The impact of experience on decisions based on pre-choice samples, and the face-or-cue hypothesis

Ido Erev[1], Ofir Yakobi[2], Nathaniel J. S. Ashby[3] and Nick Chater[4]

1. Technion

2. University of Waterloo

3. Harrisburg University of Science and Technology

4. University of Warwick

Nov 8, 2021

Abstract. The growing literature on how people learn to make decisions based on experience focuses on two types of paradigms. In one paradigm, people are faced with a choice, and must retrospectively consult past experience of similar choices in order to decide what to do. In the other paradigm, people are faced with a choice, and then have the opportunity prospectively to gather new experiences that might help them make that choice. The current paper examines the joint impact of both retrospective and prospective experiences. Two experiments reveal strong interactions. In Study 1, repeated experience with new samples appears to reduce sensitivity to the average outcome in the samples and enhances underweighting of rare events. Study 2 shows that repeated experience with pre-choice samples can reverse the impact of the new information (and decrease the tendency to select the alternative that provides the best outcome in the new sample). The results suggest that prospectively gathering new samples can have two, potentially contrasting, influences on choice: The first focuses on the sample's face value and selects the option with the higher value in the new sample; by contrast, the second treats the new sample as a cue to recall similar prior experiences, which in turn drive choice. The paper concludes with a discussion of the possibility that part of the descriptive value of prospect theory reflects the fact that it summarizes the joint impact of similar "face-or-cue" processes.

[*]Correspondence concerning this article should be addressed to: Ido Erev, erevido@gmail.com.

Acknowledgements: Ido Erev was supported by the Israel Science Foundation (grant no. 535/17). Nathaniel J. S. Ashby was partially supported by a Technion fellowship. Nick Chater was supported by ERC grant 295917-RATIONALITY, the ESRC Network for Integrated Behavioural Science [grant number ES/K002201/1], the Leverhulme Trust [grant number RP2012-V-022], and RCUK Grant EP/K039830/1.

In many natural choice tasks, ranging from buying sneakers to important policy decisions, the information available to the decision-makers comes from three main sources: descriptions of the possible alternatives, retrospectively consulting experience of similar decisions in the past, and prospective exploration of new experiences to help inform the decision. For example, while choosing between pairs of sneakers in the store, the decision-makers can rely on verbal descriptions of the pairs, use past experiences with similar pairs, and collect new experiences by trying the available pairs on. Similarly, when world leaders considered alternative reactions to the 2008 subprime crisis, they relied on reports concerning the risk that different financial institutions will collapse, recalled the outcomes of previous crises, and searched for new information, including expert judgments and public opinion, about the current crisis to help decide which policy to choose.

Most experimental studies of choice behavior concentrate on the impact of only one of these three sources of information. Mainstream studies of the impact of description (e.g., Kahneman & Tversky, 1979) used the “description” paradigm presented in the left-hand side of Figure 1. The clearest analyses of the impact of past experiences used the “consequential full feedback clicking” paradigm described in the center of Figure 1 (see review in Erev & Haruvy, 2016), where each new decision about which button to click is informed by retrospective recall of what reward was received when the buttons were clicked in the past. The clearest analyses of the impact of new experiences used the sampling paradigm described in the right-hand side of Figure 1 (see review in Wulff et al., 2018), in which, when faced with a decision, people can prospectively gather evidence about which button to click by costlessly sampling reward values from each button. In this exploration phase, they do not actually receive any reward. The purpose of the sampling is purely to inform the decision to be made.

Figure 1. Typical screens in the Description, Consequential Full Feedback Clicking (Clicking for short), and Sampling paradigms used in previous research.

