Information asymmetry in decision from description versus decision from experience

Liat Hadar[*]

The Interdisciplinary Center Herzliya

Craig R. Fox

University of California at Los Angeles

Abstract

In this paper we investigate the claim that decisions from *experience* (in which the features of lotteries are learned through a sampling process) differ from decisions from *description* (in which features of lotteries are explicitly described). We find that the experience-description gap is not as robust as has been previously assumed. We argue that when this gap appears it is driven to a large extent by asymmetries in information concerning which events are possible and which are certain. First, we find that, when experience-based decision makers sample events without error and then are told what outcomes are associated with each possible event, they are risk seeking for low-probability gains and risk averse for high-probability gains, as in description-based decision making. Second, we find that the experience-description gap for low-probability outcomes appears when rare outcomes are never experienced but disappears when: 1) all distinct outcomes are experienced at least once or 2) never-experienced outcomes are described as possibilities. Third, we find that the experience-description gap for high-probability outcomes is pronounced when decision makers previously experience lotteries that both offered the possibility of a zero outcome (which presumably makes them doubt that an always-experienced outcome is certain), but disappears when they have not previously experienced such lotteries.

Keywords: decision from experience, experience-description gap, uncertainty, risk, information asymmetry.

1 Introduction

A major thrust of behavioral decision research has been to identify situations in which people act boldly or timidly in the face of risk and uncertainty. In studies of decision under *risk* people choose between chance gambles with known probabilities and outcomes (e.g., receive $100 with probability .5 versus $30 for sure). These studies have found that people typically exhibit a fourfold pattern of risk attitudes, seeking risk for low-probability gains and high-probability losses (e.g., most prefer a .01 chance of $100 over $1 for sure) and avoiding risk for high-probability gains and low-probability losses (e.g., most prefer $99 for sure over a .99 chance of $100), as predicted by prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992).[1] In studies of decision under *uncertainty* participants choose between prospects contingent on natural events (e.g., receive $100 if the home team wins versus $30 for sure). These studies don't

generally lend themselves to analysis of "risk" attitudes because they rely on singular natural events for which there are no "objective" probabilities. However, when analyzed relative to *judged* probabilities, several studies reveal a similar pattern of risk-preferences (Tversky & Fox, 1995; Fox & Tversky, 1998; Wu & Gonzalez, 1999; Kilka & Weber, 2001; Fox & See, 2003).

For many decisions, potential outcomes are not known in advance but must be learned from experience. For instance, a commuter might choose a preferred route to work only after sampling alternative routes on several occasions. In an influential paper, Hertwig, Barron, Weber and Erev (2004; henceforth HBWE) modeled such situations by asking participants to learn about a pair of lotteries (e.g., A: gain 32 points with probability .1 and gain 0 otherwise; B: gain 3 points for sure) by sampling outcomes (with replacement), as many times as they wished, from unlabeled buttons associated with these payoff distributions (see Table 1 for a list of all lottery pairs). After this sampling phase participants chose a lottery to play once for real money. Such experience-based decision making (henceforth EBDM) appears to reverse the fourfold pattern of risk attitudes of description-based decision making (henceforth DBDM), with people avoiding risk for low-probability gains and high-probability losses and seeking risk for high-probability gains and low-probability losses (compare columns 5 and 6 of Table

[1]Risk-aversion is defined as a preference for a sure outcome over a risky outcome with equal or greater expected value; risk seeking is defined as a preference for a risky outcome over a sure payment equal to or greater than its expected value.

Table 1: Percent of participants choosing the H (high expected value) option in Hertwig, et al. (2004). The first column indicates decision problem number. The second and third indicate the high expected value and low expected value lotteries, respectively (the first number in each column refers to the possible nonzero outcome and the second number refers to its objective probability). The fourth column characterizes the nature of the risky lottery or lotteries. The final two columns indicate the percentage of participants choosing the high expected value option in description-based decision making and experience-based decision making conditions, respectively.

Decision	Options				
problem	H	L	Risky lottery type	DBDM	EBDM
1	4, .8	3, 1.0	High-probability Gain	36	88
2	4, .2	3, .25	Low-probability Gain	64	44
3	−3, 1.0	−32, .1	Low-probability Loss	64	28
4	−3, 1.0	−4, .8	High-probability Loss	28	56
5	32, .1	3, 1.0	Low-probability Gain	48	20
6	32, .025	3, .25	Low-probability Gain	64	12

1; see also Weber, Shafir & Blais, 2004, for a replication of the reversal of the fourfold pattern of risk attitudes in decisions from experience). On the basis of such results these authors have called for "two different theories of risky choice" (HBWE, p. 534).

