

Allais from Experience: Choice Consistency, Rare Events, and Common Consequences in Repeated Decisions

JASON L. HARMAN* and CLEOTILDE GONZALEZ

Dynamic Decision Making Laboratory, Carnegie Mellon University, Pittsburgh, PA, USA

ABSTRACT

The Allais Paradox is a well-known bias in which people's preferences result in contradictory choices between two normatively identical gamble pairs. Studies have shown that these preference reversals depend on how information is described and presented. In an experiment, we investigate the Allais gambles in several formats including an experiential paradigm, where participants make selections from two blank buttons and get an outcome as a result of a draw from distributions of outcomes in the selected gamble. Results indicate that a large proportion of Allais reversals are found in the traditional descriptive format, they are reduced when gambles are presented in a descriptive table format, and they disappear when choices are made from experience. Although a majority of participants made consistent choices from experience, the proportion of individual reversals is similar to that of descriptive choices. Detailed analyses of experiential choice suggest interesting behavioral differences between participants classified as consistent and those classified as reversals: consistent participants explore and maximize more than reversal participants. Furthermore, consistent participants demonstrate a different switching behavior after experiencing a rare outcome than do reversal participants. We find that, both, overweighting or underweighting of rare outcomes may occur within the course of an individual's experience; depending on the timing of those experiences, the magnitude of the outcomes observed, and the general accumulated value of the gambles. Copyright © 2015 John Wiley & Sons, Ltd.

KEY WORDS Allais Paradox; decisions from experience; Instance-Based Learning; repeated decisions

INTRODUCTION

Modern decision making research traces its roots to a series of letters between French mathematicians, Blaise Pascal and Pierre de Fermat. Originally a discussion of various gambling problems, their collaboration resulted in the concept of mathematical expectation. The basic idea, later modified by Daniel Bernoulli (1954), is that by multiplying the probability of an outcome by its subjective value, the expected utility of that outcome can be determined. Combining the expected utility of each outcome in a gamble, you can then compare the overall expected utility of different gambles. Choosing the gamble that maximizes your expected utility became the core idea of rational choice. Von Neumann and Morgenstern (1947), and later Savage (1954), formalized these principles into a model of human decision making, the expected utility theory (EUT). Their works included a series of axioms, which quickly became the most influential theory of human decision making but was also quickly challenged. The first evidence of systematic violations of EUT was proposed by French economist Maurice Allais (1953). The now classic Allais Paradox can be illustrated by the following two gamble pairs:

Gamble pair 1:

A : 1000 ($p = 1$)

B : 1000($p = .89$), 5000($p = .1$), 0($p = .01$)

Gamble pair 2:

A' : 1000($p = .11$), 0($p = .89$)

B' : 5000($p = .1$), 0($p = .9$)

The independence axiom in EUT states that a rational decision maker should not base his or her preference on outcomes that are identical in amount and probability between gambles. Common outcomes between two gambles should cancel out. For example, in Gamble pair 1, gambles A and B result in 1000 for 89% of the time (gamble A also results in 1000 the other 11% of the time), and in Gamble pair 2, gambles A and B result in 0 for 89% of the time (gamble B also results in 0 another 1% of the time). According to EUT, the common outcome of 1000 for 89% of the time in Gamble pair 1 and of 0 for 89% of the time in Gamble pair 2 should cancel out, and both pairs of gambles would reduce to the following identical gambles:

A(A') : 1000($p = .11$) B(B') : 0($p = .01$), 5000($p = .1$)

If EUT's axioms are true, a person who prefers A to B must also prefer A' to B'. However, behavioral evidence shows that when given gamble pairs 1 and 2, most decision makers prefer A to B and B' to A' (choices inconsistent with EUT's axioms are referred to as *reversals*, 53% as per Allais, 1979).

Since its inception, the Allais gambles have been a cornerstone in decision science research. A great deal of research has sought to reduce the observed Allais reversals while defending the appropriateness of EUT's axioms. However, most attempts to reduce the Allais reversals have used modifications of experimental paradigms where participants are given all relevant information explicitly in a descriptive

*Correspondence to: Jason L. Harman, Dynamic Decision Making Laboratory, Carnegie Mellon University, 5000 Forbes Avenue, PH 208 Pittsburgh, PA 15213, USA. E-mail: jharman@cmu.edu

form. Some of these attempts have shown that by modifying the way descriptive information is presented (e.g., table format), it is possible to reduce the reversals, although not fully eliminate them (Incekara-Hafalir & Stecher, 2012). On the other hand, a growing body of research suggests that systematically different choices are observed when explicit information is removed and decisions are made purely from experience with immediate feedback (decisions from experience (DFE); Hertwig, Barron, Weber, & Erev, 2004).

In this research, we study the Allais gambles in multiple presentation formats: a traditional descriptive format, a table format, and a DFE format. We review literature attempting to eliminate Allais reversals through manipulations of descriptive information and then introduce the DFE paradigm. Next, we present an experiment using these different formats. In our analyses of results, we rely on predictions from an Instance-Based Learning (IBL) model (e.g., Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012). The model's predictions provide insight into the process of reversals and consistent behavior from experience: while consistent participants attempt to maximize their experienced expected value by exploring more, reversal participants explore less and experience rare outcomes in ways that generate a surprise and a switch to the gamble, like the hot-stove effect (Denrell & March, 2001). Choice reversals depend on the recency of rare outcome experiences and the frequency of experiencing outcomes common in the two gambles. Data from a simple recency model (Sliding Window of Information Model (SWIM); Ashby & Rakow, 2014) provide additional support to the behavioral phenomena.

ALTERNATIVE DESCRIPTIVE REPRESENTATIONS TO ELIMINATE AND EXPLAIN ALLAIS REVERSALS

Many defenders of EUT have argued that Allais reversals are the result of poor understanding, thoughtless subjects, or incompetent investigators (Allais, 1990; Amihud, 1979a, 1979b; Morgenstern, 1979). Savage (1954, p. 103) presented the Allais gambles in a way that highlights the independence axiom. In the Savage presentation, the gambles are arranged in a table with the probabilities matched to tickets from a lottery:

Gamble pair 3:			
	Ticket number		
	1	2–11	12–100
A	1000	1000	1000
B	0	5000	1000

Gamble pair 4:			
	Ticket number		
	1	2–11	12–100
A'	1000	1000	0
B'	0	5000	0

This presentation makes the common consequences of each gamble and the relevance of the independence axiom

explicit. By using this presentation, the number of reversals decreases compared with what is seen in traditional presentations from 60% to 40% (Incekara-Hafalir & Stecher, 2012), although they do not entirely disappear.

Alternative representations have been shown to decrease the number of reversals between the Allais gambles. Gamble pairs 5 and 6 have the same structure as the original Allais gambles but with an added value of 9000 to each non-zero outcome (we will call these the “high outcome Allais gambles”):

Gamble pair 5:

A+ : 10000($p = 1$)
 B+ : 10000($p = .89$), 14000($p = .1$), 0($p = .01$)

Gamble pair 6:

A'+ : 10000($p = .11$), 0($p = .89$)
 B'+ : 14000($p = .1$), 0($p = .9$)

The independence axiom still holds in this case; however, the relative difference between non-zero outcomes is smaller. Increasing the magnitude of the outcomes has been shown to reduce but not eliminate reversals (Weber, 2008).

Kahneman and Tversky (1979) developed an alternative version of the Allais gambles and referred to their paradigm as the *common ratio* effect. This format was intended to test Prospect Theory's predictions: the 1% probability of gaining 0 in gamble B of Gamble pair 1 is treated as if it were much more likely than it is (i.e., it is “overweighted”), making the safe gamble A appear more attractive. The common ratio Allais gambles are as follows:

Gamble pair 7

C : 30000($p = 1$)
 D : 60000($p = .5$), 0($p = .5$)

Gamble pair 8

C' : 30000($p = .1$), 0($p = .9$)
 D' : 60000($p = .05$), 0($p = .95$)

Tversky and Kahneman found similar reversals where people prefer C to D, but D' to C' again, representing systematic deviations from EUT.

