University of Groningen

Framing From Experience

Gonzalez, Cleotilde; Mehlhorn, Katja

Published in:
Cognitive Science

DOI:
[10.1111/cogs.12268](https://doi.org/10.1111/cogs.12268)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2016

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Gonzalez, C., & Mehlhorn, K. (2016). Framing From Experience: Cognitive Processes and Predictions of Risky Choice. *Cognitive Science*, 40(5), 1163-1191. <https://doi.org/10.1111/cogs.12268>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Framing From Experience: Cognitive Processes and Predictions of Risky Choice

Cleotilde Gonzalez, Katja Mehlhorn

*Dynamic Decision Making Laboratory, Department of Social and Decision Sciences,
Carnegie Mellon University*

Received 8 April 2014; received in revised form 19 April 2015; accepted 20 April 2015

Abstract

A framing bias shows risk aversion in problems framed as “gains” and risk seeking in problems framed as “losses,” even when these are objectively equivalent and probabilities and outcomes values are explicitly provided. We test this framing bias in situations where decision makers rely on their own experience, sampling the problem’s options (safe and risky) and seeing the outcomes before making a choice. In Experiment 1, we replicate the framing bias in description-based decisions and find risk indifference in gains and losses in experience-based decisions. Predictions of an Instance-Based Learning model suggest that objective probabilities as well as the number of samples taken are factors that contribute to the lack of framing effect. We test these two factors in Experiment 2 and find no framing effect when a few samples are taken but when large samples are taken, the framing effect appears regardless of the objective probability values. Implications of behavioral results and cognitive modeling are discussed.

Keywords: Framing effect; Gains and losses; IBLT; Decisions from experience; Decision bias

1. Introduction

We often make decisions relying on explicit information about the probabilities of outcomes. For example, a patient is given probabilities regarding the outcomes of two treatments for lung cancer (e.g., radiation or surgery). In principle, such information should help the patient to maximize the chances of survival. However, after many decades of research in decision sciences, we know that humans suffer from a large set of “cognitive biases” affecting the way that they evaluate probabilities and outcomes from explicit

descriptions. These biases often lead to inaccurate judgments and suboptimal choices (Kahneman, Slovic, & Tversky, 1982; Tversky & Kahneman, 1981).

A well-known bias that demonstrates human “irrationality” is the *framing effect*. A prototypical example of this bias has been shown using the Asian Disease Problem (ADP) from Tversky and Kahneman (1981, 1983, 1986), where participants are presented with a descriptive scenario about an epidemic disease and are then asked to choose between options that are objectively equivalent (i.e., they have identical expected values) but are framed as gains or as losses.

ADP Scenario:

Imagine that the United States is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the programs are as follows:

Positive frame (Gains):

If Program A is adopted, 200 people will be saved

If Program B is adopted, there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved.

Negative frame (Losses):

If Program C is adopted, 400 people will die

If Program D is adopted, there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die.

A decision maker is said to be “risk averse” if she prefers a safe over a risky prospect of equal or higher expected value, and she is said to be “risk seeking” if she prefers a risky over a safe one of equal or higher expected value (Fox & Poldrack, 2009). In the ADP, a majority of respondents under gains prefer program A, the safe option; and a majority of respondents under losses prefer program D, the risky option—despite these programs being equivalent in terms of their expected value: Of the 600 people, 200 are expected to be saved and 400 are expected to die. In Tversky and Kahneman (1981), for example, 72% of the respondents selected the *safe* option when the problem was framed as gains, but 78% selected the *risky* option when the problem was framed as losses.

The framing effect in risky choice has generally been explained by Prospect Theory’s value and probability functions (Tversky & Kahneman, 1981). The value function suggests that for gains and losses, the subjective difference between two small outcomes is greater than the subjective difference between two large values (a phenomenon called *diminishing sensitivity*), and that the displeasure associated with a loss is generally greater than the pleasure associated with a gain of the same magnitude (a phenomenon called *loss aversion*). The subjective probability function suggests an overweighting of small probabilities and an underweighting of high probabilities. These two functions combined give rise to loss avoidance, the tendency to avoid losses of high probability and pick

actions that might result in a gain even with low probability (Cachon & Camerer, 1996; Erev, Ert, & Yechiam, 2008; Yechiam & Hochman, 2013). In the introductory example above, more patients would prefer a riskier treatment (e.g., radiation) when the information is framed as mortality rates, but a more definite treatment (e.g., surgery) when information is framed in terms of chances of survival; even when the objective survival probabilities for each treatment were the same. Although a patient may only have access to explicit descriptions of possible outcomes and their probabilities, doctors often rely on their own experiences with the different treatments (Groopman, 2007). For example, an experienced oncologist has seen the end results of many patients choosing one or the other treatment. Based on those experiences, would doctors be just as susceptible to biases like the framing effect? Would decision makers be susceptible to the framing bias when relying on their own experience to make a choice, rather than on a description?

1.1. Gains versus losses in decisions from experience

While most demonstrations of the framing bias have used descriptive scenarios as in the ADP, a growing body of research suggests that contrastingly different risk preferences are observed when decisions are made from experience. While in descriptive decisions, people often act as if low-probability outcomes were *more* probable than they really are (i.e., they are overweighted), people in experiential decisions act as if low probability outcomes were *less* probable than they really are (i.e., they are underweighted), resulting in a phenomenon coined as the *description-experience gap* (Barron & Erev, 2003; Camilleri & Newell, 2009; Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009). More generally, recent research suggests that well-known biases originally demonstrated with descriptive scenarios may not exist or might be weaker when people make decisions from experience (Dutt, Arlo-Costa, Helzner, & Gonzalez, 2013; Gonzalez, 2013; Harman & Gonzalez, 2015).

In descriptive decisions, a four-fold pattern of risky choice emerges from the probability (high, low) and the domain (gains, losses) of a decision problem. Prospect Theory explains this pattern by proposing risk aversion when the probability of winning is high or when the probability of losing is low, and risk seeking when the probability of winning is low or when the probability of losing is high (Tversky & Fox, 1995; Tversky & Kahneman, 1992). In contrast, a *reversed* four-fold pattern has been observed in decisions from experience (Hertwig, 2012, 2015; Hertwig et al., 2004). For example, in Hertwig (2012), a majority of respondents were risk seeking when the probability of winning was high or when the probability of losing was low; and they were risk averse when the probability of winning was low or when the probability of losing was high (Table 1). This reversed pattern has generally been explained by the assumption that people behave as if they underweight small probabilities and overweight both moderate and high probabilities (Hertwig, 2012; Hertwig & Erev, 2009).

However, the four-fold pattern refers to a *reflection effect*, where the sign of outcomes is reversed, resulting in *different* expected values (Fagley, 1993), in contrast to the *fram-*

Table 1
A typical set of problems that demonstrate a four-fold pattern in decisions from experience (from Hertwig, 2012); R (risky option), S (safe option), and EV (expected value)

Probability	Gain	Loss
	R: (outcome1, prob1, outcome2, prob2) S: Outcome, prob = 1	R: (outcome1, prob1, outcome2, prob2) S: Outcome, prob = 1
Low	R: 32, .1; 0, .9 (EV = 3.2) S: 3, 1.0 Risk aversion: 20%	R: -32, .1; 0, .9 (EV = -3.2) S: -3, 1.0 Risk seeking: 72%
High	R: 4, .8; 0, .2 (EV = 3.2) S: 3, 1.0 Risk seeking: 88%	R: -4, .8; 0, .2 (EV = -3.2) S: -3, 1.0 Risk aversion: 44%

ing effect, where problems in the gain and loss domains keep the *same* expected values (e.g., in the ADP problem, 200 people are expected to die, independent of the frame). For example, in the Low probability problems shown in Table 1, outcomes are mirrored from gains to losses by reversing the sign of outcomes from gains (+32) to losses (-32), thereby also changing the problems' expected values (3.2 in gains and -3.2 in losses).