The generalization that EBDM differs from DBDM is difficult to evaluate because, surprisingly, no one has yet defined "experience-based decision making." From our reading of the literature and correspondence with several major contributors[2] we conclude that these authors are referring to a single decision in which: (1) decision makers' knowledge of possible outcomes and/or respective probabilities is incomplete, and (2) this information is derived, at least in part, from a sampling process.[3] According to this definition EBDM applies to any situation in which there is uncertainty and learning through sampling. However, most of the data supporting the putative experience-description gap contrast a single DBDM paradigm (decision under risk) with minor variations of a single EBDM paradigm (the HBWE design). Note that decision under risk and the HBWE paradigm vary not only in terms of *how* information is learned (explicit description versus sequential sampling) but also *what* infor-

---

[2]Email correspondence with Ido Erev (received 5/14/07 and 5/17/07), Greg Barron (received 6/11/07), Elke Weber (received 6/11/07) and Ralph Hertwig (received 6/13/07).

[3]Note that experience-based decision making should not be confused with "small feedback-based decision making" (e.g., Barron & Erev, 2003) in which participants receive the outcome of several choices made in succession. Because most models of decision under risk and uncertainty, including prospect theory, are designed to accommodate a single decision made in isolation rather than in series, we confine the present discussion to experiments in which a single choice is made for each lottery pair. We elaborate on this distinction in the general discussion.

mation is available to decision makers (i.e., information sampled from experience may diverge from information describing the corresponding objective probability distribution over outcomes). According to the *information asymmetry* hypothesis, the experience-description gap is driven by differences in information available to decision makers, and will diminish or disappear when such differences are eliminated. The purpose of this paper is to test the information asymmetry hypothesis and better circumscribe conditions under which decisions from experience differ from decisions from description. In so doing we hope to better understand psychological factors that contribute to the experience-description gap when it occurs.

1.1 Reanalysis of Fox and Tversky (1998)

To see how information asymmetry may contribute to the experience-description gap, we first ask whether this gap will persist when such asymmetry is minimized. In particular, we examine a traditional decision under uncertainty paradigm, in which participants decide among prospects whose outcomes depend on events that might occur, with the probabilities of these target events learned through error-free sampling. This paradigm qualifies as "decision from experience" according to the definition above (there is uncertainty and learning through sampling) but the information presented is more equivalent to decision under risk (i.e., all necessary information to determine the objective probability distribution over outcomes is available).

Fox and Tversky (1998, Study 2; henceforth FT) asked participants to observe the direction in which economic indicators moved (up or down) relative to the previous
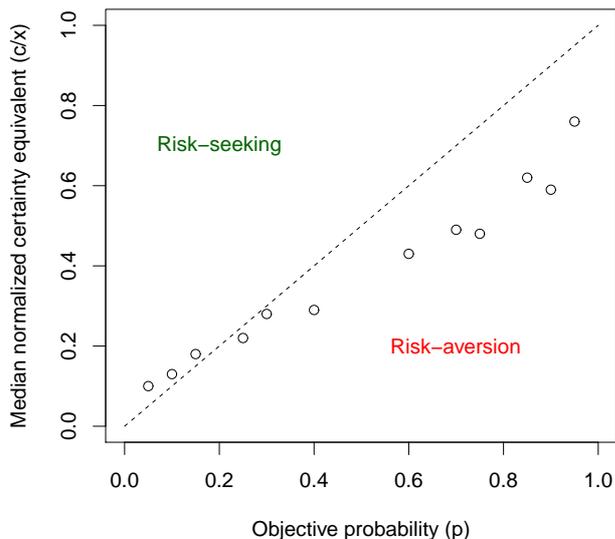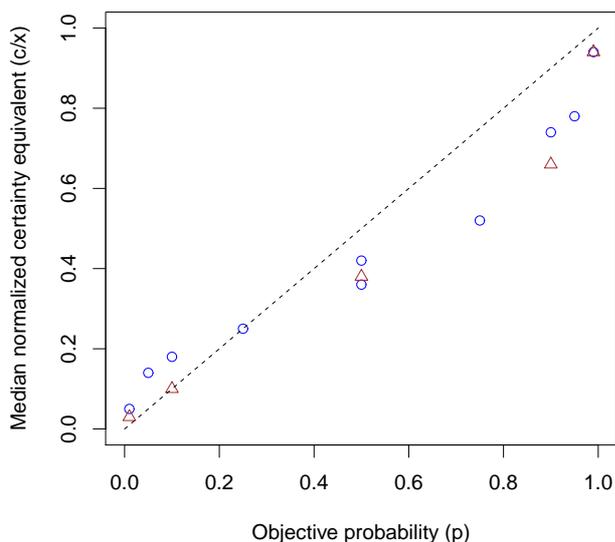
Figure 1: A. Normalized certainty equivalents as a function of corresponding objective probabilities based on a reanalysis of Fox & Tversky (1998, Study 2).