Researchers have also used Allais reversals to support decision making theories that fundamentally depart from EUT. For example, Regret Theory (Bell, 1982; Fishburn, 1982, 1987; Loomes & Sugden, 1982, 1987) proposes that gambles are analyzed based on an action/state outcome matrix and a choice is made to minimize possible regret. The Proportional Difference model (Gonzalez-Vallejo, 2002) assumes that gambles are compared by their attributes such that a decision is pulled towards gamble B when outcomes of gambles A

and B are compared, and a decision is pulled towards gamble A when probabilities are compared. Both Regret Theory and the Proportional Difference model use psychological mechanisms very different from EUT to account for Allais reversals.

Finally, the use of real monetary payoffs as opposed to hypothetical payoffs has also been shown to reduce some reversals (Burke, Carter, Gominiak, & Ohl, 1996; Harrison, 1994). Based on experimental results, Harrison concluded that monetary incentives and playing out the gambles would influence the number of reversals. Burke et al. (1996) tested the robustness of Harrison’s results and likewise found that monetary incentives reduced but did not eliminate the number of reversals in Allais gambles. Finally, in a comprehensive survey of the effect of monetary incentives in decision making experiments, Camerer (1995) cited Harrison’s Allais experiment as examples of how monetary incentives reduce EUT reversals.

DECISIONS FROM EXPERIENCE AND THE ALLAIS GAMBLES

Decisions from description have been the “fruit-fly” of decision making research (Lopes, 1983) since the seventies and have led to the demonstration of many decision biases, beginning with the Allais Paradox. DFE, on the other hand, were studied by decision scientists in the 1950s and 1960s (i.e., Edwards, 1961; Luce & Suppes, 1965), but fell out of favor in the mid 1970s (Hertwig & Erev, 2009). Recently, decision scientists have renewed their interest in DFE, where systematic and robust differences between decision behavior from experience and from description have been found (*the description–experience gap*; Hertwig et al., 2004; Hertwig & Erev, 2009). This gap suggests systematic differences in choice behavior when relying on description compared with experience (Hertwig, in press).

An experimental paradigm used in DFE is a computerized “money machine” where participants are presented with two buttons, each representing a gamble with an initially

unknown payoff distribution (Barron & Erev, 2003). Clicking on a button results in a random draw from the respective payoff distribution, and the payoff distributions are learned through repeated choices with feedback. Multiple variants of this experimental setup have been used (for reviews of differences between these variants, see Fantino & Navarro, 2012; Gonzalez & Dutt, 2011; Hertwig & Erev, 2009; Lejarraga et al., 2012), but we use a repeated consequential choice paradigm where participants receive the outcome of the selected gamble with no other information (Figure 1). Over the course of many trials (100 trials in the current work), preferences between the two gambles are evident in the proportion of choices made from one gamble over the other.

The *description–experience gap* suggests that in contrast to descriptive choices, people behave as if they underweight small probabilities in experiential choices. For example, when choosing between a safe and a risky gamble with a high or low outcome, participants prefer the risky gamble when the high outcome is likely, but they prefer the safe gamble when the high outcome has a probability around .2 or less. This pattern is reversed when the same gamble pair is presented with full written descriptions (Erev et al., 2010).

Although the Allais gambles have not been tested in DFE, two studies support these expectations. First, Barron and Erev (2003) used a pair of gambles similar to Kahneman and Tversky’s common consequences (gambles 5 and 6 in the introduction) in a repeated consequential choice paradigm. The first gamble pair was a sure thing with a 3000 point outcome and a gamble with a .8 probability of 4000 or 0 otherwise. The second gamble pair had the same outcomes with the probabilities divided by a common ratio (.25 probability of 3000 and .2 probability of 4000). Barron and Erev found that participants’ preferences were consistent across the two gamble pairs, preferring the gamble with the 4000 outcome in each case, while similar choices made from description have consistently shown preference for a smaller outcome with certainty and a reversal when the probabilities are divided by a common ratio (Kahneman & Tversky, 1979).

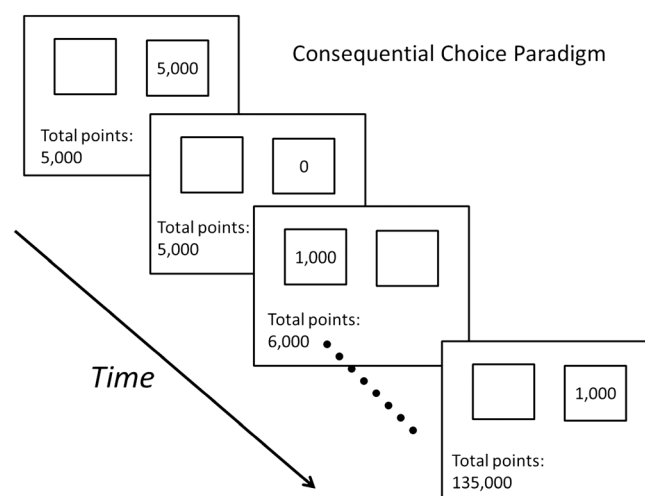


Figure 1. Illustration of the repeated consequential choice paradigm. Participants make repeated choices between two unmarked buttons. The result of each choice is shown and added to a total payoff

Although comparable with the Allais gambles, the stimuli used by Barron and Erev did not include a gamble with three outcomes like in the Allais Gamble pair 1. Research in DFE typically defines a rare outcome as one with a probability of less than .2, and most of the gamble pairs include only one rare outcome (Erev et al., 2010). Current research has not addressed the issue of how multiple rare outcomes may be underweighted. Specifically, in Gamble pair 1, one outcome is identical to the sure thing (1000, $p = .89$) and the other two outcomes are rare in terms of DFE (5000, $p = .1$ and 0, $p = .01$). General underweighting of both of these rare outcomes would predict little or no preference between the gamble pairs. Also, the gambles used by Barron and Erev differ from the Allais gambles in terms of common consequences. Their observation that participants prefer the gamble with a .2 probability of 4000 in the second gamble pair indicates that the common consequence of zero 75% of the time did not unduly influence choices (the first gamble pair uses unique outcomes for both gambles). We extend these findings to the Allais gambles, which are expected to influence people's ability to discount common consequences in DFE, given that the proportion of common consequences in both gambles is higher (.89 and .9 compared with .75) and the common consequences in the first Allais gamble pair is a non-zero positive value.

A second study suggests that experiential information may reduce or eliminate Allais reversals. Van de Kuilen and Wakker (2006) had participants make 15 choices from Allais-type gambles presented in the Savage table format (similar to gamble pairs 3 and 4 earlier). One group made all 15 choices, and at the end of the experiment, one of their choices was played for real with a 20-sided die. A second group was given feedback by rolling a 20-sided die after each choice and writing down the outcome they would receive at the end if they selected that gamble for the real reward. The first group without feedback made choices inconsistent with EUT across all 15 choices. The second group that received feedback showed a *decrease* in reversals as trials progressed. While Van de Kuilen and Wakker showed that receiving feedback can change preferences in Allais gambles, participants were still given descriptive information about the gambles, and it is not clear that learning solely through feedback would have the same effect.

Given these findings, we expect that the Allais reversals would disappear when outcome distributions are learned from experience, presumably largely due to extreme underweighting of rare outcomes (Hertwig et al., 2004) and the high probability of the common outcomes in the two gambles of both Allais pairs (Barron & Erev, 2003).