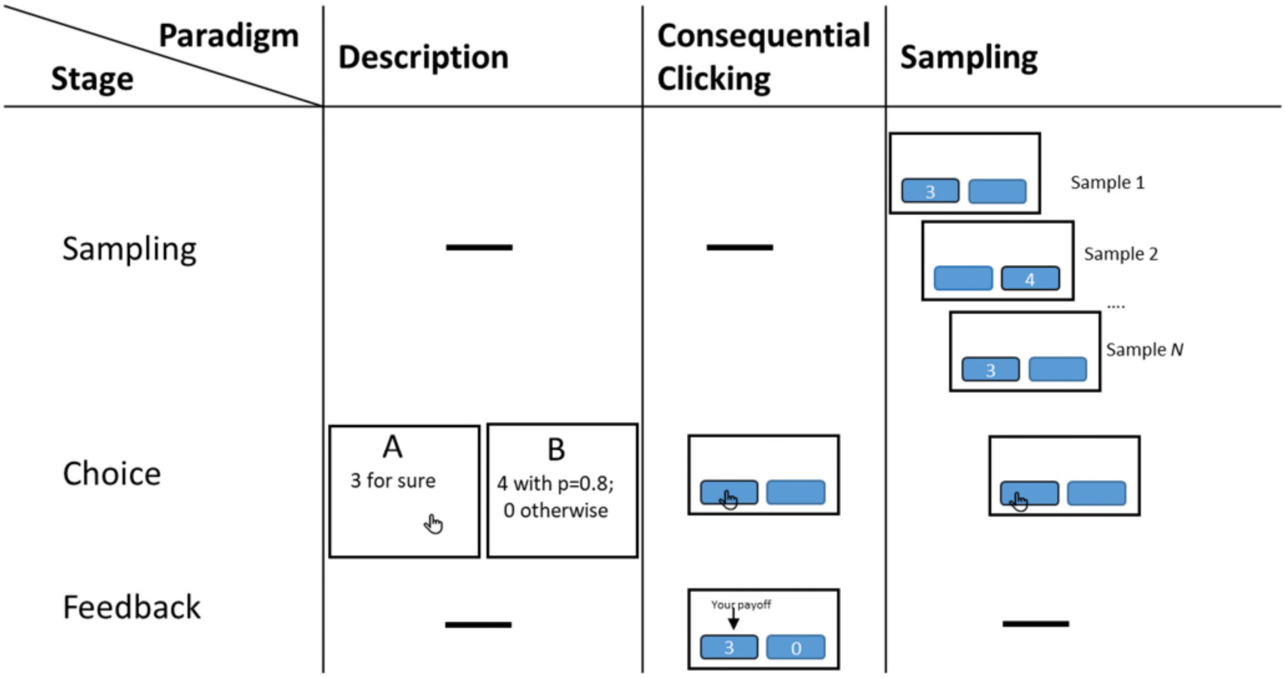

Comparison of the three lines of research has highlighted a "description-experience gap" (Hertwig & Erev, 2009): Studies that used the description paradigm document significantly higher sensitivity to low probability (rare) outcomes than studies that used the clicking and sampling paradigms (Barron & Erev, 2003; Hertwig et al., 2004). In addition, previous research documents significant differences between the two types of "decisions from experience" (de Palma et al., 2014). In order to clarify these differences, it is useful to distinguish between the measurement of the decision-experience gap, and the definition of "underweighting of rare events." In binary choice, the gap is the difference between the choice rate of the option that "pays more with higher probability" when the decision is made from experience and from description. The term underweighting of rare events refers to a tendency to prefer that option that pays more with higher probability when this option does not maximize expected return.[1] The

[1] These definitions imply that the existence of the gap does not imply underweighting of rare events in decisions from experience. For example, consider an experiment that studies decisions between R "10 with probability .9, 0 otherwise" and S "9 for sure" using the sampling and the description paradigms. Assume that the R-rate (the choice rate of the option that pays more with higher probability) is 40% in the sampling paradigm, and 10% in the description paradigm. The gap in this example is large (40% - 10% = 30%), but the results do not exhibit

main difference between the two experience paradigms involves the conditions that trigger underweighting of rare events. Studies that used the consequential full feedback clicking paradigm suggest that when the feasible prospects are of similar EV, the tendency to underweight rare events is rather robust.[2] This bias occurs even when the decision makers receive accurate information concerning the feasible payoffs (e.g., Erev et al., 2017), and when the observed rate of the rare outcomes is larger than the expected rate. In contrast, studies that used the sampling paradigm document clear underweighting of rare events only when three additional conditions hold: The decision makers are not presented with a description of the possible outcomes (Erev et al., 2008; Abdellaoui et al., 2011), the rare events are underrepresented in the sample (Wulff et al., 2018), and only one of the options is risky (Glöckner et al., 2016).[3] Meta-analysis (Wulff et al., 2018) shows that the main properties of decisions from prospectively drawing samples (rather than sampling retrospectively from memory) can be summarized by the assertion that people use the "natural mean" heuristic (Hertwig & Pleskac, 2008) and select the option with the higher sample mean.

To derive practical implications from these experimental results it is important, of course, to understand how these different sources of information interact. It is known, for example, that the impact of the description can be modified by past experiences (Yechiam et al., 2005; Jessup et al., 2008; Lejarraga & Gonzalez, 2011; Marchiori et al., 2015; Erev et al., 2017; Cohen et al., 2020). In one demonstration of this effect (Marchiori et al., 2015) participants faced 60 distinct decisions from description tasks involving a safe prospect (e.g., "5 with certainty") and a long-shot gamble ("100 with p = .05; 0 otherwise"). After each choice, participants received full feedback concerning the realized payoffs of each option. The results reveal an initial tendency to prefer the gamble that implies overweighting of rare gains (as predicted by Prospect theory; Kahneman & Tversky, 1979; Wakker, 2010). Yet, after a few trials most participants preferred the sure payoff, and behaved as if they underweighted the rare event. These results can be

underweighting of rare events in the sampling paradigm (as the choice rate of the option that pays more with higher probability is lower than 50%).

[2] When the feedback is limited to the obtained payoff, the tendency to underweight negative rare outcomes can be masked by the hot stove effect (Denrell & March, 2001).

[3] While Wulff et al. (2018), did not find underweighting of rare events when the proportion of the rare outcome in the samples equaled its probability, they did find a weak description-experience gap even in this case. That is, the implied weighting of the rare outcomes was lower in decision from experience than in decisions for description.

explained with the assumption that the descriptions are used as memory cues, and lead people to select the strategies that provided the best outcomes given similar cues in the past.