B. Median normalized certainty equivalent for positive prospects of the form (x, p; 0, 1−p) (where x is the outcome of the gamble and p is its probability) as a function of corresponding objective probability for decision under risk. Triangles and circles, respectively, correspond to values of x that lie above or below 200. (Adapted from Figure 1 from Tversky & Kahneman, 1992).

quarter in a hypothetical economy, over sixty quarters. Participants were informed that the probabilities of target events remained the same for each quarter. Next, they were asked to choose between prospects that offered $1600 if indicators moved in a particular direction (up or down) during the subsequent quarter (e.g.,

gain $1600 if inflation goes up) versus various certain amounts of money. From these responses they determined for each prospect its "certainty equivalent," that is, the sure amount of money that was deemed equally attractive to the prospect. Figure 1A displays the median certainty equivalent divided by the $1600 prize for each prospect (i.e., the median "normalized certainty equivalent"), c/x, plotted against the corresponding objective probability. This reanalysis provides an index of median risk attitudes for prospects with varying objective probabilities. The figure shows that participants tended to be risk seeking (points above the identity line) for low-probability gains and risk averse (points below the identity line) for high-probability gains, echoing the pattern commonly observed in studies of decision under risk (i.e., from description, e.g., Tversky & Kahneman, 1992), as reproduced in Figure 1B.

The foregoing analysis contradicts the generalization that EBDM necessarily differs from DBDM and raises the question: Under what conditions do risk preferences in decision from experience differ from those typically observed in decision under risk? Comparing the traditional decision under uncertainty with learning paradigm (employed by FT), in which risk preferences accord with the conventional pattern, and the dominant EBDM paradigm (introduced by HBWE), in which this pattern is reversed, suggests at least two important differences that may account for previous observations of an experience-description gap: (1) FT participants sampled *without replacement* the entire distribution of events so that they were guaranteed unbiased information about outcomes and probabilities, whereas HBWE participants sampled *with replacement* and therefore were likely to obtain a biased sample; (2) FT participants sampled a distribution of *events* (whether inflation and/or interest rates went up or down) and then were told associated outcomes (e.g., "gain $1600 if inflation goes up"), whereas HBWE participants directly sampled a distribution of *outcomes* (e.g., "3 points"). We explore the impact of each of these factors in turn.

Sampling with replacement versus without replacement can make a difference due to sampling error: small samples of binary events that are drawn with replacement tend to under- (over-) represent very rare (common) events. For instance, most participants sampling from the lottery (32,.025) that offers a 2.5% chance of 32 (Problem 6) never observed the 32 outcome. In a reanalysis of HBWE's data we (Fox & Hadar, 2006) examined choices as a function of the probability distribution over outcomes that participants actually sampled (rather than so-called "objective probabilities" that they never directly observed) and found that choices accord with DBDM. Subsequent studies found that the experience-description gap substantially diminishes with larger (less

error-prone) samples (Hau, et al., 2008) and when the EBDM condition is compared to a DBDM condition that describes the probability distribution that was actually sampled (Rakow, Demes & Newell, 2008). Likewise, Ungemach, Chater & Stewart (2009) found that the experience-description gap diminished markedly when participants sampled each of the HBWE lotteries forty times without replacement (so that there was no sampling error), though it did not disappear completely. All of these results suggest that sampling error substantially contributes to the experience-description gap, consistent with information asymmetry hypothesis.