EXPERIMENTAL STUDY: ALLAIS GAMBLES IN MULTIPLE PRESENTATION FORMATS

One hundred participants (50 of whom were female, mean age = 32.8 years) completed this experiment on Amazon Mechanical Turk (MTurk). MTurk is an increasingly popular online participant recruitment tool that allows for rapid and convenient data collection from a large pool of participants.¹

Participants were restricted to US IP addresses and were compensated after completion of the study through their MTurk accounts.

After filling out demographic information, participants saw the following instructions:

"In today's experiment, you will be making choices between different options presented in multiple ways. Some portions will be choices between gambles or lotteries that could earn hypothetical money. In other portions, you could earn points that will be translated into real money. You will get \$1 for completing the experiment and you could earn up to \$5 more based on your performance."

The experiment was composed of four randomly presented blocks, all within-subjects. Two of the four blocks were DFE representations of the two gamble pairs in the Allais Paradox (Gamble pairs 1 and 2), for which we used the repeated choice paradigm explained earlier (Figure 1). Once a choice was made, the number of points accrued was displayed and added to a total that appeared throughout the block over 100 trials. Participants were paid a bonus based on the number of points they accumulated in the two DFE blocks (one cent for every 1000 points). Before completing each of the DFE blocks, participants saw the following instructions:

"The following portion of the experiment includes 100 trials. Your task, in each trial, is to click on one of the two buttons presented on the screen. Each button press may lead to winning points. Upon pressing a button you will be told how many points you won with that choice, and that number will be added to your total number of points. Your goal in the experiment is to win as many points as possible. You will be paid 1 additional cent for every 1,000 points you earn. So, if you win 150,000 point in this portion of the experiment you will receive \$1.50 in addition to what you have made so far."

The remaining two blocks were decisions from description (DFD): one block presented gambles with described probabilities and outcomes, and the second block presented gambles in a table format consistent with the Savage presentation. In each block, three sets of gambles were presented. A first set of gambles was the original Allais (Gamble pairs 1 and 2, same ones presented in DFE); a second set of gambles was the high outcome Allais gambles (Gamble pairs 5 and 6); and a third set of gambles was the common ratio (Gamble pairs 7 and 8). In each block, an additional gamble pair was included in which one gamble was dominant and clearly better² than the other one. This final gamble pair was added to help us identify whether participants understood the paradigm and were paying attention; thus, it is not included in

¹Paolacca, Chandler, and Ipeirtis (2010) investigated possible concerns with collecting data using MTurk and concluded that US MTurk participants are at least as representative as traditional participant pools, and results of a comparative study using classic experiments in judgment and decision making found no differences in magnitude of effects obtained from MTurk participants and those using traditional subject pools.

²The dominant gamble used was A: 5 ($p = .01$), 10 ($p = .1$), 15 ($p = .89$) or B: 15 ($p = .11$), 18 ($p = .89$). Eighty-four percent of participants correctly identified the dominant option in the descriptive condition, and 94% of participants correctly identified the dominant option in the Savage presentation. Analysis omitting these participants does not change the results, and the participants were included in the reported analysis.

the analysis. Each DFD block was preceded by instructions that included an explanation and example of the respective presentation mode.

Additional unincentivized decisions from experience group

In agreement with traditional research in DFE, participants of the study described earlier were paid a bonus based on the number of points accumulated over the 100 trials in each of the two DFE blocks. Also in agreement with the tradition of research in DFD, however, they were not incentivized for their performance in the DFD blocks. This imbalance may be problematic given some evidence that real monetary incentives may reduce the Allais reversals in DFD (Burke et al., 1996). For this reason, we ran an additional 99 participants group (47 female, mean age = 36.4 years) in the DFE blocks, who were not incentivized with the performance bonus. In addition to providing a fair comparison between DFD and DFE blocks, the unincentivized group allows us to compare the same gamble pairs in DFE with and without incentives and determine whether incentives may have an effect in the reduction of the Allais reversals in DFE.

BEHAVIORAL RESULTS

In the following analyses, we refer to *reversals* as choices that are inconsistent with EUT's independence axiom, such as the traditional findings in the Allais paradox (i.e., choosing A in Gamble pair 1 and B' in Gamble pair 2, or B in Gamble pair 1 and A' in Gamble pair 2). Choice patterns that satisfy the independence axiom are referred to as being *consistent* with EUT (A and A', or B and B'). We first present results from the DFD blocks, where we expected to replicate traditional Allais results, and then compare results from the DFE blocks against the DFD blocks in the original Allais gambles. Finally, we will present detailed analyses of the DFE blocks over time and at the individual level.

Table 1 shows the choice proportion for the three gamble pairs (original, high magnitude, and common ratio Allais problems) in each DFD presentation condition (Description and Savage table), along with the number of participants (out of 100) whose choices were reversals. Consistent with previous research, the mean choice proportion in the traditional descriptive presentation favoring gamble A (.58) in the Allais gambles was significantly higher than gamble A' (.11), shown using a related samples McNemar test ($\chi^2(1)=43.18, p < .01$). This was also true for the high outcome Allais gambles ($\chi^2(1)=40.01, p < .01$) and the common ratio Allais gambles ($\chi^2(1)=13.89, p < .01$). These results indicate the persistence of the traditional findings in the Allais paradox in DFD. The choice patterns reversed regardless of the type of gamble pairs, although the number of reversals is lower for the common ratio gambles (26/100 in Gamble pairs 7 and 8) compared with the other two (49/100 Gamble pairs 1 and 2, and 60/100 in Gamble pairs 3 and 4).

Table 1. For both decisions from description (DFD) presentation types, choice proportions and the number of individuals whose choices were inconsistent with expected utility theory (reversals) are shown

DFD presentation format	Original Allais Gamble pairs 1 and 2			High outcome Allais Gamble pairs 5 and 6			Common ratio Allais Gamble pairs 7 and 8		
	Gamble pair 1 A choices	Gamble pair 2 A' choices	Reversals Number of participants (out of 100)	Gamble pair 3 A+ choices	Gamble pair 4 A+ choices	Reversals Number of participants (out of 100)	Gamble pair 5 C choices	Gamble pair 6 C' choices	Reversals Number of participants (out of 100)
Description	.58	.11*	49	.76	.26*	60	.95	.75*	26
Savage table	.54	.36*	36	.73	.67	26	.89	.85	16

* $p < .05$.

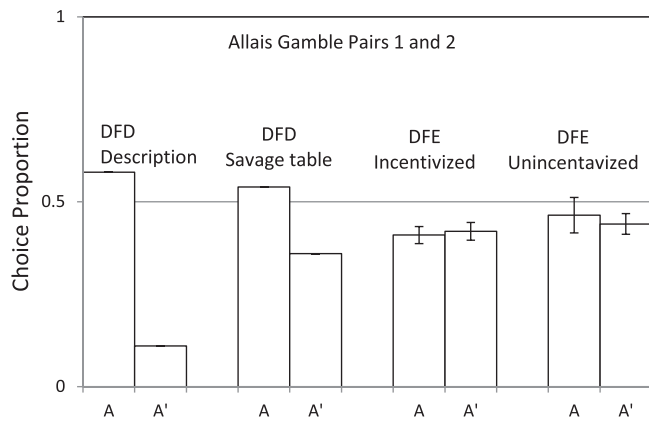


Figure 2. Allais gamble pairs across presentation modes. The choice data for the two Allais gamble pairs are plotted across the three presentation modes. In the traditional and savage presentations, the proportion of participants who chose gamble A and gamble A' are shown. For the experience presentation, the mean proportion of choices across 100 trials is shown. DFD, decisions from description; DFE, decisions from experience

When the gambles were presented using the Savage table presentation, which has been shown to decrease the effect, choice proportions in the original Allais gambles were reversed ($\chi^2(1)=8.03, p < .01$), but not in the high outcome Allais gambles ($\chi^2(1)=.962, p=.327$) or the common ratio Allais gambles ($\chi^2(1)=.562, p=.454$). In other words, the reversal patterns in terms of the mean choice proportion continue to be significant in the Allais gambles but not in the other gambles when presented in the table format. Furthermore, the number of individuals whose choices reversed is reduced in all gambles when they are presented in a table format.