The current paper extends this research to consider whether new samples can also cue memory for past experiences, which in turn influence choice. Our analysis rests on the assumption that people tend to select the alternatives that led to the best outcomes in similar situations in the past (Skinner, 1953; Plonsky et al., 2015; Chater et al., 2020). But which similar situations are retrieved from memory may depend, in part, on which new samples are experienced. Thus, it is possible that new samples that favor one of the options, may trigger memories of bad experiences with that option that will in turn decrease the tendency to select it. For instance, a pleasant feeling while trying on a new pair of sneakers can remind the consumer of an old pair that felt similarly comfortable in the store but lost its comfortable feeling after a few weeks.

Our experimental analysis reveals strong interactions between prior experience and new samples. In Study 1, repeated experience with pre-choice samples appears to reduce sensitivity to the average outcome in the samples and enhances underweighting of rare events. Study 2 shows that repeated experience with pre-choice samples can reverse the impact of the new information (and *decrease* the tendency to select the alternative that provides the best outcome in the new sample). The paper concludes with a discussion of the common elements of decisions from experience and decisions from description.

Study 1

The current study compares the joint impact of new and old experiences by studying repeated choice between the two payoff distributions presented in Figure 2 using the experimental paradigm presented in Figure 3. Each participant faced this choice task in 100 trials and could base their choices on new evidence (a pre-choice sample of 12 pairs of outcomes) and on the experience obtained in previous trials.

Figure 2: The choice problems examined in Study 1, The error term at time, e_t, is drawn from a normal distribution with a mean of 0 and standard deviation of 3.

Options	Risky (R)	Safe (S)
	$19 + e_t$ with p= .9; $30 + e_t$ with p =.1	$20 + e_t$

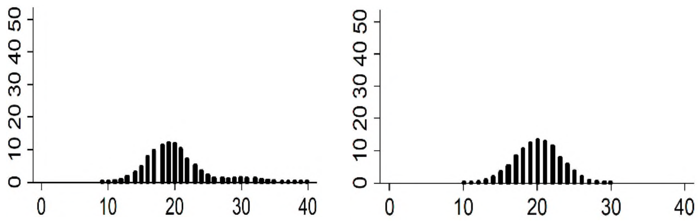

Figure 3: (a) The instructions presented to participants at the beginning of the experiment and (b) the experimental paradigm.

(A)

The current academic experiment includes 100 rounds.
In each round you will be asked to select between two alternatives, and your choice will lead to some earning in points. The feedback after each choice, presented on the buttons, will include the payoff of the chosen option and the payoff from the option that you did not choose.
After pressing continue, a stream of samples from each alternative will be presented, in a rate of 500 milliseconds a sample.
These samples represent simulation of playing each of these alternatives, thus, they will give you some idea on the possible outcomes and their likelihood of occurrence.
When the stream of 12 samples stops, the buttons will become active (blue) and you can make your choice by pressing one of the buttons.
Important: After the buttons (A and B) become active (changed their color to blue), you have only 1.2 seconds to click one of them. If you will not click in time, you will not get points for this round.

(B)

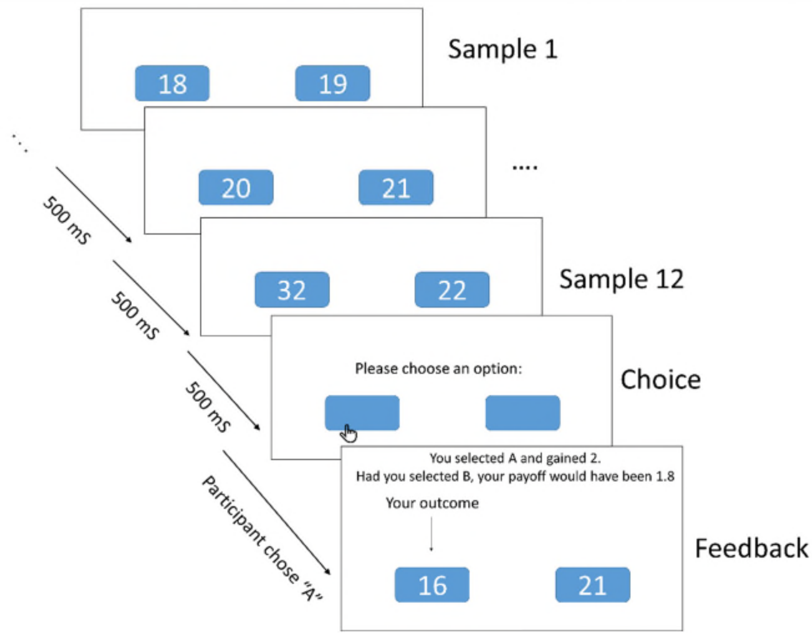

Notice that the paradigm described in Figure 3 is a variant of the consequential full feedback clicking paradigm in Figure 1, but crucially now with the addition of a costless sampling stage. It differs from the sampling paradigm presented in Figure 1 in three ways: The availability of feedback, the fact that the sampling was controlled by the experimental design, and the fact the samples are presented in pairs.