As for sampling events versus outcomes, this factor has not yet been investigated. Note that in the traditional EBDM paradigm (as in HBWE) participants do not learn about all of the distinct possible outcomes unless they sample them, whereas in the traditional decision under uncertainty with learning paradigm (as in FT) participants are explicitly told that they are choosing between prospects that offer the possibility of particular outcomes (e.g., $1600 if inflation goes up and 0 otherwise), which are sometimes described as certain (e.g., $1200 for sure). Thus, the information asymmetry hypothesis predicts that decisions from experience will coincide more closely with decisions from description when: (1) a never-experienced outcome is explicitly described as a possibility, and (2) an always-experienced outcome is described as certain. However, sampling events versus outcomes should not otherwise make a difference on decisions because information provided about the probability distribution over outcomes is otherwise the same.

2   Study 1: Information asymmetries involving never-sampled outcomes and events

In Study 1 we independently manipulate: (a) what information is sampled (*events* versus *outcomes*) and (b) completeness of information described (*incomplete* information in which we mention only the outcomes/events that were sampled versus *complete* information in which we mention all possible outcomes/events whether or not they were sampled). We predict that the experience-description gap, present when rare outcomes/events are not sampled or described, will be significantly attenuated when all outcomes/events are sampled, or when they are all described. We predict further that choices will not be significantly affected by whether participants sample outcomes or events.[4]

---

[4]We note that in a previous study, Erev et al. (2008) contrasted a "blank" EBDM condition in which participants sampled freely from lotteries by clicking two unlabeled buttons, with a "mere presentation"

2.1   Method

We randomly assigned 111 UCLA students to one of four variations of the HBWE paradigm in a 2×2 factorial design (sampled information: outcomes vs. events; information completeness: complete vs. incomplete information). In the *outcome sampling* conditions participants randomly sampled outcomes (e.g., 4 or 0) in the learning phase. In the *event sampling* conditions participants randomly sampled symbols (e.g., star, circle) in the learning phase and were later told the outcomes associated with each shape (e.g., star = 4, circle = 0).

On each trial participants sampled outcomes or events 10 times from each of two unlabeled lotteries in any order they wished (the same lottery pairs used by HBWE — see Table 1), then chose which lottery they would like to play once if real money were at stake. Setting the number of outcomes/events sampled to 10 for each lottery (20 for each pair) holds the amount of information sampled constant across experimental conditions, and also allows for substantial sampling error as has been observed in previous EBDM studies that allowed participants to sample as many times as they wished.[5] During the choice phase in the *incomplete information* conditions, buttons were labeled only with the outcomes/events that a participant had sampled from each button; in the *complete information* conditions buttons were labeled with all possible outcomes/events regardless of which had been sampled.

2.2   Results

Combining information conditions, we found no significant differences in choices under the *event* sampling and *outcome* sampling conditions for any of the six decision problems ($p > .4$ in each of six two-tailed Fisher's exact tests), as predicted.

Note that, from the perspective of participants, the *incomplete* information conditions differ from the *complete* information conditions only when a rare event is never

---

condition in which participants sampled from lotteries by clicking buttons that were labeled with all possible outcomes. While they found that button labeling led to a shifting of choices in the direction of the traditional DBDM pattern, their design does not allow a clean test of the information asymmetry hypothesis because possible outcomes were presented throughout the sampling phase, and the sample size varied systematically with the information provided — participants in the "mere presentation" condition took larger samples from the button associated with the risky lottery but fewer draws from the button associated with a sure outcome, in comparison to the "blank" condition (see their Table 3). Also, this study used only variations of HBWE problems (1) and (2) for which few participants in the "blank" condition would never experience an outcome and therefore receive unique information about what outcomes were possible. Thus, the study speaks more directly to the question of whether participants know that an always-experienced outcome is certain, which we will take up in Study 2.

[5]For example, the median number of draws taken from each pair of lotteries was 15 in Hertwig et al. (2004), 19 in Ungemach, et al. (2009), and the average number was 20 in Weber et al. (2004).

Table 2: Internal Analysis of Study 1 for the three lottery pairs in which at least 20% of participants did not sample all possible outcomes. The first two columns indicate the high and low expected-value options. Columns 3 and 4 indicate which distinct outcomes participants actually sampled from the high and low expected-value options. The remaining columns indicate the number of participants in each condition who experienced the corresponding outcomes and the percentage of these participants choosing the high expected value option.