An overall comparison of DFE choices against DFD in the original Allais Gamble pairs 1 and 2 was carried out using the mean proportion of A and A' choices over all participants and across 100 trials for the incentivized ($N=100$) and unincentivized ($N=99$) DFE groups. Figure 2 shows the proportion of A and A' choices in the original Allais gambles in each presentation mode: Description, Savage table, and DFE (incentivized and unincentivized). It is clear that the Allais effect is reduced from the description format to the Savage table and that it disappears in DFE with and without incentives. There was no significant difference in mean choice proportion of gambles A and A' for either the incentivized ($t(99)=.36, p=.72$) or the unincentivized ($t(98)=.605, p=.547$) groups, and these were significantly below .5.³ Thus, at the overall mean level of analysis, preferences in the Allais gambles in DFE are consistent across Gamble pairs, and they suggest preferences for gambles B and B', the ones with the maximum expected value.

Individual reversals in decisions from experience

In DFE, each participant made 100 choices, and thus, determining the reversals at the individual level is less

³Incentivized condition: gamble A in Gamble pair 1, Mean = .41, SD = .24, $t(99)=3.59, p < .01$; gamble A' in Gamble pair 2, Mean = .43, SD = .22, $t(99)=3.32, p < .01$. Unincentivized condition: gamble A, Mean = .46, SD = .22, $t=1.55, p=.124$; gamble A', Mean = .44, SD = .17, $t(98)=3.06, p < .01$.

straightforward than in the DFD representations. As an approximation of the analyses of reversals at the individual level, we averaged the individual choice proportions across 100 trials in Gamble pairs 1 and 2, and we classified participants as consistent or reversal according to the .5 cut-off. For example, if a participant chose gamble B in more than 50 of the 100 trials and gamble B' in more than 50 of the 100 trials, he or she would be classified as consistent. If a participant chose gamble B in more than 50 of the 100 trials and gamble A' in more than 50 of the 100 trials, he or she would be classified as reversal. Figure 3 shows each of the individuals' average A and A' choice proportions for the incentivized and unincentivized groups. Based on the .5 rule, we obtain two consistent groups depicted in Figure 3: A and A' (top right) and B and B' (bottom left); and two reversal groups: B and A' (bottom right) and A and B' (top left).

As observed in Figure 3, the distributions and types of reversal and consistent choices vary widely. Yet at this individual level of analyses, a majority of participants (56%) made consistent choices (44% reversals) in DFE across Gamble pairs. Furthermore, most consistent participants show preferences for gambles B and B' (bottom-left corner); these have the maximum expected value. However, the numbers of individual reversals in DFE are not significantly different than those found in the original Allais Gamble pairs in DFD (incentivized, $\chi^2(1)=2.46, p=.11$; unincentivized, $\chi^2(1)=.181, p=.67$) or the Savage table formats (incentivized, $\chi^2(1)=.08, p=.76$; unincentivized, $\chi^2(1)=1.32, p=.25$) (see Table 1, Original Allais Gamble pairs 1 and 2). To better understand the reversal and consistent choices in DFE, we analyze these individual choices in more detail later while using the IBL cognitive model to explain the process by which reversal and consistent choices emerge from experience.

Analyses of reversal and consistent individuals in decisions from experience

We explore the behavior of participants classified as reversal and consistent in DFE (using the incentivized and

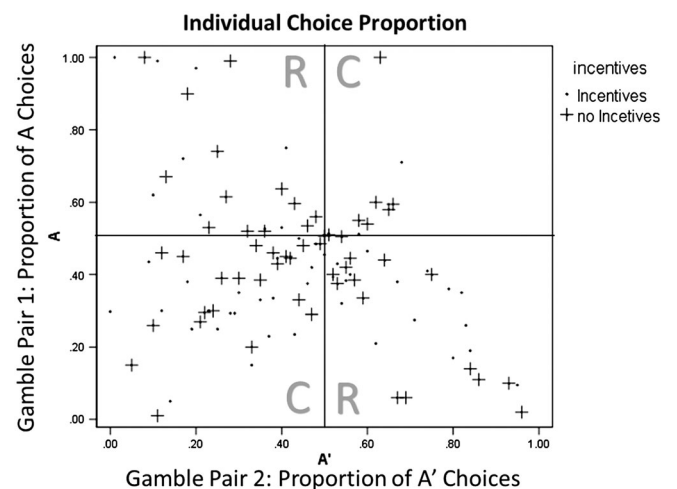


Figure 3. Choice proportions by participant across gambles in the experiential conditions. The overall choice proportions for A and A' are plotted for each individual. C and R indicate quadrants where participant's choices are considered consistent or reversals

unincentivized groups together, because there is no difference between these groups). Table 2 summarizes the main behavioral results. First, we analyzed the proportion of A and A' choices. In both Gamble pairs, consistent participants made significantly less A choices than reversal participants. This confirms the observations that consistent participants select the B and B' gambles more often at the average and individual levels and that consistent participants were able to distinguish the gamble with the maximum EV. Second, we looked at the proportion of alternations (how often a participant switched from one gamble to another), which has been proposed as a measure of exploration (Gonzalez & Dutt, 2011) and a source of individual differences in DFE (Hills & Hertwig, 2010). Overall, on average, there is no significant difference in the amount of exploration that consistent and reversal participants engaged in. In both Gamble pairs, however, the correlation between the proportion of maximizing choices and the amount of exploration is significant for consistent participants and non-significant for reversals. For consistent participants, more exploration corresponded to higher maximization (i.e., the correlation between A and A' choices, and alternation in Gamble pair 1 or Gamble pair 2, respectively, is negative and significant), a relationship previously reported under different paradigms and gambles (Gonzalez & Dutt, 2012).

Given that a major explanation of the differences between DFD and DFE is the underweighting of rare outcomes in DFE, we also analyzed rare outcome experiences for consistent and reversal participants. Table 2 details the experience of rare events in the two gambles and the choice following a rare outcome experience (proportion of switches to the alternative gamble after experiencing a rare outcome). The consistent and reversal groups explored the rare outcomes

with similar frequencies, but their primary difference is found in Gamble pair 2. Compared with reversals, consistent participants are more likely to switch to B' after experiencing a 1000 outcome in gamble A', and they are less likely to switch to A' after experiencing a 5000 outcome in gamble B'. Again, this suggests an exploratory behavior that is led by the search for the gamble with the maximum expected value.

PREDICTED AND OBSERVED DYNAMICS OF CHOICE IN DECISIONS FROM EXPERIENCE

An advantage of the DFE paradigms is the possibility of analyzing how preferences emerge through experience and, more concretely, how participants classified as reversals or consistent develop their preferences in each Gamble pair. The emergence of the aforementioned differences between consistent and reversal participants may be explained with the use of a cognitive computational model of binary choice, the IBL model (Gonzalez & Dutt, 2011), derived from a well-known theory of dynamic DFE, the IBL theory (Gonzalez, Lerch, & Lebiere, 2003). The IBL model has produced very accurate predictions in the absence of human data in a large diversity of choice tasks. It provides cognitive explanations of the learning process that mostly rely on mechanisms such as frequency and recency of experienced outcomes and the desire to maximize experienced expected value (e.g., Gonzalez, 2012; Gonzalez & Dutt, 2011; Lejarraga et al., 2012; Lejarraga & Gonzalez, 2011).