Participants and Design

Forty-five participants were recruited from Amazon Mechanical Turk. All participants had a no-response rate below 10% (i.e., they made a decision within the 1.2s time frame in 90 trials or more). The sample included 30 females (15 males), with a mean age of 35.54 (SD_{AGE} = 10.3). We planned to run 40 participants; the studies that first documented the phenomena we try to replicate and clarify ran fewer subjects. For example, the relevant conditions in Barron and Erev (2003) and Tsetsos et al. (2012) included 24 and 18 participants respectively.

Participants received payment contingent on their decisions during the experiment plus a fixed participation fee. The total payment ranged from \$3.57 to \$4.04 with a mean of \$3.81. An attention check was included (i.e., participants were instructed to write the word “thanks” in a text field labeled “comments” on the instructions page) and 12.7% potential participants who failed it were not allowed to take part in the study. The experiment was programmed in the oTree platform (Chen, Schonger & Wickens, 2016).

Materials and Procedure

Figure 3 presents the experimental task. The main details of the experimental procedure were adapted from Tsetsos et al. (2012). After pressing “Continue”, a sequence of 12 sample pairs were presented on the two buttons. Each pair was drawn from the corresponding distribution and presented for 500ms. The fact that the same e_t term is used in each pair of draws implies positive correlation between the outcomes of the two options. The choice keys became active only after the presentation of the set was finished, and participants had 1.2s to make a choice or the round would be skipped. The choice resulted in the presentation of a new pair of draws, the payoff of the selected option was the trial’s payoff (and the conversion rate from points to money was 1 US cent for 10 points). The payoff for skipped rounds was zero, and participants were informed that if they did not respond in time and were paid zero for this round:

Time pressure was introduced to keep participants engaged. The total sample presentation time was 6 seconds (0.5s x 12).

Results and Discussion

The left-hand side of Figure 4 displays the mean rate of choosing the risky option (R-rate) in 10 blocks of 10 trials by the number of high payoffs from R ($30 + e_t$, 11 more than S) in the sample. Comparison of the current results to previous studies of decisions based on pre-choice samples (reviewed in Wulff et al., 2018) highlights one similarity and one difference. A repeated measures ANOVA with the number of rare events in the sample and block as two within-subject factors, suggests that the similarity involves the existence of high sensitivity to the number of rare events in the sample, $F(3,132) = 32.34$, $p < 0.001$. The R-rates given zero, one, two, and three or more rare events (high payoff from R) are 0.19, 0.31, 0.38 and 0.48 respectively, and they are all significantly different from each other at the $p < 0.001$ level.

Figure 4 – Study 1: Risky-choice rate by experience (blocks of 10 trials) and number of high payoffs from R in the sample.

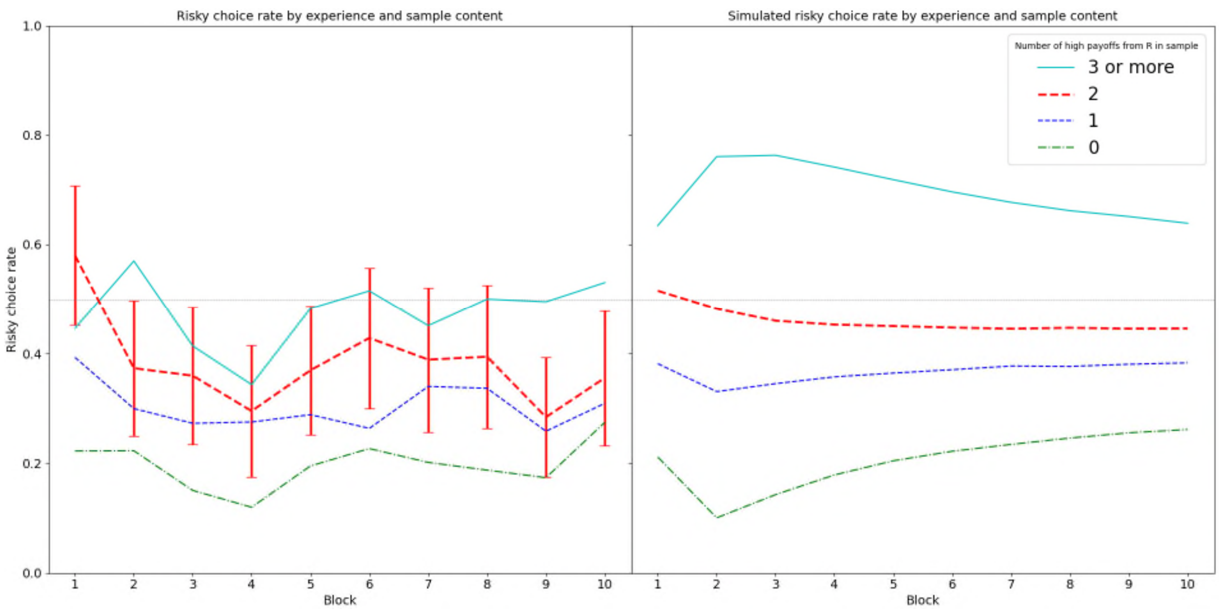

Note: The left-hand side presents the experimental results, and the right-hand side shows the rates under the face-or-cue model presented below.