				Sampled events (shapes)				Sampled outcomes ($ amounts)			
				Incomplete information (n=31)		Complete information (n=27)		Incomplete information (n=30)		Complete information (n=23)	
H	L	Outcomes sampled									
		H	L	n	%H	n	%H	n	%H	n	%H
−3, 1.0	−32, .1	−3	0	8	13	11	46	8	0	15	60
		−3	0, −32	23	65	16	75	22	73	8	63
32, .1	3, 1.0	0	3	11	0	6	66	11	9	13	62
		0, 32	3	20	85	21	62	19	63	10	70
32, .025	3, .25	0	0, 3	21	5	21	62	21	14	13	77
		0, 32	0, 3	9	89	5	80	6	100	8	88

sampled. Thus, we restrict the rest of our analysis to decision problems that include very low-probability events so that a nontrivial proportion of participants (at least 20%) never experienced at least one of the possible outcomes (problems 3, 5 and 6).[6] Table 2 lists the proportion of participants choosing the higher expected value lottery in each experimental condition, separately for each set of distinct outcomes that a participant sampled. This internal analysis reveals that choices were not affected by whether participants sampled outcomes or sampled events and then had the outcomes described to them ($p > .10$ for each of twelve two-tailed Fisher's exact tests), echoing the results of the combined analysis mentioned above. More interestingly, every decision problem reveals a dramatic pattern in which: (1) Most participants who experienced all possible outcomes/events chose the high expected value lottery (consistent with DBDM; $p < .05$ by sign test for every decision problem); (2) Most participants who never experienced the lowest-probability outcome/event and never had it described to them chose the low expected value option (consistent with HBWE, $p < .01$ by sign test in each case); (3) Participants who never sampled the lowest-probability outcome/event but were informed that it was a hypothetical possibility were

much more likely than those not told about it to choose the high expected value option (consistent with DBDM, $p < .002$ in each case by Fisher's exact test).

Combining responses of participants who sampled all possible outcomes/events (i.e., participants for whom the complete and incomplete information conditions coincide), the proportion of participants choosing the high-expected value option for each relevant decision problem is displayed in Figure 2. These results support our hypothesis that describing possible outcomes makes a difference only when these outcomes have never been sampled so that the description provides new information. Importantly, merely informing participants that never-sampled outcomes are a hypothetical possibility causes these outcomes to have as much impact on choices as if they had actually been sampled.

2.3 Discussion

The results of Study 1 support the information asymmetry hypothesis for never-sampled events and outcomes. Thus, the tendency to "underweight" low-probability outcomes in EBDM can derive from not considering them as a possibility; when participants are informed that these outcomes are possible then responses accord with DBDM.[7]

---

[6]Overall results for Problems 1 and 4, which involve moderately high-probability events, qualitatively replicate the dominant pattern of choices reported for EBDM by HBWE. Study 2 speaks to these results. Results for Problem 2, which involves two moderately low-probability events, in fact accord with the usual pattern found in DBDM. This is not surprising as HBWE's trend in the opposite direction was not statistically significant and has not consistently been replicated by other researchers.

[7]Naturally, the present data do not allow us to distinguish the extent to which the impact of information asymmetry on choice is due to differences in the judged likelihood that these events will occur versus differences in the weighting of those probabilities (see Fox & Tversky, 1998, for the distinction between probability judgment and probability weighting).
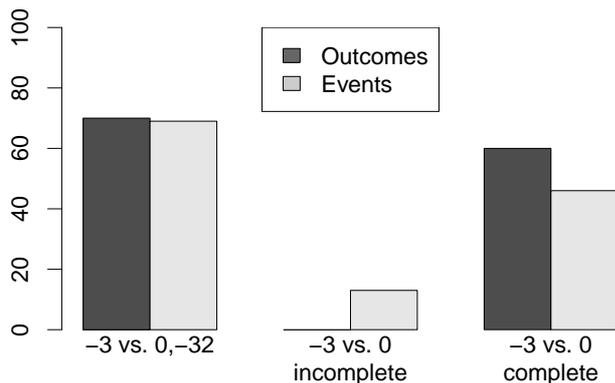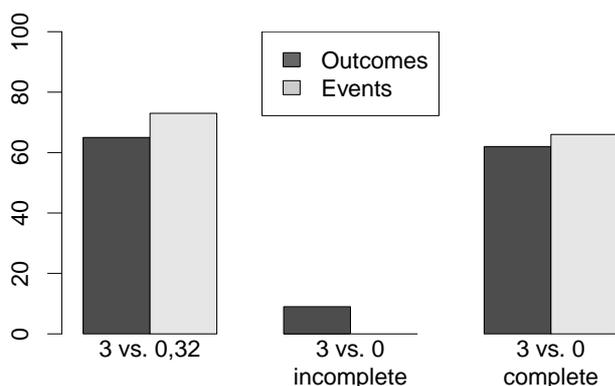
Figure 2: A. Percentage of participants in Study 1 choosing H: $(-32,.1)$ over L: $(-3,1)$ in the conditions in which all outcomes/events were sampled, the least probable outcome/event was never sampled or described, and the least probable outcome/event was never sampled but was described, respectively, for the outcome sampling and event sampling conditions.