Within the literature of decisions from experience, a large number of studies have tested the reflection effect (Barron & Erev, 2003; Ert & Yechiam, 2010; Hertwig et al., 2004; Ludvig, Madan, & Spetch, 2013; Ludvig & Spetch, 2011). Findings consistently replicate the contradicting risky choices in the gain and loss conditions (the description-experience gap), supporting the reversed four-fold pattern (Hertwig, 2012). Furthermore, differences also appear at the level of search (during a sampling process that precedes the consequential choice between the options): Participants search longer when facing possible losses relative to gains in a large number of problems similar to those shown in Table 1 (Lejarraga, Hertwig, & Gonzalez, 2012; Mehlhorn, Ben-Asher, Dutt, & Gonzalez, 2014). In these studies, loss aversion (i.e., people are thought to weigh losses more heavily than gains of the same magnitude) is a typical explanation for the framing effect in descriptive choices, but explanations for the reflection effect in decisions from experience are unclear. Loss aversion, albeit controversial, does not appear to be a plausible explanation for the contrasting choices in experiential decisions (Erev et al., 2008; Ert & Erev, 2013; Hochman & Yechiam, 2011). For example, participants in Erev et al. (2008) were similarly indifferent when selecting between the status quo (payoff of 0) and an equal chance to win 1,000 and lose 1,000. More strongly, recent evidence suggests that it is possible to manipulate experiences to obtain loss aversion, loss neutrality, or even a reversal of loss aversion through manipulating the range of possible gains and losses that people experience during experiments (Walasek & Stewart, 2015).

In summary, demonstrations of the framing effect in problems that maintain the same expected values across domains have not been conducted; and cognitive explanations of observed behavior in problems framed as gains and losses in decisions from experience are unclear.

1.2. Goals and overview of the current research

To clarify whether people are susceptible to the framing bias in decisions from experience, we first present an experiment comparing traditional description-based choice and experiential-based choice in the sampling paradigm using the ADP (Experiment 1). Then, to provide cognitive explanations of behavioral results in decisions from experience, we run simulations of participants' expected behavior in the ADP from experience using a computational cognitive model that has successfully predicted choice from memory in a large variety of decision-making domains (IBL [instance-based learning], derived from Instance-Based Learning Theory; Gonzalez, Lerch, & Lebiere, 2003). Analyses of values and memory retrieval probabilities in the IBL model compared to observed behavior suggest that the lack of framing bias emerges from under-experiencing the high-value outcome associated with a low probability in the risky option, making the safe option more attractive in both gains and losses.

This effect may be explained by the fact that high outcomes occur less often in the environment (*low objective probability*, given that the probability of the high-value outcome in the ADP is 1/3), or by the well-known phenomenon of *under-exploration*, by which there is a human tendency to limit their information search, even when they are given the opportunity to search extensively (Gonzalez & Dutt, 2011, 2012; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hills & Hertwig, 2010; Mehlhorn et al., 2014, 2015). We further investigate these two explanations using the IBL model ("out of the box") to make predictions in conditions where we vary the probabilities systematically and the number of samples taken before a choice is made. Based on these predictions, we select problems to use in Experiment 2, to verify the low objective probability and under-exploration phenomena that expand our understanding of the framing effect in decisions from experience.

2. Experiment 1

In the laboratory, experiential choices are often investigated in interactive paradigms, where participants learn about the possible outcomes from feedback, by selecting options represented with blank buttons on the screen and receiving the outcome resulting from a random draw of the probability associated with that button (Barron & Erev, 2003; Hertwig et al., 2004). The paradigm shown in Fig. 1 is an experiential equivalent (e.g., sampling paradigm) of the descriptive form of the ADP that we use in this research. Participants are presented with the ADP scenario and are told to sample the outcomes of the two programs for as long as they want and in whatever order they want, by pressing the respective buttons (labeled A and B). Participants do not know which is a risky and a safe option (the two options were randomized per participant). When a button is selected, an outcome (summary statistic according to the ADP) is drawn from the respective distribution and is displayed. Once participants feel sufficiently informed about the options' outcomes, they can proceed to the final choice phase in which they make a consequential choice between the two programs.

Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. You may first explore the possible scientific estimates of the outcomes of the two programs by clicking on the two buttons on the screen for as long as you want. Then, when you are ready to make a choice between the two programs, click on the “Go to Decision” button.

Fig. 1. An illustration of the sampling paradigm in the Asian Disease Problem (ADP) with the loss frame. The scenario is described, and the participant explores the possible outcomes of both options until she feels ready for the final decision. In the decision phase, she is asked to select one of the options for real.

In the loss frame exemplified in Fig. 1, a click in the risky option B displays the low-value outcome (a high loss of 600 people die) with a 2/3 probability or the high-value outcome (low loss of 0 people die) otherwise (1/3 probability); while a click in the safe option displays the medium-value outcome (400 people die) with a probability of 1. These values make the expected value of the two options equivalent. Participants are free to sample, without any consequences until they choose to proceed to the decision phase where one final consequential choice is made based on personal experiences from

sampling. The preferences in the final consequential choice are often compared to those made from description to identify differences in the two choice modes (Hertwig et al., 2004).

2.1. *Participants and design*

In total, 227 paid volunteers (58% male) with a mean age of 27.0 years ($SD = 10.8$) were recruited through Amazon Mechanical Turk and completed one ADP problem for a fixed payment of \$0.25. Participants were randomly assigned to one of four conditions: Description-Gain (51), Description-Loss (51), Experience-Gain (60), and Experience-Loss (65); and gave informed consent before starting. In the experience condition, participants were presented with the problem scenario in the sampling paradigm (Fig. 1). In the description condition, participants were given a choice between two options described in the ADP presented in the introduction. In the experience condition, eight additional subjects were tested, but they were excluded from all analyses because they did not follow instructions to sample at least once from each option.

2.2. *Choice behavior*

Table 2 shows the overall proportions of risky choices (Prisky) for each of the four groups. Note that in experiential decisions, only the choice at the decision phase is analyzed in Table 2. In agreement with the literature on descriptive choice, we observe a significant framing bias in decisions from description: We find a lower Prisky in the gain frame than in the loss frame. Additional comparisons against chance levels show that Prisky was significantly lower than .5 in gains, but it did not differ from .5 in losses. In contrast, we observe no significant framing bias in experience-based choice: Prisky does not differ between gains and losses.

Our results replicate the framing effect in the ADP in descriptive decisions but show no framing effect in experiential decisions. When making decisions from experience, humans were indifferent between the risky and safe options in both gains and losses; but the difference in the choice proportions between descriptive and experiential decisions did not lead to a significant description-experience gap.

2.3. *Sampling behavior*

As expected, participants under-explored; they took a small number of samples before making a choice (Table 3). Although there was large variability in the number of samples, there is no significant difference between the sample size of gains and losses ($t(113.21) = .304, p = .761$).