The difference between the current results and the results reviewed by Wulff et al. (2018) involves the conditions that trigger deviations from maximization that imply underweighting of rare events. The experiments reviewed by Wulff et al. (2018) show a clear underweighting of rare events only when the frequency of the rare events in the sample was lower than the objective probability. Analysis of the R-rate given a sample with two rare events (when observed frequency, 2/12 = 0.167, is larger than the objective probability of 0.10) shows the emergence of the behavior predicted by underweighting of rare events (R-rate below 50%) with experience. The average R-rate over all 100 trials (38%) is significantly lower than 50%, $t\ (44) = 2.782$, $p = 0.008$. Figure 4 shows that starting from the second block (after 10 trials) the median choice reflects underweighting of rarely observed outcomes. The difference between the R-rate given a sample with two rare events in the first block (58%) and over all the other blocks (37%) is significant, $t(40) = 4.10$, $p = 0.0002$).[4]

According to one explanation of the effect of repeated experience, documented in Figure 4, the results reflect the joint impact of two "reactions to samples" rules (see a related idea in Achtziger & Alós-Ferrer, 2014): The first treats the new samples as a set of estimates of the final payoff, and implies a choice of the option with the highest average estimate. That is, it depends purely on the samples' face value. The second rule treats the samples as cues that can be used to estimate the current state of the world, and implies a choice of the option that led to the best outcomes under similar states in the past. This "face-value or samples-as-memory-cues" ("face-or-cue" for short) hypothesis implies that the decrease in the R-rate in response to a sample with two rare events reflects an increase in the tendency to treat the sample as a cue with experience (Plonsky et al., 2015).

A second explanation for the decrease in the R-rate with experience even after observing a sample with 2 high payoffs from R assumes that boredom, or fatigue, decreases the size of the subsamples used by the participants. Under this "decreased-sample-size" hypothesis, the decline from near 60% to around 40% in the R-rate when the full sample includes two attractive rare payoffs from the risky alternative, can be the product of a decrease in the size of the subsample

[4] This analysis was not planned in advanced, it was conducted in response to a reviewer's request. The degree of freedom reflects the fact that only 41 participants experienced and responded in time to a sample with two rare events in the first block, and our analysis used a within person test. Additional analyses lead to similar conclusions. For example, the R-rate in the very first trial with a sample containing two rare events (62.5%) is significantly higher than the rate in all other trials with two rare events (37%), $t(39) = 3.86$, $p = 0.0004$. Similarly, the R-rate in the very first trial with a sample with two or more rare events (55%) is significantly higher than the rate in all other trials with two or more rare events (40%), $t(37) = 2.10$, $p = 0.043$.

from 5 to 3. The predicted R-rates given subsamples of size 5 and 3 are 59% and 42% respectively.[5]

Study 2

Study 2 was designed to compare the two explanations to the results of Study 1. It examines a dynamic environment in which the two hypotheses predict contradicting patterns. Specifically, it studies the incentive structure described in Table 1 in which the payoff distribution associated with the risky option depends on the state of nature. In state L, the risky option provides “20 with certainty.” In state H the payoff distribution of the risky option is “10 in 10% of the trials, 20 in 88%, and 150 in the remaining 2%”; thus, the expected value (EV) is 22. The payoff distribution associated with the safe option is normal with mean of 21 independently of the state. The State was randomly determined in the beginning of each trial, and then the participants were presented with a sample of 12 random draws from the two distributions.

The current payoff rule implies that the existence of extreme outcomes (at least one ‘10’ or ‘150,’) in the sample can be used to estimate the probability of State H, and find the option that maximizes expected return. The probability of State H is 1 when the sample includes extreme outcome, and .177 otherwise (see the note in Table 1). The implied EV maximizing strategy is “Risk if and only if the sample includes at least one ‘10’ or ‘150’.” Crucially, if a participant encounters a ’10’, this will reduce the average value of the sample; but will cue people to recall large positive outcomes (because the state of nature of H, rather than L).

Table 1: The payoff distributions associated with the two options in Study 2.

Safe: N (21, 3), (normal with mean of 21, standard deviation of 3, EV = 21)
Risk: The trial’s state of nature is L with probability 0.5, H otherwise If State = L: 20 with certainty, (EV = 20) If State = H: 20 with p = .88; 10 with p = .1; 150 otherwise (p = 0.02), (EV = 22)

Note: Table 1’s incentive structure implies that the probability of State H given a sample from Risk that includes at least one 10 or 150 (or both 10 and 150) is 1. The probability of State H given a sample of size 12 that does not includes 10 or 150 is $.88^{12}/(.88^{12} + 1) = .177$

[5] The decrease in the risk rate given a sample with two rare outcomes could also be explained by assuming an increase in risk aversion with experience. However, this assertion cannot capture the flat learning curves observed when the sample included zero, one, or more than two rare outcomes.

Thus, the face-or-cue hypothesis can predict an increase in the rate of expected payoff maximizing responses to a sample that includes a 10 outcome with experience. In contrast, the decreased-sample-size hypothesis predicts a decrease in the maximization rate in response to a sample that includes a 150 outcome.