B. Same analysis as Figure 2A for H: $(3,1)$ over L: $(32,.1)$
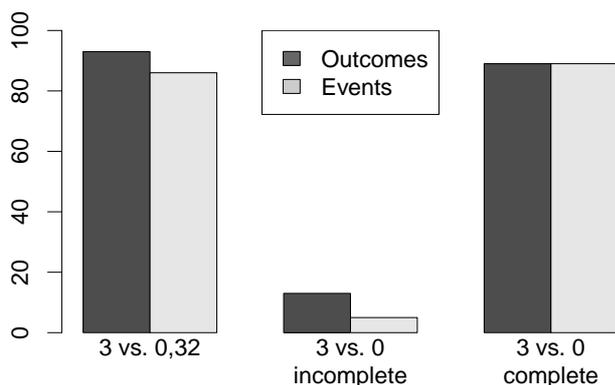
C. Same analysis as Figure 2A for H: $(32,.025)$ over L: $(3,.25)$

The observation that decisions are not affected by whether one samples events or outcomes is important because models of experience-based decision making typically assume sequential learning of sampled outcomes, which is impossible if one instead samples events and learns of associated outcomes only after sampling. For example, the value-updating model (developed by Hertwig et al., 2006, to account for the results of Hertwig et al., 2004) assumes that decision makers update their estimates of the value of an option after each new draw from that lottery by computing a weighted average of the previously estimated value and the value of the most recently experienced outcome. Similarly, the fractional-adjustment model (March, 1996), invoked by Weber et al. (2004), assumes reinforcement learning in which the initial propensity to choose each of the two options in each decision problem is 0.5 and is updated after each draw based on the magnitude of the sampled outcome and its valence: it increases after a favorable outcome but decreases after an unfavorable outcome. Neither model can readily accommodate the results of our event sampling conditions.

3 Study 2: Information asymmetries concerning always-sampled outcomes

Describing possible outcomes can add information not only when an outcome has never been experienced (as in Study 1) but also when an always-sampled outcome is explicitly described as certain. In a previous study, Erev et al. (2008, Study 2) assigned participants to a standard HBWE condition with unlabeled buttons or to a modified condition in which the buttons were labeled with all possible outcomes (e.g., "4 or 0" versus "3"). Thus, participants in the labeled condition presumably interpreted the "3" outcome as certain because it was the only outcome mentioned. Indeed, when the buttons were unlabeled, only 28% of participants chose the button that always paid 3 (as with HBWE, problem 1); however, when the buttons were labeled, 60% preferred the button that always paid 3 and was labeled only with a 3.[8]

The foregoing result suggests that preference for an always-sampled outcome is enhanced when it is explicitly labeled as certain. This information asymmetry between EBDM and DBDM should be especially conse-

---

[8]Note that this finding suggests the choice of (4,.8) over (3,1) in traditional EBDM studies can be reconciled with prospect theory if outcomes that have always been experienced are nevertheless believed to be (slightly) less than certain (say, a 90–95% chance of occurring). For instance, assuming prospect theory parameters from Tversky & Kahneman (1992), a typical participant would prefer an 80% chance of receiving $4 to a 94% chance of receiving $3.
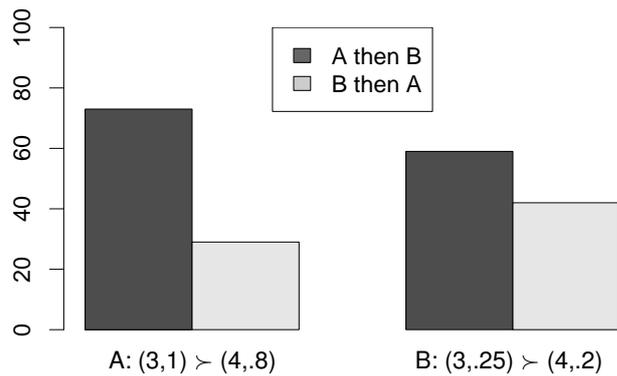
Figure 3: Results of Study 2. Percentage of participants who chose the low expected value option in each decision problem, displayed by the order in which decision problems were presented.

quential when decision makers' prior experience primes them to treat always-sampled outcomes as less than certain. Such might be the case, for instance, if a decision maker has previously encountered trials in which all lottery choices offer both nonzero and zero outcomes. We test this possibility in Study 2.