To analyze the actual sampling experiences, we calculated the experienced expected value (EEV) of each risky and safe option for each participant. The EEV is the accumulation of the product of each outcome's experienced probability and each outcome's value, over all the samples taken by the participant (t), as follows:

Table 2
Proportion of risky choices (Prisky) in the description and experience paradigms, separately for the different frames in Experiment 1

	Description		Experience		Description–Experience Gap (Description – Experience)
	Mean (SE) Prisky	Comparison Against Chance Level (.5)	Mean (SE) Prisky	Comparison Against Chance Level (.5)	
Gains	.33 (.07)	χ^2 (1, N = 51) = 5.667, p = .017 , ϕ = .34	.38 (.06)	χ^2 (1, N = 60) = 3.267, p = .071, ϕ = .24	–.05 (N = 111) = –.55, p = .585, ϕ = .05
Losses	.57 (.07)	χ^2 (1, N = 51) = .961, p = .327, ϕ = .14	.51 (.06)	χ^2 (1, N = 65) = .015, p = .901, ϕ = .02	+ .06, $z(N$ = 116) = .65, p = .514, ϕ = .06
Framing effect (Gains – Losses)		–.24, χ^2 (1, N = 102) = 5.70, p = .017 , ϕ = .24		–.13, χ^2 (1, N = 125) = 1.95, p = .162, ϕ = .13	

Note. SEs are standard errors for mean proportions. The p -values below .05 are indicated in bold.
The values in Italics in the Framing Effect are the difference of mean proportions between Gains and Losses. The values in Italics in the Description–Experience Gap are the difference of mean proportions between Description and Experience.

Table 3

Number of samples (median [range], mean) from the risky, the safe, and both options in each frame in the experience condition for Experiment 1

	Safe Option	Risky Option	Total
Gains	2 [1–18], 3.5	2 [1–18], 3.5	4 [2–36], 7.0
Losses	2 [1–25], 3.6	2 [1–30], 3.8	4 [2–46], 7.4

$$EEV_t = \sum_{i=1}^t \text{experienced probability}_i * \text{outcome}_i$$

EEV is an approximation of the value of each option used as the basis for choice in past research (in the IBL model, the option with the highest Blended value, similar to the EEV, is selected, see Gonzalez & Dutt, 2012; Mehlhorn et al., 2014; see also Hertwig & Pleskac, 2008 for a similar calculation). The experienced probability for each outcome is calculated by the frequency of occurrences divided by the number of samples.

Fig. 2a shows the average EEV of the safe and risky options in gains and losses, and Fig. 2b shows the average preferences (difference between the EEVs of the risky and safe options, $EEV_{\text{risky}} - EEV_{\text{safe}}$) in each frame. The average EEV of the risky option is lower than the safe option in gains (136 and 200), and it is only slightly higher in losses (–393 and –400), but their variability is quite large as seen by the standard deviations. Similarly, the average difference between their EEVs (Fig. 2b) suggests a tendency to

Fig. 2. (a) Average (± 1 SD) experienced expected values (EEVs) for the safe and risky options in gain and loss problems after all samples. (b) Average (± 1 SD) preference for the safe (risk aversion) or risky option (risk seeking), as indicated by the difference between EEVs in gain and loss problems.

Fig. 3. *Experienced probabilities* after all samples for the high and low outcomes in (a) gains (600, 0) and (b) losses (0, -600) in Experiment 1. In both plots, small crosses indicate the experienced probabilities for each individual participant. Large crosses indicate the mean experienced probabilities across participants. The numbers next to large crosses show the exact values of the mean probabilities, along with their standard deviation.

prefer the safe option in gains, and indifference between the safe and risky values in losses; again the standard deviations of these differences are very large.

To further investigate how these preferences emerged, we analyzed the experienced probabilities. Fig. 3 shows the resulting probabilities of each of the outcomes in the risky option in the gain (Fig. 3a) and loss (Fig. 3b) frames as experienced by individuals. Although there is large individual variability (small crosses), the average experienced probabilities (large crosses) are very close to the objective probabilities in the loss domain (.35 instead of .33; and .65 instead of .66), while small probabilities are slightly under-experienced (.23 instead of .33) and large probabilities are slightly over-experienced (.77 instead of .66) in gains. These results suggest that participants experienced the values of safe and risky options to be about the same in the loss frame, which resulted from a close to accurate experienced probability of the outcomes; while in the gain frame, participants experienced the safe option to be slightly more valuable than the risky option, probably due to an under-experience of the high-value outcome (600) and an over-experience of the low-value outcome (0) in the risky option.

2.4. From sampling to choice

To test the relationships between participants' experiences during sampling and their choices, we correlated each participant's (risky-safe) EEV differences during sampling with his or her final choice: The choice was coded as 1 (risky choice) or 0 (safe choice). We found strong positive correlations for gains: $r(58) = .70$, $p < .001$ and for losses: $r(63) = .59$, $p < .001$. At the aggregate level, preferences during sampling (Fig. 2b)

reflect the subsequent choice in experiential decisions very closely (Table 1): Indifference or slight risk aversion in gains and losses, with neither preferences being statistically different from chance. However, there was no significant correlation between participants' sample sizes and their final choice for gains: $r(58) = -.07$, $p = .614$, or for losses: $r(63) = .02$, $p = .884$, perhaps due to the fact that most participants take small samples in both gains and losses.

2.5. Summary

In descriptive decisions, we replicate the well-known framing bias, but we find risk indifference (i.e., no framing bias) in decisions from experience. Participants sampled very little (median of 4 samples), and there was no difference in the number of samples taken in gains and losses. Based on these samples, we find that participants' experienced probabilities are generally close to the objective probabilities in the ADP, more so for losses than for gains, resulting in very similar EEVs for the safe and risky options. We find that these experiences turned out to be a good predictor for participants' final choices at the individual and at the aggregate level. Furthermore, we find no description-experience gap (Table 2), neither in gains nor in losses. This is interesting, as the reflection effect finds that risk preferences in description and experience are reversed and it is often assumed that this due to a reversal of the experienced probabilities: While small probabilities are overweighted in decisions from description, they are underweighted in decisions from experience (Hertwig et al., 2004). As our results demonstrate, this was not the case for the ADP. From experience, participants built a close-to-accurate representation of the probabilities and were thereby able to escape the framing bias; thus, their choices were not reversed when compared to descriptive choices.

Many questions emerged from this experiment: What cognitive mechanisms explain indifference between the safe and risky options in gains and losses from experience? How does the number of samples taken influence the experienced probabilities? How do objective probabilities influence the framing effect? How general is the risk indifference in framing from experience? We turn next to cognitive modeling and further experimentation to answer these questions.

3. Explaining sampling and choice in framing tasks using cognitive models

Many models may be able to capture the sampling and choice behavior found in Experiment 1 (see Gonzalez & Dutt, 2011 for a summary of these models). In fact, a common approach in the cognitive sciences is to follow behavioral phenomena with the development of a cognitive model that is able to reproduce that behavior (Cassimatis, Bello, & Langley, 2010). This has also been a tradition in decision sciences, where task-specific models (not necessarily cognitive) are often put into competition to fit and predict human choices in one particular task (Erev, Ert, & Roth, 2010; Erev, Ert, Roth et al., 2010). A problem with this approach is that new tasks lead to the design of new models,

resulting in a multitude of highly task-specific models that fail to reproduce behavior in other closely related tasks (Cassimatis et al., 2010; Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012; Newell, 1973).

