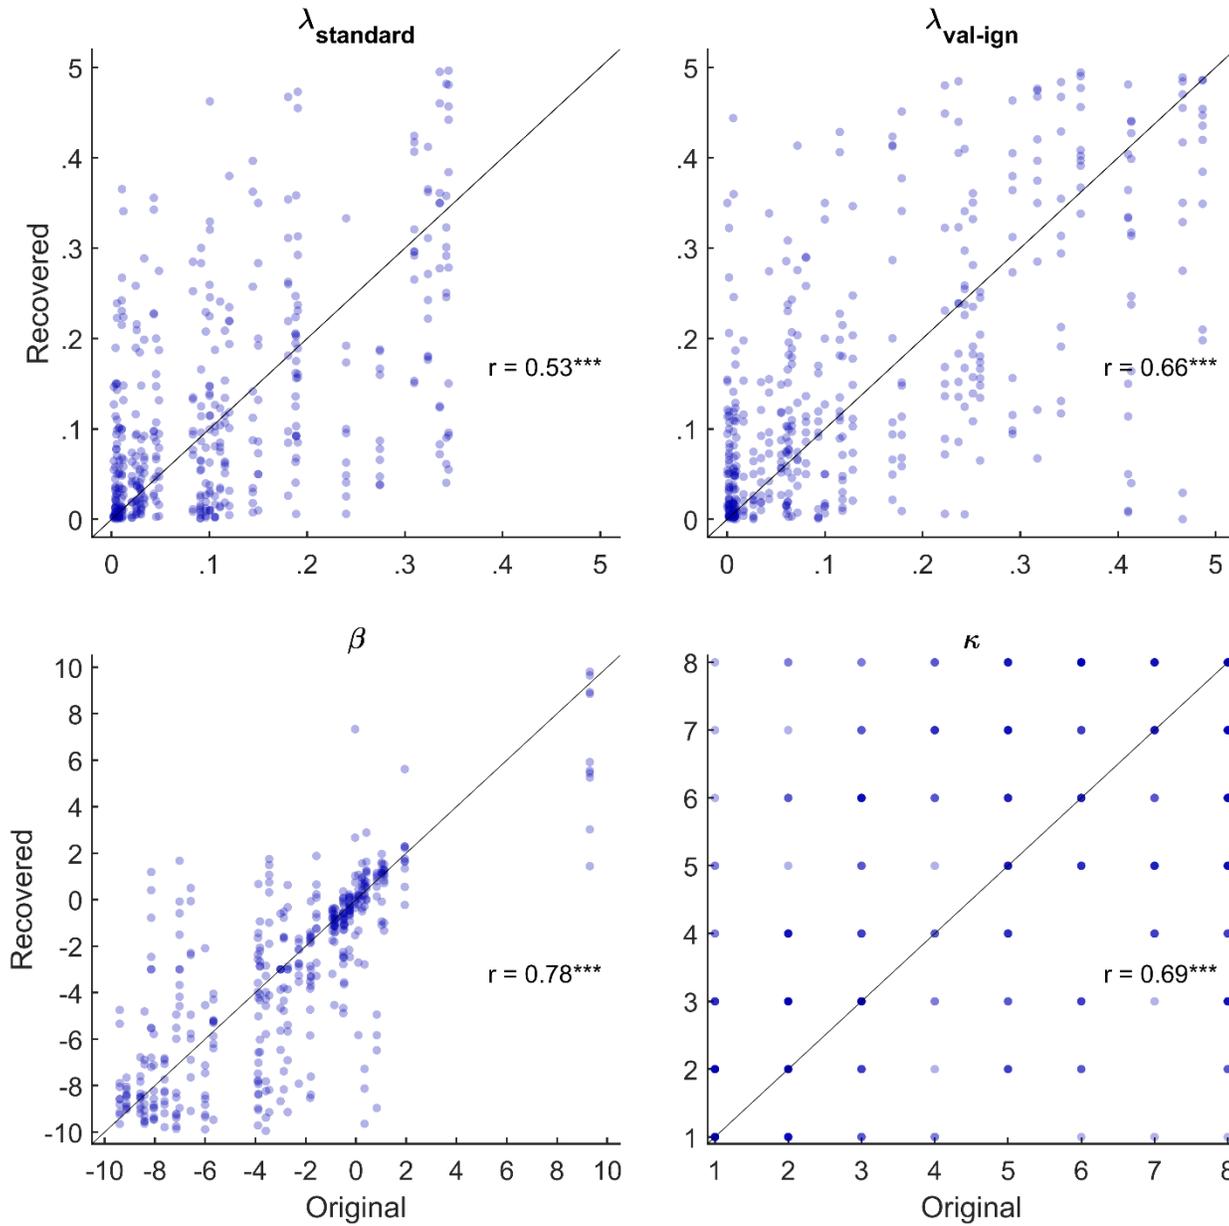


Figure B2. Comparison of original and recovered parameter values for MEM-EX in Experiment 3. $***p < .001$