3.1 Method

We presented 68 UCLA students with two decision problems, (A) (4,.8) vs. (3,1), (B) (4,.2) vs. (3,.25), in one of two orders (A then B or B then A). Before making a choice, participants sampled 10 times from each unlabeled lottery in any order they wished. We predicted that participants would be less likely to choose the sure option "3" in Decision A if they had previously encountered Decision B in which both options offered both the possibility of a zero outcome.

3.2 Results

The first two bars of Figure 3 present the percentage of participants who chose the (3,1) lottery over the (4,.8) lottery for Decision A in Study 2, listed by the order in which decision problems were presented. The results accord with our prediction: although 73% of participants who were first presented Decision A chose (3,1) over (4,.8), which matches the pattern commonly found in DBDM, only 29% of participants who were presented Decision A after Decision B chose (3,1) over (4,.8), ($p < .001$ by Fisher's exact test). This result is consistent with our interpretation that the apparent pattern of "risk seeking" for high-probability gains in EBDM might be driven partly by the tendency to treat always-experienced outcomes as less than certain.

3.3 Discussion

Results of Study 2 are consistent with the information asymmetry hypothesis. An alternative interpretation is that the difference in the (expected sampled) probabilities .8 and 1 appears smaller after completing Decision B, because they are considered in the context of a wider range of probabilities — from as low as .2 to as high as 1 (as predicted by the "decision by sampling" model of Stewart, Chater and Brown, 2006). Thus, probability differences would receive less weight relative to outcome differences when Problem A comes second. Note that this account also predicts that the difference between the (expected sampled) probabilities .20 and .25 in Problem B would also appear smaller when viewed after completing Problem A. In fact, there was a nonsignificant trend in the opposite direction, as can be seen in the last two bars of Figure 3: 42% of participants chose the (3,.25) lottery over the (4,.2) lottery when it was presented first, but 59% chose the (3,.25) over (4,.2) when it came second ($p = .17$ by Fisher's exact test). Thus, our result cannot be explained by range effects as in the decision by sampling model.

4 General discussion

In this paper we explored conditions under which decisions from experience differ from decisions from description. We defined "decisions from experience" broadly as any situation in which a person makes a single decision based on incomplete information about the probability distribution over outcomes that is learned, at least in part, through some sort of sampling process. Using this definition we presented an example in which decisions from experience accord with decisions from description (reanalyzing Fox & Tversky, 1998). Contrasting paradigms, we argued that a major source of the putative experience-description gap is informational asymmetries between the description-based decision task (usually risk) and experience-based decision task (usually variations of HBWE). In two studies we provided evidence that experience-based decisions are especially likely to diverge from description-based decisions when: (1) never-experienced outcomes/events that might occur are never described as possibilities and/or (2) context or prior beliefs stir some doubt that an outcome/event that has always been experienced will necessarily occur. Both of these conditions require at least the possibility of sampling error — so that rare events are not necessarily experienced and/or certain events are not necessarily believed to be certain.

We have confined our attention in this paper to paradigms in which a single choice is made for each lottery pair. We did this because most models of decision

under risk and uncertainty, including prospect theory, are designed to accommodate a single decision made in isolation rather than in a series. We note, however, that an experience-description gap has also been found in a related paradigm called "small feedback-based decision making," (Barron & Erev, 2003; see also Barron, Leider & Stack, 2008; Jessup, Bishara & Busemeyer, 2008; Wu, Delgado & Maloney, 2009). In a typical small feedback-based decision experiment, participants do not go through a sampling phase then decide once, but rather make a large number of repeated decisions (usually 100 or more), each involving a very small amount of money (usually pennies or less). Following each push of a button, participants observe the outcome of their decision, which is then added to their accumulated payoff counter. For example, in one study (Barron & Erev, 2003, Experiment 2) participants were asked to make 400 decisions. They were presented with two buttons associated with either (A) *80% of winning 4 points and nothing otherwise* or (B) *3 points for sure*, where aggregate points were later exchanged for a small amount of money (typically, a total of a few dollars). In this case choices alternated quite a bit, but on average participants chose (A) most of the time, the opposite preference normally observed when people make a single decision under risk involving the same (explicitly described) probability distributions over outcomes (e.g., Kahneman & Tversky, 1979).