In the current research, we chose an existent instance-based learning (IBL) model that has shown to be a generalist rather than a specialist (Gonzalez, 2013; Hertwig, 2015). This model, originating from IBLT (Gonzalez et al., 2003), provides accurate predictions of choice behavior in a large diversity of sequential and dynamic decision-making tasks, particularly across decisions from experience paradigms (e.g., Gonzalez & Dutt, 2011; Lejarraga, Dutt, et al., 2012). In large-scale model comparisons, the IBL model has been shown to account for sampling and human choice processes better than the best task-specific models (Gonzalez & Dutt, 2011). Furthermore, the IBL model has also been shown to provide similar predictions to a reinforcement-learning model (from Yechiam & Busemeyer, 2005) that was used in similar repeated choice tasks (Lejarraga & Gonzalez, 2011). Although both reinforcement-learning and IBL models can predict choices accurately, the IBL model is preferred here because of its foundations in cognitive and decision sciences and its memory mechanisms taken from a well-known cognitive architecture (ACT-R, Anderson & Lebiere, 1998, 2003); and because it formulates value in a way that connects to traditional models of choice (i.e., expected value), allowing us to perform analyses of main elements of choice (probability and values).¹

The IBL model assumes that observed outcomes are stored in memory. The availability (i.e., the activation) of these instances in memory decays over time, and it is a function of the recency and frequency with which an instance was observed. Consequently, more frequent and more recent observed outcomes influence choices more heavily than more distant and rare ones. Most important, the IBL model, like many decision theories, proposes that the valuation of an option (called *blending*, in IBLT) depends on each outcome's likelihood (i.e., probability of retrieval of an instance from memory in IBLT) and value. This formulation is similar to the basic concept of expected value, which is essential for comparing to experienced probabilities and EEVs in the ADP and to be able to understand how over- or underweighting of probabilities may emerge from experience. The Appendix provides the mathematical formalization of the IBL model, as it has also been reported in several past publications (Gonzalez, 2013; Gonzalez & Dutt, 2011; Lejarraga, Dutt, et al., 2012).

Using this IBL model "out of the box" (i.e., without fitting parameters and using the default parameters values from ACT-R; for decay: $d = .5$, and noise $\sigma = .25$), we ran 1,000 simulations to make predictions in the ADP problem in the absence of human data. To investigate the effects of the sample size in the framing effect, we ran two groups of fixed sampling sizes (5, small; and 100, large) in the ADP problems in gains and losses. The data from the IBL model are analyzed in ways similar to human data. We calculated the average blended values (Eq. 1 in Appendix) and the retrieval probabilities for the risky outcomes (Eq. 2 in Appendix) after sampling.

Figs. 4a and c show the resulting average blended values (Eq. 1 in Appendix) for the safe and the risky options in gains and losses after 5 and 100 samples. Figs. 4b and d illustrate the predicted choice (the difference between the corresponding blended values).

Fig. 4. (a and c) Average (± 1 SD) blended values for the safe and risky options in gain and loss problems after a total of n samples (3a: $n = 100$; 3c: $n = 5$). (b and d) Average (± 1 SD) preference for the safe (risk aversion) or risky option (risk seeking), as indicated by the difference between blended values in gain and loss problems after a total of n samples (3b: $n = 100$; 3d: $n = 5$).

The model predicts a risk aversion in both gains and losses, but this preference seems stronger after 100 samples than after 5 (c.f. Fig. 4b and d), though the variability is large. Therefore, the IBL model predicts no framing effect in the ADP from experience, as it was found in Experiment 1, regardless of the number of samples.

Fig. 5. (a and b) Instance-based learning (IBL) predictions of the *retrieval probability after 100 samples* for the high and low outcomes in gains (600, 0) and losses (0, -600). (c and d) IBL predictions of the *retrieval probability after 5 samples* for the high and low outcomes in gains (600, 0) and losses (0, -600). In each plot, small crosses indicate the predicted probabilities for each of the 1,000 simulated participants. Large crosses indicate the mean predicted probabilities across those 1,000 simulated participants. The numbers next to the large crosses show the exact values of the mean probabilities, along with their standard deviation.

To explain these preferences in the model, we investigated the retrieval probabilities for the outcomes in the gain (Fig. 5a and c) and loss (Fig. 5b and d) frames. The predictions indicate that, on average, after 100 samples, the lower probability in the risky option will be strongly underweighted in gains (.16 vs. .33) and losses (.15 vs. .33), and the higher probability will be strongly overweighted in gains (.83 vs. .66) and losses (.85

vs. .66). However, after 5 samples, the model predicts low probabilities that are slightly underweighted in gains (.24 vs. .33) and losses (.21 vs. .33) and higher probabilities that are close to the objective probabilities in gains (.62 vs. .66) and losses (.64 vs. .66).

3.1. Explanations and insights emerging from the IBL model

Consider the ADP in the gain domain. Samples from the safe option produce a single instance in memory, which is reinforced with each sample from this option because the only possible outcome is 200. Samples from the risky option produce two unique instances in memory, one with a 600 outcome and one with a 0 outcome, which are reinforced when the respective outcomes are observed. Each of these instances has an activation value reflecting the frequency of the experienced outcome, the recency with which the outcome was experienced, and noise (Eq. 3 in Appendix). This activation value is used to calculate the probability of retrieval for each instance (Eq. 2 in Appendix), which is then used to calculate the blended value of each option (Eq. 1 in Appendix). The blended value for the safe option will be near 200 and grow ever closer to 200 with larger samples (see blended values for the safe gain in Fig. 4a and c). However, the blended value of the risky option depends on the frequency and recency of experiences for the 600 and the 0 outcomes. With each sample, the option with the highest blended value up to that specific sample is chosen. If the risky option was selected and “600” was experienced in an initial sample, the blended value of the risky option is likely to be higher than the blended value of the safe option, leading to a risky choice. Given that the higher outcome of the risky option (600) is generally under-experienced relative to its normative value, it ends up being underweighted (see Fig. 5a and c). The lower outcome of the risky option (0) generally becomes overweighted after 100 samples (see Fig. 5a) and slightly underweighted after 5 samples (see Fig. 5c). Consequently, the blended value of the safe option (200) is often higher than the blended value of the risky option (see Fig. 4). This difference is stronger after 100 samples (Fig. 4a) than after 5 samples (Fig. 4c). Therefore, a preference is developed for the safe option (stronger preference after 100 samples, Fig. 4b; than after 5 samples, Fig. 4d).

A very similar process applies to the loss domain. Samples from the safe option produce a single instance in memory as the only possible outcome is -400 . Samples from the risky option produce two unique instances in memory, one with a 0 outcome and one with a -600 outcome. Because the higher outcome of the risky option (0) tends to be under-experienced relative to its normative value, it ends up being underweighted with the lower outcome of the risky option (-600) becoming overweighted (after 100 samples, Fig. 5b; and slightly underweighted after 5 samples, Fig. 5d). Consequently, the resulting blended value of the safe option (-400) is often higher than the blended value of the risky option (see Fig. 4a and c), meaning that a preference is likely developed for the safe option.