The appendix presents a formalization of the IBL model, which is a reproduction of what has been reported in past

Table 2. Behavioral differences between participants classified as reversal and consistent (incentivized and unincentivized groups have been joined together, significant results in bold)

	Consistent participants (n = 112)	Reversal participants (n = 87)	Statistical test
Gamble pair 1			
Proportion of A choices	.391 (.193)	.497 (.279)	t(197) = 3.16, p < .01
Proportion of alternations	.247 (.216)	.266 (.205)	t(197) = .646, p = .519
Correlation A choice and alternation	r = -.398, p = .002	r = .06, p = .141	
Proportion of participants who experienced 0 outcome in B	.357 (40/112)	.54 (47/87)	
Proportion of A choices after experiencing 0 in B	.392	.517	t(105) = 1.30, p = .196
Proportion of participants who experienced 5000 outcome in B	.875 (98/112)	.908 (79/87)	
Proportion of A choices after experiencing 5000 in B	.232	.192	t(1089) = 1.64, p = .10
Gamble pair 2			
Proportion of A' choices	.390 (.171)	.490 (.232)	t(197) = 3.49, p < .01
Proportion of alternations	.284 (.172)	.284 (.172)	t(197) = .002, p = .998
Correlation A choice and alternation	r = -.558, p < .01	r = .057, p = .566	
Proportion of participants who experienced 1000 outcome in A'	.901 (101/112)	.931 (81/87)	
Proportion of B' choices after experiencing 1000 in A'	.381	.277	t(966) = 3.42, p < .01
Proportion of participants who experienced 5000 outcome in B'	.964 (108/112)	.839 (73/87)	
Proportion of A' choices after experiencing 5000 in B'	.197	.271	t(1073) = 2.87, p < .01

publications (e.g., Gonzalez & Dutt, 2011; Lejarraga et al., 2012). The IBL model stores every choice–outcome association (referred to as an instance) at the moment it is experienced. For example, in the Allais Gamble pair 1, the model stores up to four instances (A-1000; B-1000; B-5000; and B-0) when those outcomes are experienced. Each instance has a value of activation, which reflects the frequency and recency with which each instance has been experienced (see Equation 3 in the Appendix). Every time a selection is made for gamble A, the activation of the instance A-1000 is strengthened, and when it is not chosen, the activation decays. Activation is a concept from the ACT-R cognitive architecture (Adaptive Control of Thought-Rational; Anderson & Lebiere, 1998), and it reflects how readily available an instance is in memory. Each instance has a probability of being retrieved from memory, which is calculated as a function of its activation relative to the activation of all other instances existent in memory (see Equation 2 in the Appendix). At each trial, the model selects the gamble with the highest blended value (BV; i.e., Utility) (see Equation 3 in Appendix). The BV is calculated as an experienced expected value for each gamble, where the outcomes experienced are multiplied by the probability of retrieval.

We ran this IBL model in Gamble pairs 1 and 2 and produced choices for 1000 simulated participants for 100 trials, using values of $d=5$ (decay) and $s=2.5$ (noise) parameters found in past studies. These values of d and s reflect high degrees of recency and noise (compared with the ACT-R default values); recently experienced outcomes have a large effect on the next choice, and the process is noisy.

Averaged over 100 trials, the predicted proportion of choices for gamble A is .30 ($SD=.07$), and the predicted proportion of choices for gamble A' is .39 ($SD=.08$). The model predicts no difference in the proportion of A and A' choices. At the overall mean level of analysis, the model's predictions are in agreement with the mean choice proportions observed from participants, although the IBL model predicts slightly lower proportions than those observed in human participants (.30 for the IBL model instead of .41 and .46 for A, and .39 for the IBL model instead of .43 and .44 for A').⁴

Predicted and observed choices in decisions from experience across trials

Figure 4 shows the mean proportion of predicted and observed choices over the course of experiencing each of the two Gamble pairs in the incentivized and unincentivized groups. The observed proportions of A and A' choices are consistent with the predictions from the model.⁵ Human participants gradually *decrease* their proportion of choices for A and A' over time, building preferences towards B and B'.

Table 3 shows the numerical comparison between IBL and the observed choices. We calculated the mean absolute difference (MAD) between the data and model predictions at both

the group and individual levels, using observed outcomes from humans to provide trial-by-trial IBL predictions for each participant (as in Lejarraga et al., 2014).⁶ The individual-level MAD in Table 3 shows the difference between each observed choice (coded 0 or 1) and the model's prediction for that trial and participant. For each participant, we summed the difference between the predicted choice on each trial (1 or 0) and the observed choice (1 or 0), and compared the result to a random model that predicts .5 on each trial (summarized in Table 3). The mean difference between observed choice proportions and predicted proportions at the group level is .05, and the mean difference at the individual level is .42. Additionally, IBL predicted the observed data better than chance for 72% of the participants (286 of 399). Note that on each measure, IBL performs better for Gamble pair 1 than Gamble pair 2, which is to be expected, as choice behavior in Gamble pair 2 is closer to equal preference between gambles.

Initial overweighting and late underweighting of rare outcomes

Using the IBL model has the advantage of also providing predictions at the process level. Consider Gamble pair 1 as an example. Choices from A produce a single instance in memory (A-1000), and given the formulation of activation and probability of retrieval, the BV for gamble A is near 1000 in every trial (varying only according to noise and initial setting expectations in the model; see Appendix). Therefore, to examine the emergent dynamics of choice in Gamble pair 1, we can focus on gamble B. This gamble produces three unique instances in memory (B-1000; B-5000; and B-0). Disregarding the effects of recency and noise for now, the model would predict indifference between gambles A and B (A-1000 and B-1000), at least until the other two rarer outcomes in gamble B are experienced. When any of the other outcomes in gamble B are experienced, they would produce a momentary *overweighting* of this rare outcome given the recency effect (Appendix, Equation 3), resulting in an effect similar to the well-known hot-stove effect (Denrell & March, 2001). However, the momentary overweighting of rare outcomes may be stronger or weaker, depending on the interplay between frequency and recency and the magnitude of the outcome.

For example, experiencing the very rare 0 outcome for the first time early on will likely increase preference for gamble A, while experiencing the 0 outcome for the first time later on will likely have little effect on preferences. This is because the increase in activation produced by the recency of the 0 would be overwhelmed by the larger activation produced by the greater frequency of 1000 experienced in earlier

⁴It is important to note that model's data are produced in the *absence of human data* and that the model's parameters were not calibrated to the observed choices.

⁵The IBL model does not account for any effects of incentives. Thus, the same predictions were used for the incentivized and unincentivized groups.

⁶The IBL model assumes a deterministic choice rule (the option with the highest BV on a given trial is chosen); however, the specific BV of an option is based on the probability of memory retrieval that is stochastic, modulated by a random draw from a uniform distribution ($\gamma_{j,t}$). To control for this stochasticity, we ran 1000 simulations with each participant's observed feedback to create a choice probability prediction at each trial that was then transformed into a deterministic prediction.