A second goal of Study 2 is to contrast two possible formal instantiations of the face-or-cue hypothesis. One instantiation assumes reliance on all the (objectively) similar past experiences, and the second implies reliance on a small sample of similar past experiences. In order to compare these, we explored two ways in which the payoff distributions determined the trials' payoff. In Condition Draw, the trials' payoff, like each of the samples, was a single draw from the relevant payoff distribution. In Condition EV, the trials' payoff was the EV of the relevant distribution (while the samples were still single draws). That is, the final payoff in each trial in Condition EV is "21 with certainty" from Safe, and "22 if the State is H, 20 otherwise" from Risk.[6]

Assuming that the cue strategy reflects reliance on small sets of past experiences, the participants are likely to learn to approximate the payoff-maximizing rule in Condition EV, but not in Condition Draw. In contrast, if the cue strategy is driven by the average payoffs, the participants are predicted to learn to approximate the expected payoff maximizing rule in both conditions. These predictions were pre-registered (https://osf.io/5fgwr/).

The instructions to the participants, presented in Table 2, did not describe the final payoff rule and did not differ between the two conditions.

Participants and Design

One-hundred and sixty new participants were recruited from Amazon Mechanical Turk. The experiment used a two group between-subject design. The sample included 90 females (69 males, one chose not to disclose), with a mean age of 37.72 (SD_{AGE} = 11.26). Participants received payment contingent on their decisions during the experiment. The payment ranged from $2.72 to $3.3 with a mean of $2.97. As in Study 1, the final sample includes only participants

[6] Notice that Condition EV simulates natural situations, like the sneakers example, in which the sampling (e.g., trying the sneakers in the store) provide less information than the feedback for the actual decision (e.g., buying and wearing the sneakers).

who passed the attention check (175 out of 215 participants) and completed all 100 trials (160 out of 175).

Table 2: The instructions used in Study 2 (same instructions in the two payoff-rule conditions).

Instructions: The current academic experiment includes 100 rounds. In each round you will be asked to select between two alternatives, and your choice will lead to some earning in points. After pressing continue, a stream of samples from each alternative will be presented, in a rate of 0.5 second a sample. These samples are drawn from the payoff distributions associated with the two alternatives. When the stream of 12 samples stops, the two buttons will become active (blue) and you can make your choice by pressing one of them. The feedback after each choice, presented on the buttons, will include the payoff of the chosen option and the payoff from the option that you did not choose. These payoffs are determined by the payoff distribution associated with the selected alternative.

Materials and Procedure

The experiment used a between-subject design. Seventy-eight participants were assigned to Condition Draw, and the other 82 participants faced Condition EV. With the exception of the incentive structure (Table 1), time-limit (there was no 1.2s time-limit in Study 2), and instructions (Table 2), the material and procedure used here were identical to those used in Study 1.

Results and Discussion

The left-hand side of Figure 5 displays the observed R-rate as a function of experience and the composition of the samples. The results favor the face-or-cue hypothesis, based on drawing on a small samples. Table 3 summarizes the results of a repeated measures ANOVA with block and sample as within-subject factors, and condition as a between-subject factor. Post-hoc analyses with Tukey corrected p-values demonstrated that experience in Condition EV increased the R-rate in response to low sample mean (at least one ‘10,’ no ‘150’) from 27% to

63% ($p < .001$); over trials, the low sample in Condition EV increased the R-rate from 29% (given a sample without rare events) to 51% ($p < .001$). Another demonstration of the advantage of the small sample hypothesis is provided by the observation that in Condition Draw, low sample means reduced the R-rate in Condition Draw from 35% to 24% ($p = .004$).

Figure 5 – Risky-choice rate by experience (blocks of 10 trials) and sample content