In summary, the IBL model predicts a difference in how the probabilities are experienced according to the number of samples taken, but it also predicts no difference in choice between gains and losses, regardless of the sample size. The model's predictions

after five samples are closer to the observed behavior in Experiment 1 than the predictions after 100 samples. This seems reasonable, given the small sample sizes observed in participants (Table 3). According to the IBL model, large samples would result in underweighting of low probabilities and overweighting of high probabilities, while small samples would result in slight underweighting of low and high probabilities. Since the low probability is associated with the higher outcome in the ADP, the blending value predicts a slight risk aversion in gains and losses. These predictions are in agreement with results from Experiment 1, which shows no framing effect from experience.

There are, however, some slight differences in the observed experienced probabilities between gains and losses in Experiment 1 (Fig. 3), which the model does not capture (see Gonzalez, 2013, for some of the limitations of this IBL model), although ultimately the model's predictions agree with the lack of framing effect observed. Instead of attempting to fit the IBL model's parameters to human data (a very common approach expected to lead to lower discrepancies between the IBL model's predicted and observed behavior), we decided to produce a larger set of predictions from the IBL model to help determine the generality of the lack of framing effect observed in Experiment 1. For example, the IBL model makes interesting predictions regarding the number of samples, and the relationships between the samples taken and the probabilities associated to the outcomes in the risky options. More exploration is expected to lead to more underweighting of high outcomes and overweighting of low outcomes. Furthermore, the predicted experience depends directly on the objective probabilities associated to those outcomes. If the probability of the high outcome decreases, this might result in more underweighting of these outcomes (cf. Ludvig & Spetch, 2011; Ludvig et al., 2013). New predictions from the IBL model are presented next.

3.2. Predictions of framing from experience: Effects of number of samples and probability values

We ran the IBL model in seven different variations of the ADP, where we systematically varied the probabilities of the risky option and the corresponding value of the safe outcome, while keeping the expected values equal across both. Table 4 presents the list of seven problems, with Problem 6 being the standard ADP used in Experiment 1. Note that these problems were created so that the probability of the high outcome in both gains and losses was systematically decreased (while the probability of the low outcome in the risky option increased accordingly), and the value of the outcome in the safe option was adjusted to maintain the essential condition of equal expected values of the framing task.

The IBL model with default ACT-R parameters: $d = .5$, $\sigma = .25$, ran 1,000 simulated participants for each of the seven problems in groups that varied in the number of samples, 5 or 100 samples, and the frame, gains or losses.

Fig. 6 shows the resulting predicted average preferences in gains and losses after 100 samples and after five samples. These preferences are again calculated as the difference between the blended values of the risky and safe options. The model consistently predicts

Table 4

Overview of the seven problems for which predictions were generated in the gain and the loss frames. These problems vary the probabilities of the outcomes in the risky option systematically. Two of these problems are used in Experiment 2 (in bold): Problem 6, the standard ADP used in Experiment 1; and Problem 6, in which the probability of the high outcome is low (a “rare event”)

Frame	Gain		Loss	
Problem Number	Risky Option: High Outcome (pH), Low Outcome (pL)	Outcome of the Safe Option (Probability = 1)	Risky Option: High Outcome (pH), Low Outcome (pL)	Outcome of the Safe Option (Probability = 1)
1	600 (1/10), 0 (9/10)	60 (1)	0 (1/10), –600 (9/10)	–540 (1)
2	600 (1/8), 0 (7/8)	75 (1)	0 (1/8), –600 (7/8)	–525 (1)
3	600 (1/6), 0 (5/6)	100 (1)	0 (1/6), –600 (5/6)	–500 (1)
4	600 (1/5), 0 (4/5)	120 (1)	0 (1/5), –600 (4/5)	–480 (1)
5	600 (1/4), 0 (3/4)	150 (1)	0 (1/4), –600 (3/4)	–450 (1)
6	600 (1/3), 0 (2/3)	200 (1)	0 (1/3), –600 (2/3)	–400 (1)
7	600 (5/12), 0 (7/12)	250 (1)	0 (5/12), –600 (7/12)	–350 (1)

risk-averse preferences (on average) for all values of probabilities in both gains and losses after 100 samples (Fig. 7a and b). Also, the results suggest that as the probabilities of the risky outcomes become less extreme (e.g., going from Problem 1 to Problem 7, see Table 4), the predicted average risk aversion becomes stronger but also more variable. These preferences, however, are close to indifference after five samples and they are highly variable, regardless of the probability values and the frame (Fig. 6c and d).

According to our discussion above and within the scope of this paper, we decided to test two predicted phenomena in a second experiment: the effects of the number of samples taken before choice and the effect of the objective probability associated with the high outcome in the risky option.

4. Experiment 2: Framing effects with small and large sample size and with high and low probabilities

In Experiment 2, we used Problem 3 to contrast to Problem 6 (i.e., the standard ADP) framed as gains or losses, in decisions from experience conditions in which participants were asked to sample *exactly* 5 or 100 times before making a choice. Problem 6 uses the probabilities of the ADP for the high outcome (1/3), whereas Problem 3 uses extreme probabilities where the high outcome occurs more rarely (1/6).

In total, 800 paid volunteers (58% male) with a mean age of 31.6 years ($SD = 10.0$) were recruited through Amazon Mechanical Turk and completed one problem for a fixed number of trials for a fixed payment of \$0.25. Participants were randomly assigned to one of eight conditions constructed by the probability of the high outcome (1/3 or 1/6), number of samples before making a choice (5 or 100), and frame (Gain or Loss). There were 100

Fig. 6. Predicted preferences for each of the seven problem types for gains (left) and losses (right). Problems are sorted by the probability of the high outcome in gains (600) and the low outcome in losses (-600). Preferences are based on the average (± 1 SD) difference between the blended values of the risky and safe options after 100 samples (a and b) and after 5 samples (c and d). P3 indicates the predictions for the standard Asian Disease Problem (ADP) as tested in Experiment 1 (also shown in Fig. 3). P6 and P3 were both tested in Experiment 2.

participants in each condition. The sampling paradigm and procedure was similar to Experiment 1, except for the fixed number of samples. Although participants were not told the number of samples in advance, they were advised that the number of samples would be fixed, and a counter indicating the number of samples was displayed in every trial. The paradigm automatically advanced participants to make a final choice after a fixed number of samples. They were still able to explore the buttons in the order they desired.

Fig. 7. Average (± 1 SD) preference for the safe (risk aversion) or risky option (risk seeking), as indicated by the difference between experienced expected values for the two problem types tested in Experiment 2. (a) Gain problems, 100 samples. (b) Loss problems, 100 samples. (c) Gain problems, 5 samples. (d) Loss problems, 5 samples.