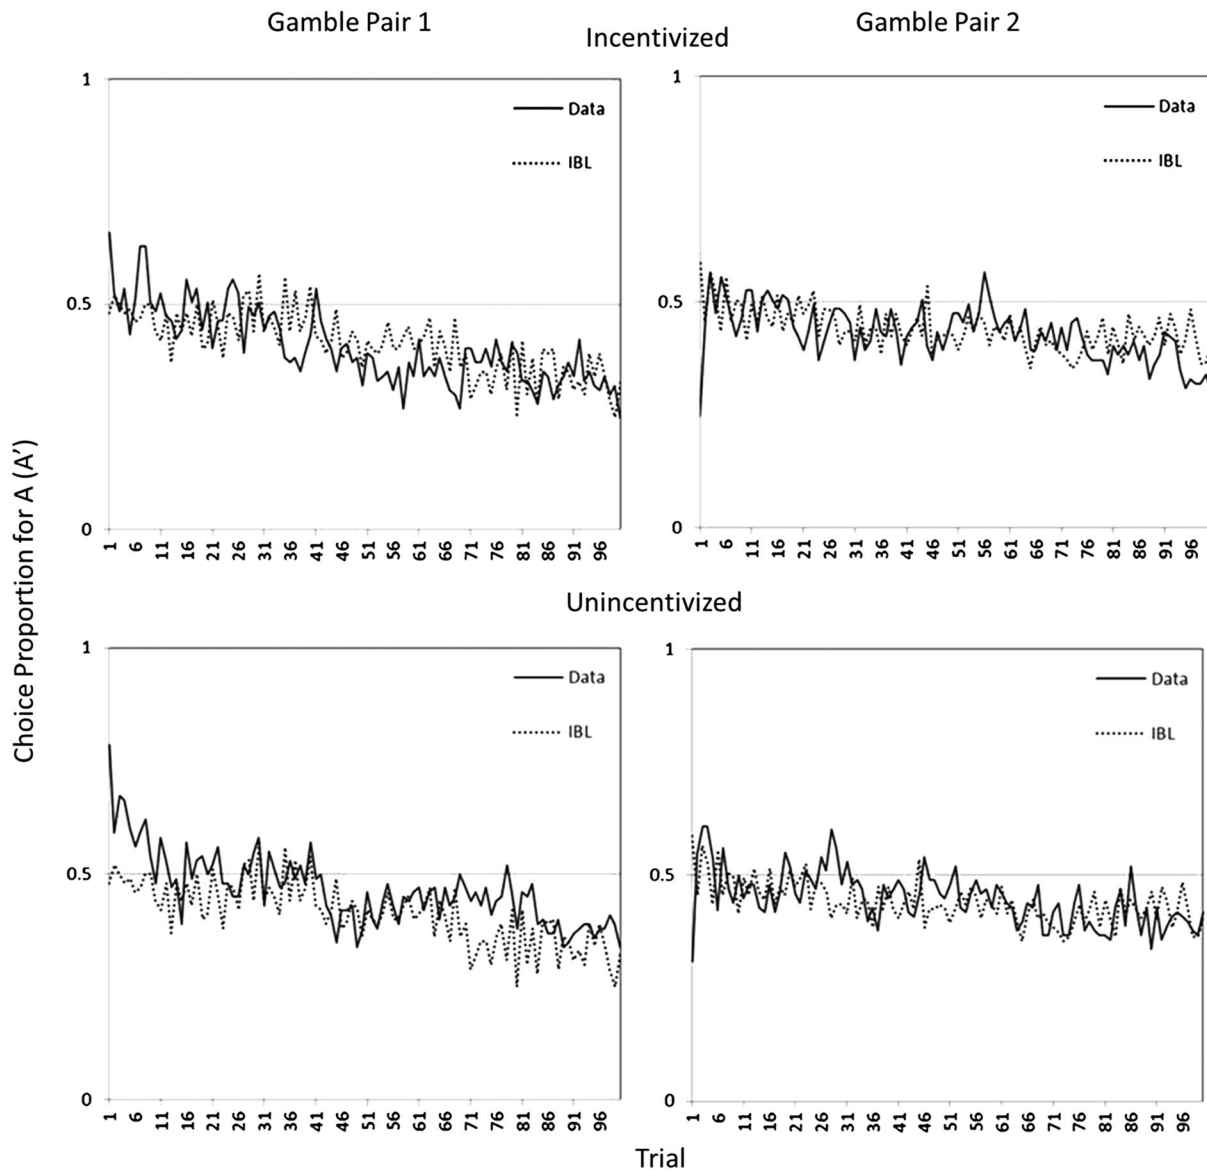


Figure 4. Model predictions for each of the two gamble pairs in the incentivized and unincentivized conditions. The choice proportions are plotted by trial from the data and the Instance-Based Learning model

trials. Thus, with more experience, underweighting of the very rare outcome emerges, and if the 5000 outcome is also experienced, a preference for B becomes more likely. Table 4 lists trial-by-trial data from two simulated participants.

Table 3. Instance-Based Learning (IBL) comparison results

	Incentives (<i>n</i> = 100)		No incentives (<i>n</i> = 99)		Means
	Gamble pair 1	Gamble pair 2	Gamble pair 1	Gamble pair 2	
MAD group	.061	.049	.065	.050	.05
MAD individual	.466	.362	.476	.406	.42
vs. Random (.5)	64	82	68	72	286/ 399

The mean absolute distance (MAD) between the observed data and model predictions for each gamble pair and condition are presented in the first two rows. The group level compares mean proportion per trial, while the individual level compares trial-by-trial choices. The final row lists the frequency of participants' data in which IBL outperformed a baseline model predicting .5 on every trial.

Simulated participant 1 experienced the 0 outcome early on (in Trial 16 left side) and shows a drop in the BV of gamble B, producing a shift in choice to A, which decays over time, given that the 0 outcome is rare. Simulated participant 2 experienced the 0 outcome for the first time on trial 91 (right side) and shows a BV for B that drops only slightly on trial 92 as the frequency of previous experiences overwhelms the recency effect of the 0 outcome. Also, the BVs of each gamble are very close. This similarity in BVs reflects the overlap of outcomes between the gambles (common consequences) and indicates that preferences, while moving towards gamble B over time, will be somewhat similar between A and B, with noise playing a large role in choice on any single trial.

To go beyond a single example, we analyzed a random sample of 100 simulated participants in Gamble pair 1. Of the 100 simulations, 62 experienced the 0 outcome. For these simulations, we correlated the trial number at which the 0 outcome was experienced for the first time and whether the

Table 4. Individual simulation data for Gamble pair 1

Simulated participant 1: Early observation of rare outcome					Simulated participant 2: Late observation of rare outcome				
Trial	BV A	BV B	Choice	Outcome	Trial	BV A	BV B	Choice	Outcome
10	1008.4	1081.2	A	1000	84	1000.1	1005.7	B	1000
11	1037.8	1036.7	B	1000	85	1000.5	1010.8	B	1000
12	1002.4	1012.7	A	1000	86	1000.5	1000.2	A	1000
13	1003.0	1010.1	A	1000	88	1000.8	1217.5	B	1000
14	1007.9	1000.9	B	1000	89	1000.0	1003.5	B	1000
15	1124.7	1001.4	A	1000	90	1000.6	1016.4	B	1000
16	1002.1	1010.6	B	0	91	1000.0	1006.3	B	0
17	1005.4	992.09	A	1000	92	1000.3	1002.2	B	1000
18	1000.5	976.78	A	1000	93	1000.2	1018.8	B	1000
19	1000.2	444.32	A	1000	94	1000.1	1004.2	B	1000
20	1001.1	943.73	A	1000	95	1000.1	1002.2	B	1000
21	1000.2	1007.7	B	1000	96	1000.5	1010	B	1000
22	1000.5	1056.8	B	1000	97	1000.7	1001.9	B	1000
23	1002.4	999.74	A	1000	98	1000.0	1000.2	B	1000
24	1004.7	1019.6	B	1000	99	1000.1	1004.9	B	1000
10	1008.4	1081.2	A	1000	84	1000.1	1005.7	B	1000