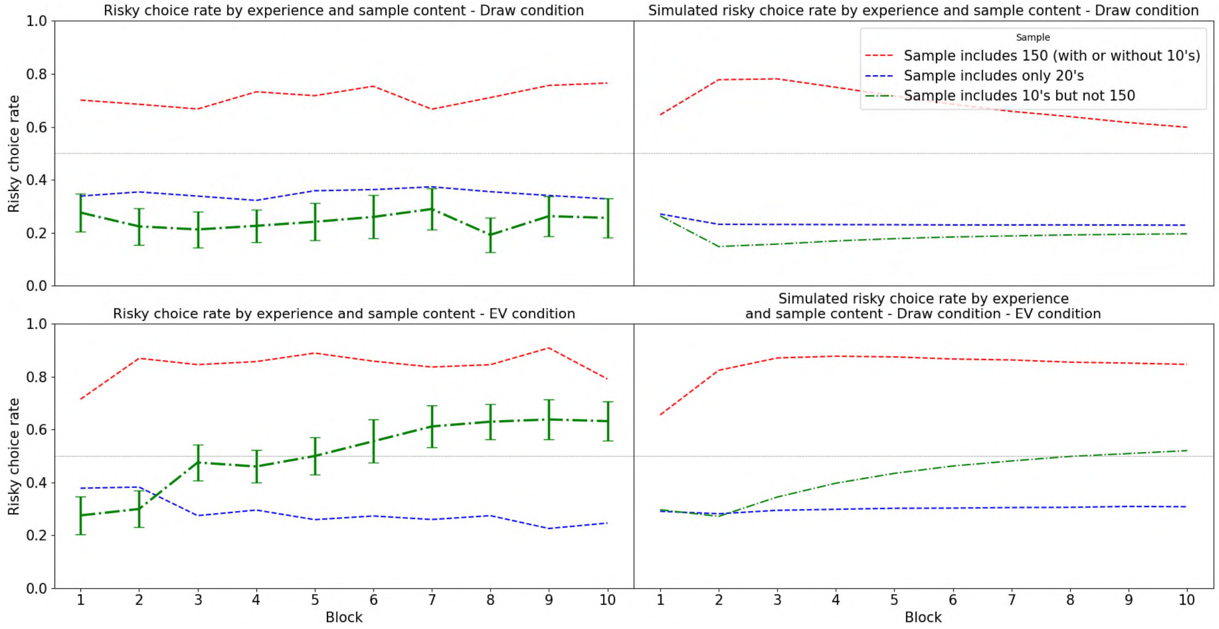

Table 3 – Summary of effects in Study 2

	df	*F*	*P*
Sample	2	4.32	0.014
Sample ∗ Condition	2	35.79	< .001
Residual	234		
Block	9	1.33	0.215
Block ∗ Condition	9	1.21	0.283
Residual	1053		
Sample ∗ Block	18	6.69	< .001
Sample ∗ Block ∗ Condition	18	7.73	< .001
Residual	2106		

A face-or-cue model

To clarify the face-or-cue hypothesis supported above, we outline perhaps the simplest possible formal instantiation of the model compatible with our findings, which assumes that each choice is made based on one of two “reaction to samples” strategies: Face or Cue.

Strategy Face is a generalization of the natural mean heuristic (Hertwig & Pleskac, 2008). It implies a focus on the difference between the payoff from Risk and Safe (V[Risk]-V[Safe], the “the difference score”) in the sample’s 12 pairs, computing the weighted average, and a choice of Risk if and only if the weighted average of difference scores is positive. All the 12 scores are weighted at least once, and some pairs are overweighted. The additional weighting is determined by mental sample k_j (k_j is a property of agent *j*) difference scores, drawn with replacement from the observed sample. For example, if a particular item was drawn twice in the mental sample its final weight is 3/(12+ k_j), and the weight of an item that was not in the mental sample is 1/(12+ k_j). Notice that with the constraint $k_j = 0$, Strategy Face is identical to the natural mean heuristic.

Strategy Cue treats the sample as a cue, and evaluates the option based on the final outcomes observed after obtaining similar cues. It implies drawing a set of k_j past trials that occurred after samples similar to the current sample, and selecting the option with the higher mean payoff in this set. Similarity is defined by the distinction between four similarity classes: *Low* (no rare event in Study 1, at least one “10” and no “150” in Study 2); *Medium* (one rare events in Problem 1, no “10” or “150” in Study 2); *High* (two or more rare events in Study 1, at least one “150” in Study 2), and *Very high* (three or more rare events in Study 1). For example, after a low sample, the agent is assumed to draw k_i final outcomes after a low sample.

The model assumes that repeated experience with each similarity class (cue), increases the tendency to use Strategy Cue. The exact probability is the product of the number of past experiences with Cue *c* (denoted N_c), the presented sample size (N_s = 12 in the current study), and $0 \leq f_j \leq 1$ is a property of agent *j*:

$$P(\text{Strategy Face at the } N_c\text{+1 encounter with Cue } c) = f_j^{\left[N_c/(N_s+N_c)\right]}$$

This equation implies that the probability to use Strategy Cue is 0 when agent *j* first encounters Cue *c*, and it increases toward 1-f_j with encounters. In addition, the model assumes random choice in the first experience with each similarity class.

The distribution of the two traits in the population are assumed to be uniform: $f_j \sim U(0, \varphi)$, and k_j is uniform in the set $\{1, 2, \ldots, \kappa\}$. Thus, the model has two free parameters φ and κ. The right-hand sides of Figures 4 and 5 present the behavior of virtual agents, programmed to behave in accordance with the current model, in the current experiments with the parameters that best fit the data: $\varphi = .4$, and κ = 10. The curves present the mean choice rate over 10,000 virtual agents, and show that the current model captures the main results. In Study 1, it reproduces the slow decrease in the sensitivity to the sample, and the emergence of underweighting of rare events even after observing two rare events in the sample. In Study 2, it reproduces the learning to select the option with the lower sample mean in Condition EV (when the sample from Risk includes the outcome 10), and the flat curves plus underweighting of the extreme rare events in Condition Draw.

It is important to emphasize that the current abstraction is a simple illustrative instantiation of the face-or-cue hypothesis, and more data (more experimental conditions) is needed to allow careful comparison of the feasible alternatives.

General Discussion

Behavioral decision research tries to identify underlying processes by focusing on simple experimental paradigms that concentrate on the way people react to one of the sources of information that drive decisions in natural settings. It uses different paradigms to study reactions to descriptions, past experiences, and new samples. Recent research (e.g., Marchiori et al., 2015) highlights one shortcoming of this convention; it shows that experience can reverse the way people use descriptions of the incentive structure, and suggests that ignoring this interaction can lead to incorrect conclusions concerning the impact of descriptions in natural settings.