4.1. Choice behavior

Table 5 shows the overall Prisky in the final choice for each of the eight groups. We find a significant framing effect when the number of samples is large (100), but not when the number of samples is small (5). Results for Problem 6 and small number of samples replicate the lack of framing effect of Experiment 1. All of the choice proportions indicate either indifference between the two options or risk aversion, except for one group:

Table 5
Proportion of risky choices (Prisky) in the eight conditions of Experiment 2

	Mean (SE) Prisky	Comparison Against Chance Level (.5)
Problem 6-Probability (1/3) and samples (100)		
Gains	.45 (.05)	$\chi^2(1, N = 100) = 1.0, p = .317, \phi = .10$
Losses	.62 (.05)	$\chi^2(1, N = 100) = 5.76, p = .016, \phi = .24$
Framing effect	-.17, $\chi^2(1, N = 200) = 5.81, p = .016, \phi = .17$	
Problem 6-Probability (1/3) and samples (5)		
Gains	.40 (.05)	$\chi^2(1, N = 100) = 4.0, p = .046, \phi = .20$
Losses	.36 (.05)	$\chi^2(1, N = 100) = 7.84, p = .005, \phi = .28$
Framing effect	+.04, $\chi^2(1, N = 200) = .34, p = .560, \phi = .04$	
Problem 3-Probability (1/6) and samples (100)		
Gains	.30 (.05)	$\chi^2(1, N = 100) = 16.0, p < .001, \phi = .40$
Losses	.50 (.05)	$\chi^2(1, N = 100) = 0, p = 1, \phi = 0$
Framing effect	-.20, $\chi^2(1, N = 200) = 8.33, p = .004, \phi = .20$	
Problem 3-Probability (1/6) and samples (5)		
Gains	.24 (.04)	$\chi^2(1, N = 100) = 27.04, p < .001, \phi = .52$
Losses	.25 (.04)	$\chi^2(1, N = 100) = 25, p < .001, \phi = .50$
Framing effect	-.01, $\chi^2(1, N = 200) = .027, p = .869, \phi = .01$	

Note. SEs are standard errors for proportions.
The *p-values* below 0.05 are indicated in bold. The values in Italics for the Framing Effect are the difference in mean proportions between Gains and Losses.

Problem 6 (standard ADP) in the loss domain with 100 samples, which showed risk seeking.

4.2. Sampling behavior

We analyzed the distribution of samples taken from the safe and the risky options in gains and losses before making a choice. Table 6 shows the results. Again, there was large heterogeneity in the samples taken from the safe and the risky options, but none of the comparisons between the gain and loss frames or the safe and risky option were significant ($p > .05$).

As in Experiment 1, we calculated the EEV of each option after the sampling phase for each participant. Fig. 7 shows the average preferences (EEV risky—EEV safe) for each condition. The observed behavior indicates indifference or slight risk aversion (on average) after sampling, regardless of the frame or the number of samples. Furthermore, these results also show stronger average risk aversion in Problem 6 than in 3; and wider variance in behavior after 5 samples than after 100 samples. These results are in agreement with the predicted preferences from the IBL model (c.f. Problems 3 and 6 in Fig. 6). However, the IBL model predicts stronger average risk aversion after 100 than after 5 samples, and this pattern is not observed in human data.

Table 6

Number of samples (median [range], mean) from the risky and the safe options in both gains and losses for Experiment 2

	Safe Option	Risky Option
Problem 6-Probability (1/3) and samples (100)		
Gains	49 [2–99], 49.8	51 [1–98], 50.2
Losses	48 [17–99], 50.5	52 [1–83], 49.5
Problem 6-Probability (1/3) and samples (5)		
Gains	3 [1–4], 2.6	2 [1–4], 2.4
Losses	2 [0–4], 2.45	3 [1–5], 2.55
Problem 3-Probability (1/6) and samples (100)		
Gains	50 [10–99], 56.5	50 [1–90], 43.5
Losses	50 [13–96], 51.1	50 [4–87], 48.9
Problem 6-Probability (1/6) and samples (5)		
Gains	3 [1–4], 2.7	2 [1–4], 2.3
Losses	3 [0–4], 2.6	2 [1–5], 2.4

Table 7

Number of samples (median [range], mean) from the risky and the safe options in both gains and losses taken by the instance-based learning (IBL) model simulations. Values are based on 1,000 simulated participants in each condition

	Safe Option	Risky Option
Problem 6-Probability (1/3) and samples (100)		
Gains	82 [29–96], 78.8	18 [4–71], 21.2
Losses	83 [24–96], 79.2	17 [4–76], 20.8
Problem 6-Probability (1/3) and samples (5)		
Gains	3 [1–4], 2.6	2 [1–4], 2.4
Losses	3 [0–4], 2.6	2 [1–5], 2.4
Problem 3-Probability (1/6) and samples (100)		
Gains	91 [53–96], 88.1	9 [4–47], 11.9
Losses	91 [51–96], 87.9	9 [4–49], 12.1
Problem 6-Probability (1/6) and samples (5)		
Gains	3 [1–4], 2.7	2 [1–4], 2.3
Losses	3 [1–4], 2.7	2 [1–4], 2.3

To better understand these differences, we analyzed the number of samples taken by the IBL model from the safe and from the risky option in the various conditions. Table 7 reveals that in agreement with observed sampling patterns, the IBL model shows no difference between gains and losses. However, in contrast to observed sampling patterns, the IBL model samples significantly more from the safe than the risky option, particularly with large samples (100) than small samples (5), explaining the stronger risk aversion developed in the model after 100 than after 5 samples.

Table 8
Correlation between participants' experienced expected value (EEV) differences (safe-risky) and their final choice in each of the eight conditions

Condition	Correlation
Problem 6-Probability (1/3) and samples (100)	
Gains	$r(98) = .34, p = .001$
Losses	$r(98) = .44, p < .001$
Problem 6-Probability (1/3) and samples (5)	
Gains	$r(98) = .61, p < .001$
Losses	$r(97) = .68, p < .001$
Problem 3-Probability (1/6) and samples (100)	
Gains	$r(98) = .27, p = .007$
Losses	$r(98) = .42, p < .001$
Problem 3-Probability (1/6) and samples (5)	
Gains	$r(98) = .61, p < .001$
Losses	$r(97) = .61, p < .001$

Note. We had one participant each in two conditions that never sampled from the safe option, which explains the df of 97. For these participants, no EEV difference could be calculated.
The *p-values* below 0.05 are indicated in bold.

4.3. From sampling to choice

Average preferences show indifference and risk aversion in all conditions which generally agree with the subsequent choice, except for choice behavior in Problem 6 with 100 samples, which indicate significant risk seeking in losses. Interestingly, all the correlations of the individual participant's EEV differences after sampling with their final choice are strong and positive in each of the eight conditions (see Table 8), and higher in the conditions with lower sample sizes (5 samples) than large sample sizes (100 samples). This result suggest a differential effect of risk preferences after sampling with respect to the average choice behavior at the average and individual levels, an issue that has been discussed in past research (Hills & Hertwig, 2010; Gonzalez & Dutt, 2012).

5. Discussion

Most decision biases have traditionally been shown through the use of descriptive gambles where outcomes and probabilities are stated explicitly. Results from the current experiments replicate the framing bias in decisions from description, but the bias disappears in decisions from experience and when participants are free to sample information for as long as they desire before making a choice. The analyses from an IBL cognitive model explain this effect by how participants may under-experience the higher outcome in the risky option, and predict that the presence or absence of the framing effect will depend on the number of samples taken and the objective probabilities associated to the outcomes.

Results from Experiment 1 expand demonstrations of biases originally found with descriptive scenarios that become weaker or disappear when choices are made from experience (Dutt et al., 2013; Gonzalez, 2013; Harman & Gonzalez, 2015). We also find that, the description-experience gap (e.g., Gonzalez & Dutt, 2011; Hertwig, 2012; Hertwig & Erev, 2009), disappears, and instead, we find consistent risk behavior for gains and losses, in contrast with the reversed four-fold pattern (Hertwig, 2012). Our results clarify some boundaries of the four-fold pattern and the description-experience gap: in problems that have equivalent expected values people become risk-indifferent or averse from experience, in contrast to situations such as the reflection effect (Table 1), where the signs of the outcomes are reversed, changing the problem's expected values (Barron & Erev, 2003; Ert & Yechiam, 2010; Hertwig et al., 2004; Ludvig & Spetch, 2011; Ludvig et al., 2013). As suggested by the IBL model, this may be a consequence of the development of preferences for the option with the highest accumulated value (blended value) through experience. The lack of the framing effect from experience may be due to under-experience of the higher outcome of a risky option, similarly for gains and losses, making the safe option equally or slightly more valuable than the risky option.