For two simulated participants, the table lists the trial number, the blended value (BV) for gambles A and B, the choice on that trial (gamble with the highest BV), and the outcome of that choice (calculated into the BV on the next trial). Simulated participant 1 on the left experienced the rare 0 outcome on trial 16, and simulated participant 2 on the right experienced the rare 0 outcome on trial 91. The trial that resulted in the 0 outcome is bolded for each simulation. The effect of this early or late experience of the very rare outcome is shown in the trials that follow, highlighted in the gray area. The model that experienced the rare outcome early shifts from B to A choices, while the model that experienced the rare outcome later stays with the B choice.

choice in the following trial switched from B to A. Among the 62 simulated participants who experienced the 0 outcome, there was a negative relationship between the trial number in which the 0 outcome was experienced and whether choice switched to gamble A in the next trial, with fewer switches as trials progressed (Spearman’s $\rho(62) = -.295, p = .02$). This prediction is confirmed with the human data. Of the 199 participants (incentivized and unincentivized grouped together), 87 experienced the 0 outcome in Gamble pair 1. The correlation of the trial number where the 0 outcome was first experienced and the switch from B to A was found to be very similar to the one predicted by the IBL model (Spearman’s $\rho(87) = -.257, p < .01$). Thus, the IBL model predicts initial overweighting of rare outcomes experienced and later underweighting of rare outcomes, and decreased preference switches depending on the individual’s experience of rare outcomes, the timing of such experiences, and the magnitude of the rare outcome being experienced.

Performing the same analysis earlier while separating participants into reversal and consistent groups, there is a strong and significant negative correlation between the trial number in which the 0 outcome was experienced and choice switch in the next trial for the reversal group (Spearman’s $\rho(47) = -.429, p < .01$). This correlation was not significant for participants classified as consistent (Spearman’s $\rho(40) = -.121, p = .457$). These analyses suggest different reactions to rare outcomes and different patterns of exploration between consistent and reversal groups. To further support the IBL predictions and aforementioned observations, we used an additional model that can simplistically isolate one

of these mechanisms (e.g., recency). SWIM uses the ζ parameter, which represents a fixed number of most recent past trials, which are then averaged to predict choice (Ashby & Rakow, 2014). SWIM fixes the frequency to ζ and relies only on the recency, as it ignores trials behind the ζ trials.⁷ If the IBL model’s predictions are correct, then SWIM will not be able to predict the differential weighting of rare outcomes in early and late trials, and will not be able to predict the differences between consistent and reversal participants.

We used the ζ parameter (window size) of 7, the average parameter found by Ashby and Rakow (6.84, p. 5). Participants’ choice data were coded 1 (A or A’) or 0 (B or B’), and we calculated SWIM choice predictions on each trial using an average of the outcomes experienced by each participant in the previous seven trials. As expected, data generated from the SWIM model fit the overall choice data as well as the IBL model, but SWIM did not predict the correlations predicted by IBL and found in the data. The correlation between trial number and choice switch was not significant in SWIM’s predictions (Spearman’s $\rho(105) = -.048, p = .63$). Splitting participants into consistent and reversal groups, SWIM does not reliably predict the dynamics found in the data: the reversal group was incorrectly predicted as null for the 0 outcome (Spearman’s $\rho(56) = 0, p = 1$) and negative for the 5000 outcome (Spearman’s $\rho(530) = -.114, p = .008$). For Gamble pair 2, no SWIM predictions reached significance. These

⁷SWIM does not make predictions before ζ trials are experienced or when the mean of ζ previous trials in each gamble is the same. In these cases, we set SWIM predictions to .5 (i.e., equal preference between gambles).

results provide strong additional support for the dynamics of behavior among reversals and consistent participants.

DISCUSSION

We designed a comprehensive study of various formats of the Allais gamble pairs where choices are made based on descriptive information, and added novel behavioral results of choices made from experience acquired with repeated feedback. Our results replicate previous findings, where Allais reversals are found in the traditional descriptive format and reduced when gambles are presented in a descriptive table format (Allais, 1953, 1990; Incekara-Hafalir & Stecher, 2012; Savage, 1954). However, overall results of experiential choices in the classic Allais gamble pairs indicate no difference in the mean choice proportion between the two gamble pairs—a majority of participants make consistent choices from experience. Furthermore, there are interesting behavioral differences between consistent and reversal participants, such as consistent participants tend to choose the gamble with maximum expected value (B and B') more often than do reversal participants.

Our results provide some explanations for how overall choice consistency in the Allais gamble pairs emerges from experience. First, we find that consistency is not due to a difference in monetary incentives provided. The same choice consistency between the two gamble pairs was found for experiential groups even when they were not incentivized. Second, choices that participants make from experience gradually drift towards a preference for the gambles with the maximum expected value. Third, a majority of participants make consistent choices from experience, but the number of individual reversals from experience was not different from the individual reversals in the descriptive formats.

Detailed analyses of individual choices from experience and predictions obtained from the IBL model explain the differences between consistent and reversal participants and their dynamic behavior. Participants that show consistent choices across gamble pairs generally search and explore more to find the gamble with the maximum expected value: the data and the model show a direct relationship between the amount of exploration and maximization selections. The reversal participants, in contrast, direct their exploration less towards maximization and they behave differently towards rare outcomes. For example, our results show how consistent participants are more likely to switch to B' after experiencing a high rare outcome in A' than reversal participants and less likely to switch to A' after experiencing the high rare outcome in B'.

Predictions from the IBL model help to unfold the dynamics of human behavior and to understand the development of preferences over the course of experience (Gonzalez, 2013; Gonzalez & Dutt, 2011; Lejarraga et al., 2012). The model explains choices according to the desire to maximize experienced expected value (Blending, Utility; Equation 1 Appendix) but bounded by the cognitive constraints of frequency, recency, and noise for experienced outcomes and retrieval. We use the IBL model to understand the differences between consistent and reversal participants in terms of their

exploration and maximization processes, particularly in terms of switches and alternations in the presence of rare outcomes. These analyses show the two groups differ given early and late experiences of rare outcomes. The model suggests a tendency to overweight early experiences of rare outcomes and to underweight later experiences, producing initial tendencies to alternate to the other gamble and late tendencies to stick to the current gamble. Although many of the model's predictions need to be verified in controlled human experiments, these observations are important given traditional explanations of the Allais choice behavior. In decision sciences where DFD research dominates, researchers have focused on the overweighting of rare events: whether decision theories explain the Allais choice behavior through a systematic overweighting of small probabilities (such as in PT) or by focusing on the possible regret of each outcome (as in Regret Theory). But current empirical work of experiential choices contrasts descriptive-based choice by suggesting that people behave *as if* they underweight rare outcomes (e.g., Hertwig et al., 2004). We provide more accurate observations of experienced-based choice, which expand upon other studies using cognitive models (Gonzalez & Dutt, 2011, 2012; Gonzalez & Mehlhorn, 2014; Harman, Mehlhorn, & Gonzalez, 2014). People do not seem to generally underweight rare outcomes in experiential-based choice. Rather, overweighting or generally underweighting may occur depending on the timing of those experiences, the magnitude of the outcomes observed, and the general accumulated value of the gambles. Rare events may be overweighted (in agreement with the *hot-stove effect*; Denrell & March, 2001) when the frequency of other experienced outcomes in the gamble are not enough to overcome the recency of the rare outcome. But overweighting diminishes with experience to give rise to the underweighting of rare events.⁸ The interaction between frequency and recency of rare outcomes is an important aspect of DFE emerging from the current work that calls for further investigation of the dynamics of choice (Gonzalez & Mehlhorn, 2014).

ACKNOWLEDGEMENTS

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

REFERENCES

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21(4), 503–546. DOI:10.2307/1907921
- Allais, M. (1979). The so-called Allais' Paradox and Rational Decisions under Uncertainty. In A. Hagen (Eds.), *Expected Utility Hypotheses and the Allais Paradox* (p. 437–699). Dordrecht, Holland: Reidel.