The experience-description gap observed in later stages of small feedback-based choice cannot readily be attributed to information asymmetries because this paradigm involves a very large number of draws for which sampling error is likely to be minimal — participants eventually sample all possible outcomes and there is opportunity to develop confidence that always sampled outcomes are certain. However, the small feedback-based decision paradigm introduces a number of complicating factors that make it extremely difficult to ascertain the source of the experience-description gap. Among them: small monetary consequences for each of a very large number of choices which may encourage experimentation and/or lead to boredom and random responses; feedback on the total amount earned which may lead to "house money" effects (Thaler & Johnson, 1990); some participants may invoke lay theories of sequence such as negative recency (the "gambler's fallacy," Lee, 1971) or positive recency (the "hot hand," Tversky & Gilovich, 1989). Another possibility is that some participants may engage in probability matching, where the proportion of times an alternative is selected trends toward the proportion of times in which this alternative provides the best outcome (Estes, 1950). Testing such possibilities is beyond the scope of the present paper but might be explored in future research.

We assert that traditional description-based models generally perform well in decisions from experience when they account for a decision maker's subjective representation of a decision problem. For instance, before taking a long trip a driver may seem to "underweight" and/or "underestimate" the possibility of a tire blowout by failing to check tire wear and inflation because the possibility of this outcome never occurs to him. However, if the driver has experienced (personally or vicariously) a blowout or is reminded about this possibility by a companion then he may "overweight" and/or "overestimate" this outcome, going to great lengths to avoid a low-probability catastrophe (blowout). In a similar vein, an outcome that is objectively certain and has always been experienced may be treated by a decision maker as likely but not certain, and therefore appear to be "underweighted" and/or "underestimated." For instance, a business traveler might choose a room at a hotel chain that had nice rooms and wireless Internet at two of three previously visited properties, rather than a hotel chain with adequate rooms that has provided wireless Internet at three of three previously visited properties. However, if our traveler learns that the second chain explicitly *guarantees* wireless Internet at all locations, knowledge that Internet is certain might cause him to choose this chain instead. Clearly, better prediction of naturally occurring decisions will require a better understanding of the role that both description and experience play and how they interact.

References

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

Barron, G., Leider, S., & Stack, J. (2008). The effect of safe experience on a warnings' impact: Sex, drugs, and rock-n-roll. *Organizational Behavior and Human Decision Processes, 106*, 125–142.

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty, 36*, 153–177.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57,* 94–107.

Fox, C. R. & Hadar, L. (2006). "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making, 1*, 159–161.

Fox, C. R., & See, K. E. (2003). Belief and preference in decision under uncertainty. Chapter in D. Hardman & L. Macchi (eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making.* New York: Wiley.

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44,* 879–895.

Hau, R. C., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sampling size and experienced probability. *Journal of Behavioral Decision Making, 21*, 493–518

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*, 534–539.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006). The role of information sampling in risky choice. In K. Fiedler & P. Juslin (Eds.), Information sampling and adaptive cognition. (pp. 75–91). New York: Cambridge University Press.

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from Prospect Theory in descriptive choice. *Psychological Science, 19,* 1015–1022.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 4,* 263–291.

Kilka, M. & Weber, M. (2001). What determines the shape of the probability weighting function? *Management Science, 47,* 1712–1726.

Lee, W. (1971). *Decision theory and human behavior*. New York: Wiley.

March, J. G. (1996). Learning to be risk-averse. *Psychological Review, 103*, 309–319.

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

Stewart, N., Chater, N. & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53,* 1–26.

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science, 36*, 643–660.

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review, 102*, 269–283.

Tversky, A. & Gilovich, T. (1989) The Cold Facts About the "Hot Hand" in Basketball. *Chance: New Directions for Statistics and Computing, 2*, 1, 16–21.

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

Ungemach, C., Chater, N., & Stewart, N. (2009). Are Probabilities Overweighted or Underweighted When Rare Outcomes Are Experienced (Rarely)? *Psychological Science, 20*, 473–479.

Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review, 111,* 430–445.

Wu, G. & Gonzalez, R. (1999). Nonlinear decision weights in choice under uncertainty. *Management Science, 45,* 74–85.

Wu, S. W., Delgado, M. R., & Maloney, L. T. (2009). Economic decision-making under risk compared with an equivalent motor task. *Proceedings of the National Academy of Sciences USA, 106*, 6088–6093.