Furthermore, predictions from the IBL model suggest two possible factors that may influence the under-experience of the high outcome in the risky option: One is the objective probability of this outcome and a second one is the limited information search. When the rare outcome is high valued and the more common outcome is low valued (as it is the case in the ADP for gain and losses), people would develop an intuition that the risky option is less valuable than the safe option, resulting in indifference or risk aversion. In addition, given that people generally do not search extensively, the probability of retrieving an outcome from memory may be influenced by what outcomes are experienced during that small sample.

Results from Experiment 2 suggest that the amount of exploration is a key factor in finding a framing effect. When participants were forced to make five samples before making a choice, we reproduced the results from Experiment 1: Participants develop a tendency to prefer the safe over the risky option in gains and losses, regardless of the objective probabilities and the framing of the problems. In contrast, when participants were forced to make 100 samples before making a choice, the framing effect emerges: People are more risk seeking in losses and risk averse in gains. This result is puzzling; with large samples, one would expect more informed and accurate choices than with small samples, given that there are more chances of identifying the accurate probabilities in the two options. Hertwig and Pleskac (2008) explain this effect as an amplification of the experienced expected values of the two options, which renders small samples easier. This effect is observed in our results as the difference of the EEVs generated after 5 or 100 samples (Fig. 7), where after small samples, the preference for the safe option is more clear (on average) than after large samples. However, although the sampling behavior shows that participants do indeed develop a more accurate representation of the environment with large samples, their choices after 100 samples do not show indifference between the two options (safe and risky). Rather, we observe significant risk seeking in losses (Problem 6) and significant risk aversion in gains (Problem 3),

suggesting a disconnection between the sampling and the choice processes. This result is supported by the lower correlation values (albeit significant) between the EEV differences after 100 samples and the final choice compared to the correlations after 5 samples. Thus, with less samples, people rely more in the limited experience to make a final choice; with more samples people rely less on their experience to make a choice.

Clearly the IBL model although it predicts the final choice accurately, it fails to account for the amount of exploration done for each of the two options with large samples. The IBL model predicts more sampling of the safe than the risky option but humans balance the exploration between the risky and safe option equally (on average). This may be an indication of human *decay avoidance*, where humans may aim to maintain their memories of outcomes from the risky and safe options as equally accessible, particularly with large samples, in an attempt to make more accurate choices, while the IBL model is willing to forget the outcomes of the risky option by focusing on choosing the safe option more frequently. It is also possible that this effect results only in the sampling paradigm, in which the sampling process is not consequential, but in situations of repeated consequential choice (e.g., Barron & Erev, 2003), it may be expected that participants would behave more similarly to what the IBL model predicts.

Given that the IBL model was not fit to human data, it is possible that it is not accounting for the individual differences in memory decay (d parameter) accurately. In general, as discussed in past research (Gonzalez & Dutt, 2011), the IBL model cannot account for the amount of sampling and it would benefit from the development of an accumulation mechanism that would indicate a stopping rule in sampling. Such a mechanism would explain the decision of *when* to stop sampling in the sampling paradigm. The integration of sequential-sampling models (Busemeyer, 1985; Busemeyer & Townsend, 1993; Ratcliff, 1978; Ratcliff & Smith, 2004) with learning and choice models like IBL or Reinforcement Learning (e.g., Sutton & Barto, 1998) may improve our ability to account for human sampling and choice processes in an integrated theory. This and other ideas are relevant for future research. Research should investigate the boundaries and connections between the description–experience gap and the probability of different outcome values. The non-significant description–experience gap found in Experiment 1 suggests that the gap may be influenced by contextual or domain effects (such as in the ADP). Up to this point the gap has mostly been investigated in problems involving only monetary outcomes, and out of context. Also, just as the sampling paradigm in decisions from experience opens a window to the processes occurring while preferences are being formed through experience, there is a need to investigate the processes occurring during an analogous stage within decisions from description. For example, it is possible that even when participants do not physically “sample” the different options, they may “play out” the scenarios in their minds and imagine the possible outcomes. An fMRI study lends support to this possibility (Gonzalez, Dana, Koshino, & Just, 2005): Differences in brain activity levels in response to certain and risky choices were found in the parietal lobes, often associated with imagery processes. This relationship with imagery is also supported by behavioral findings, suggesting that people are more likely to imagine bad outcomes and how those outcomes could have been improved (Kahneman & Miller,

1986). Such finer process-level analyses would help us improve our understanding of the similarities and differences in the risk tendencies between experiential and descriptive decisions in the gain and loss domains.

Acknowledgments

This research is supported by the National Science Foundation Award number: 1154012 to Cleotilde Gonzalez. The authors thank Hau-yu Wong for providing insightful comments.

Note

1. In contrast to the IBL model, the RL model assumes that outcomes and their likelihood alter simultaneously each option's expectancy (or attractiveness), and they are therefore not conceived as separate concepts. Also, the RL model does not assume storage of outcomes in memory, but only the storage of each single expectancy developed for each option. Finally, the RL model gives a distinctive treatment to obtained versus other outcomes (e.g., foregone outcomes), which would not explain a "sampling" process in which outcomes from sampling are not consequential (Lejarraga & Gonzalez, 2011).

References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., & Lebiere, C. (2003). The Newell test for a theory of mind. *Behavioral and Brain Sciences*, 26(5), 587–639. doi:10.1017/S0140525X0300013X.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215–233. doi:10.1002/bdm.443.
- Bussemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 11(3), 538–564. doi:10.1037/0278-7393.11.3.538.
- Bussemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 106(3), 432–459. doi:10.1037/0033-295X.100.3.432.
- Cachon, G. P., & Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *The Quarterly Journal of Economics*, 111(1), 165–194. doi:10.2307/2946661.
- Camilleri, A. R., & Newell, B. R. (2009). The role of mental representations in experience-based choice. *Judgment & Decision Making*, 4(7), 518–529.
- Cassimatis, N. L., Bello, P., & Langley, P. (2010). Ability, breadth, and parsimony in computational models of higher-order cognition. *Cognitive Science*, 32(8), 1304–1322. doi:10.1080/03640210802455175.