⁸A related effect is melioration, where impulsivity drives choices towards options with higher immediate payoffs. See Stillwell and Tunney (2009) and Tunney and Shanks (2002) for changes in this effect over time.

- Allais, M. (1990). Allais Paradox. In J. Eatwell, M. Milgate, & P. Newman (Eds.), *Utility and probability* (pp. 3–9). Hong Kong: Macmillian Press.
- Amihud, Y. (1979a). Critical examination of the new foundation of utility. In M. Allais, & O. Hagen (Eds.), *Expected utility hypotheses and the Allais' Paradox: Contemporary discussion and rational decisions under uncertainty with Allais' Rejoinder* (pp. 149–160). Dordrecht: Springer.
- Amihud, Y. (1979b). A reply to Allais. In M. Allais, & O. Hagen (Eds.), *Expected utility hypotheses and the Allais' Paradox: Contemporary discussion and rational decisions under uncertainty with Allais' Rejoinder* (pp. 185–190). Dordrecht: Springer.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ashby, N. J., & Rakow, T. (2014). Forgetting the past: Individual differences in recency in subjective valuations from experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(4), 1153–1162. DOI:10.1037/a0036352
- 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
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5), 961–981. DOI:10.1287/opre.30.5.961
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1), 23–36. DOI:10.2307/1909829
- Burke, M. S., Carter, J. R., Gominiak, R. D., & Ohl, D. F. (1996). An experimental note on the Allais paradox and monetary incentives. *Empirical Economics*, 21(4), 617–632. DOI:10.1007/BF01180705
- Camerer, C. F. (1995). Individual decision making. In J. Kagel & A. Roth (Eds.), *Handbook of experimental economics* (pp. 587–704). Princeton, NJ: Princeton University Press.
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5), 523–538. DOI:10.1287/orsc.12.5.523.10092
- Edwards, W. (1961). Behavioral decision theory. *Annual Review of Psychology*, 12, 473–498. DOI:10.1146/annurev.ps.12.0201.61.002353
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., et al. (2010). A choice prediction competition for choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47. DOI:10.1002/bdm.683
- Fantino, E., & Navarro, A. (2012). Description–experience gaps: Assessments in other choice paradigms. *Journal of Behavioral Decision Making*, 25(3), 303–314. DOI:10.1002/bdm.737
- Fishburn, P. C. (1982). Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1), 31–67. DOI:10.1016/0022-2496(82)90034-7
- Fishburn, P. C. (1987). Reconsiderations in the foundations of decision under uncertainty. *The Economic Journal*, 97(388), 825–841.
- Gonzalez, C. (2012). Training decisions from experience with decision making games. In P. Durlach, & A. M. Lesgold (Eds.), *Adaptive technologies for training and education* (pp. 167–178). New York: Cambridge University 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, Netherlands: Elsevier.
- 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
- Gonzalez, C., & Mehlhorn, K. (2014). Gains and losses from experience: Cognitive processes and predictions of risky choice and the framing effect. Unpublished manuscript under review.
- González-Vallejo, C. (2002). Making trade-offs: A probabilistic and context-sensitive model of choice behavior. *Psychological Review*, 109(1), 137–155. DOI:10.1037/0033-295X.109.1.137
- Harman, J. L., Mehlhorn, K., & Gonzalez, C. (2014). *Frequencies in the world, frequencies in memory, and tradeoffs of value in risky choice*. Unpublished manuscript under review.
- Harrison, G. W. (1994). Expected utility theory and the experimentalists. In J. D. Hey (Ed.), *Experimental economics* (pp.43–73). Berlin: Springer-Verlag.
- Hertwig, R. (in press). Decisions from experience. In G. Keren, & G. Wu (Eds.), *Blackwell handbook of judgment & decision making*. Oxford, UK: 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
- 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. DOI:10.1177/0956797610387443
- Incekara-Hafalir, E., & Stecher, J. D. (2012). An experimental test of theories of behavior in Allais-type tasks. Unpublished manuscript. Available at SSRN: 10.2139/ssrn.2179060
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291. DOI:10.2307/1914185
- 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., Lejarraga, J., & Gonzalez, C. (2014). Decisions from experience: How groups and individuals adapt to change. *Memory & Cognition*, 42(8), 1384–1397. DOI:10.3758/s13421-014-0445-7
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368), 805–824.
- Loomes, G., & Sugden, R. (1987). Some implications of a more general form of regret theory. *Journal of Economic Theory*, 41(2), 270–287. DOI:10.1016/0022-0531(87)90020-2
- Lopes, L. L. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance*, 9(1), 137–144. DOI:10.1037/0096-1523.9.1.137
- Luce, R. D., & Suppes, P. (1965). Preference, utility, and subjective probability. In R. Luce, R. Bush, and E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 3, pp. 249–410). New York: John Wiley & Sons.
- Morgenstern, O. (1979). Some reflections on utility. In M. Allais and O. Hagen (Eds.), *Expected utility hypotheses and the Allais' Paradox: Contemporary discussion and rational decisions under uncertainty with Allais' Rejoinder* (pp. 175–183). Dordrecht: Springer.
- Paolacca, G., Chandler, J., & Ipeirtis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5), 411–419.
- Savage, L. (1954). *The foundations of statistics*. New York: Wiley.
- Stillwell, D. & Tunney, R.J. (2009). Melioration behavior in the Harvard game is reduced by simplifying decision outcomes. *Quarterly Journal of Experimental Psychology*, 62(11), 2252–2261. DOI:10.1080/17470210902765999

- Tunney, R. J., & Shanks, D. R. (2002). A re-examination of melioration and rational choice. *Journal of Behavioral Decision Making*, 15(4), 291–311. DOI:10.1002/bdm.415
- van de Kuilen, G., & Wakker, P. P. (2006). Learning in the Allais paradox. *Journal of Risk and Uncertainty*, 33(3), 155–164. DOI:10.1007/s11166-006-0390-3
- Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (Second ed.). Princeton, NJ: Princeton University Press.
- Weber, B. J. (2008). The effects of payout and probability magnitude on the Allais paradox. *Memory and Cognition*, 36(5), 1013–1023. DOI:10.3758/MC.36.5.1013

APPENDIX: INSTANCE-BASED LEARNING MODEL

The Instance-Based Learning (IBL) model (Gonzalez et al., 2003; Lejarraga et al., 2012) evaluates an option (e.g., gamble) according to its blended value (BV).

The BV of option j is

$$V_j = \sum_{i=1}^n p_i x_i \tag{1}$$

where x_i is the value of the observed outcome i and p_i is the probability of retrieval of 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 the observed outcomes k in option j

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

where τ is random noise defined as $\tau = \sigma\sqrt{2}$ and σ

is a parameter fitted to the data. At trial t , the activation (Anderson & Lebiere, 1998) of an outcome i is

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

where d is a decay parameter fitted to the data, $\gamma_{i,t}$ 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.

The IBL model chooses the option with the highest BV, $V_{j,t}$.

Authors' biographies:

Jason L. Harman is a post-doctoral research fellow in the Dynamic Decision Making Laboratory at the Department of Social and Decision Sciences at Carnegie Mellon University. He received his PhD in 2012 in Experimental Psychology from Ohio University.

Prof. Cleotilde Gonzalez is an associate research professor of Decision Sciences and the founding director of the Dynamic Decision Making Laboratory (DDMLab:www.cmu.edu/ddmlab) at Carnegie Mellon University. She is part of the editorial board psychology and decision science journals, a human factors fellow, and a principal investigator in a number of large research grants.

Authors' addresses:

Jason L. Harman and Cleotilde Gonzalez, Dynamic Decision Making Laboratory, Carnegie Mellon University, Pittsburgh, PA, USA.