- Dutt, V., Arlo-Costa, H., Helzner, J., & Gonzalez, C. (2013). The description-experience gap in risky and ambiguous gambles. *Journal of Behavioral Decision Making*. doi:10.1002/bdm.1808. Advance online publication.
- Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games*, 1(2), 117–136. doi:10.3390/g1020117.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., Hertwig, R., Stewart, T., West, R., & Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47. doi:10.1002/bdm.683.
- Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21(5), 575–597. doi:10.1002/bdm.602.
- Ert, E., & Erev, I. (2013). On the descriptive value of loss aversion in decisions under risk: Six clarifications. *Judgment and Decision Making*, 8(3), 214–2325.
- Ert, E., & Yechiam, E. (2010). Consistent constructs in individuals' risk taking in decisions from experience. *Acta Psychologica*, 134(2), 225–232. doi:10.1016/j.actpsy.2010.02.003.
- Fagley, N. S. (1993). A note concerning reflection effects versus framing effects. *Psychological Bulletin*, 113(3), 451–452. doi:10.1037/0033-2909.113.3.451.
- Fox, C. R., & Poldrack, R. A. (2009). Prospect theory and the brain. In P. W. Glimcher, E. Fehr, C. Camerer, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 145–173). San Diego, CA: Academic Press.
- Gonzalez, C. (2013). The boundaries of Instance-based Learning Theory for explaining decisions from experience. In V. S. Pammi, & N. Srinivasan (Eds.), *Progress in brain research* (Vol. 202, pp. 73–98). Amsterdam, the Netherlands: Elsevier.
- Gonzalez, C., Dana, J., Koshino, H., & Just, M. A. (2005). The framing effect and risky decisions: Examining cognitive functions with fMRI. *Journal of Economic Psychology*, 26(1), 1–20. doi:10.1016/j.joep.2004.08.004.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review*, 118(4), 523–551. doi:10.1037/a0024558.
- Gonzalez, C., & Dutt, V. (2012). Refuting data aggregation arguments and how the IBL model stands criticism: A reply to Hills and Hertwig (2012). *Psychological Review*, 119(4), 893–898. doi:10.1037/a0029445.
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27(4), 591–635. doi:10.1016/S0364-0213(03)00031-4.
- Groopman, J. (2007). *How doctors think*. New York, NY: First Mariner Books.
- Harman, J. L., & Gonzalez, C. (2015). Allais from experience: Choice consistency, rare events, and common consequences in repeated decisions. *Journal of Behavioral Decision Making*. Published online in Wiley Online Library (wileyonlinelibrary.com). doi:10.1002/bdm.1855.
- 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. doi:10.1002/bdm.598.
- Hertwig, R. (2012). The psychology and rationality of decisions from experience. *Synthese*, 187(1), 269–292. doi:10.1007/s11229-011-0024-4.
- Hertwig, R. (2015). Decisions from experience. In G. Keren & G. Wu (Eds.), *The Wiley Blackwell handbook of judgment and decision making*, 2 vols. Malden, MA: Wiley-Blackwell.
- 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. doi:10.1111/j.0956-7976.2004.00715.x.
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523. doi:10.1016/j.tics.2009.09.004.

- 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, England: Oxford University Press.
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, 21(12), 1787–1792.
- Hochman, G., & Yechiam, E. (2011). Loss aversion in the eye and in the heart: The autonomic nervous system's responses to losses. *Journal of Behavioral Decision Making*, 24(2), 140–156. doi:10.1002/bdm.692.
- Kahneman, D., & Miller, D. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93(2), 136–153. doi:10.1037/0033-295X.93.2.136.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- 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. doi:10.1002/bdm.722.
- 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. doi:10.1016/j.obhdp.2011.05.001.
- Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, 124(3), 334–342. doi:10.1016/j.cognition.2012.06.002.
- Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2013). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*. doi:10.1002/bdm.1792. Advance online publication.
- Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS ONE*, 6(6), e20262. doi:10.1371/journal.pone.0020262.
- Mehlhorn, K., Ben-Asher, N., Dutt, V., & Gonzalez, C. (2014). Observed variability and values matter: Towards a better understanding of information search and decision from experience. *Journal of Behavioral Decision Making*, 27(4), 328–339. doi:10.1002/bdm.1809.
- Mehlhorn, K., Newell, B. R., Lee, M., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. *Decision*. Advance online publication. doi:10.1037/dec0000033.
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–310). New York: Academic Press.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. doi:10.1037/0033-295X.85.2.59.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2), 333–367. doi:10.1037/0033-295X.111.2.333.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, 102(2), 269–283. doi:10.1037/0033-295X.102.2.269.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458. doi:10.1126/science.7455683.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *The Journal of Business*, 59(4), S251–S278. doi:10.2307/2352759.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323. doi:10.1007/BF00122574.

- Walasek, L., & Stewart, N. (2015). How to make loss aversion disappear and reverse: Tests of the decision by sampling origin of loss aversion. *Journal of Experimental Psychology: General*, 144(1), 7–11. doi:10.1037/xge0000039.
- Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, 12(3), 387–402. doi:10.3758/BF03193783.
- Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, 139(2), 497–518. doi:10.1037/a0029383.

Appendix

Instance-based learning model

The instance-based learning (IBL) model (Gonzalez & Dutt, 2011; Gonzalez et al., 2003; Lejarraga, Dutt, et al., 2012) evaluates an option according to its Blended Value, and the model chooses the option with the highest Blended Value $V_{j,t}$.

The *Blended Value* V of option j is

$$V_j = \sum_{i=1}^n p_i x_i \quad (1)$$

where x_i is the value of the observed outcome i , and p_i is the probability of retrieving that outcome from memory. At trial t , the *probability of retrieval* of observed outcome i is a function of the activation of that outcome relative to the activation of all other observed outcomes k in option j

$$P_{i,t} = \frac{e^{\frac{A_{i,t}}{\tau}}}{\sum_k e^{\frac{A_{k,t}}{\tau}}} \quad (2)$$

where τ is random noise defined as $\tau = \sigma\sqrt{2}$, and σ is a free parameter. In the current study, we used ACT-R's default noise value of .25. At trial t , the activation (Anderson & Lebiere, 1998) of an outcome i is:

$$A_{i,t} = \ln \sum_{t_p \in \{1, \dots, t-1\}} (t - t_p)^{-d} + \sigma \ln \left(\frac{1 - \gamma_{i,t}}{\gamma_{i,t}} \right) \quad (3)$$

where d is a decay free parameter. In the current study, we used ACT-R's default decay value of .5. $\gamma_{i,t}$ is a random draw from a uniform distribution bounded between 0 and 1 for each outcome and trial, and t_p is each of the previous trial indexes in which the outcome i was encountered.

Because memory is unlikely to be empty when starting a task, the model assumes that some initial expectation exists in memory before any choice is made (Lejarraga, Dutt, et al., 2012). Two instances with pre-populated initial expectations were set to \$600 for

each of the options in the gain domain, and to \$0 for each of the options in the loss domain. These are initially the only active instances and initial expectations for each option. As a result, the probability of retrieval of each initial expectation is 1, given that it is the only outcome active in memory for each option. Thus, the model chooses randomly in the first trial.

Explanations for the representations used in the ADP

According to the IBL model, when a choice is made and the resulting outcome is experienced, the choice-outcome association (i.e., an instance) is stored in memory. Each instance has a value of activation and when a previously experienced instance is experienced again, that activation is strengthened. Activation is a concept from the ACT-R cognitive architecture, and it reflects how readily available an instance in memory is and how easy and quickly that instance can be retrieved (Anderson & Lebiere, 1998). The activation in the IBL model of binary choice accounts for frequency, recency, and noise of the experiential process (Eq. 3). Each option (e.g., S: 200, $p = 1$ or R: 600, $p = .33$, 0 $p = .67$ in gains) can have several instances that are created when they are experienced (e.g., S, 200; R, 600; R, 0). Each instance has a probability of being retrieved from memory (Eq. 2), which is a function of its relative activation level. On each trial (t), the model calculates a *blended value* (Eq. 1) for each option (risky and safe) and selects the option with the highest blended value. The blended value is equivalent in form and function to the traditional notion of “expected value,” with the only difference being the value of each outcome is multiplied by its probability of memory retrieval rather than by its objective probability. The resulting activation levels are used in the subsequent trial ($t + 1$), which again influences the probability of memory retrieval, the blended value, and subsequently the choice.