The current investigation extends this analysis and examines if previous experience can also impact the way people use new samples. Our results suggest that the impact can be large. Study 1 demonstrates that experience can decrease the weighting of observed rare events. While initial behavior suggests that decision-makers do not underweight the rare outcomes represented in the sample, increased experience triggered underweighting of these outcomes.

Study 2 examines a dynamic environment in which the mean outcome in the pre-choice sample is a noisy predictor of the expected return, and using the composition of sample as a cue facilitates more accurate prediction and EV maximization. The results reveal that experience increased EV maximizing when the maximizing strategy leads to the best possible payoff in most trials (Condition EV). However, when the typical outcomes from following the maximizing strategy were disappointing (Condition Draw where the EV maximizing choice in State H yields

150 in 2% of the trials, and disappointing outcomes in most other trials), the participants failed to learn to use it.

These results can be captured with a model assuming the use of two “reactions to new information” strategies: Face and Cue. Strategy Face builds on the sample’s face value and implies the selection of the option estimated to be best based on the sample. Strategy Cue treats the sample as a cue to recall similar past experiences. Importantly, the best fit is obtained with the assumption that both strategies imply reliance on small samples of outcomes, and experience increases the tendency to treat the sample as a cue.

It is important to emphasize that the current results do not imply a qualitatively different reaction to new samples and past experiences. Indeed, studies of sequential dependencies in pure decisions from experience (Plonsky et al., 2015; Plosnky & Erev, 2017; Plonsky & Teodorescu, 2020) suggest a face-or-cue reaction to the most recent outcomes. The models supported in these studies suggest that in certain cases people select the option that led to the best outcome in recent trials, but in other cases they appear to use the sequence of recent outcomes as a cue and select the option that led to the best outcomes in similar situations.

Wider implications: face-or-cue and prospect theory

In order to illustrate the wider implications of the current analysis, we conclude by considering the possibility that part of the descriptive value of prospect theory (Kahneman & Tversky, 1979; Wakker, 2010) is a result of the fact that the model summarizes the joint impact of face-or-cue processes. We believe that this feasible cognitive interpretation of the descriptive value of prospect theory has two advantages. The first involves the natural abstraction of the large effect of feedback on decisions under risk discussed above. Importantly, this effect cannot be explained by assuming that experience only changes the reference point (as in Kőszegi & Rabin, 2006). The results (e.g., Marchiori et al. 2015; Erev et al., 2017; Cohen et al., 2020) suggest that experience in a specific setting, leads people to behave as if they recall and rely on small samples of past experiences in this setting.

The second advantage involves the abstraction of choice variability in decisions from description. Experimental studies show that when people are presented with the same choice task more than once, they often change their decisions. The magnitude of this tendency is clarified by Wakker, Erev and Weber (1994). Their analysis focused on 24 choice tasks (designed to test the comonotonic independence hypothesis, Wakker, 1990) and 8 “fillers”. Each participant encountered each of these 32 problems twice. The results revealed a high rate of alternations between the two replications. For example, the alternation rate between the first and the second reaction to the (filler) pair: “25% to win $1.5, 0 otherwise” or “20% to win $2, 0 otherwise” was 31%. In addition, the results reveal that the alternation rate increases with temporal distance. For example, analysis of the 24 target problems reveals that the alternation rate was 26% if the distance between the two presentations was less than 5 trials, and 39% if the distance was more than 30 trials. Annoyingly, this rate of inconsistency was larger than the estimated rate of violations of expected utility theory predicted by the comonotonic independence hypothesis.[7]

The common abstractions of choice variability assume that the choice process is noisy; the decision-makers select the prospect with the highest “temporal value,” and that the temporal value is the sum of the “true value” and a noise term drawn from a symmetric normal or logistic distribution. Moreover, the common abstractions also assume that the noise terms are independent and identically distributed (iid). While these abstractions have many attractive properties, they cannot capture the effect of temporal distance. The face-or-cue hypothesis suggests a simple explanation. Assuming that decisions are affected by cues that trigger the recall of past experiences, the probability of recalling additional experiences is likely to increase with time. For example, it is possible that the description in trial 17 of the experiment triggered the recall of experience in which the participant was lucky, and this “feeling” increases risk-seeking until the recall of a very different experience in trial 35.

Summary

The current analysis demonstrates that the common effort to study human decision making using simple experimental paradigms that concentrate on one source of information, can lead to incorrect conclusions. It can lead to exaggeration of the differences between the ways people react to descriptions, old experiences, and new samples. Analysis of the interaction between the impact of the distinct sources of information suggests that differences between the

[7] The wise idea of presenting each problem twice was introduced by Wakker. However, his initial reaction to the results was not positive. He said (smiling, more or less) “these subjects are not very smart, not sure that I want to run more experiments.” The clarification of the magnitude of within-person choice variability has contributed to Erev’s interest in the implications of noise (e.g., Erev, Wallsten & Budescu, 1994), and the impact of experience (e.g., Barron & Erev, 2003).

impact of the three sources of information may not be large. In particular, it is possible that after gaining experience people always try to select the option that led to the best outcome in a small sample of similar situations in the past.

References

Achtziger, A., & Alós-Ferrer, C. (2014). Fast or rational? A response-times study of Bayesian updating. *Management Science*, *60*(4), 923-938.

Abdellaoui, M., L'Haridon, O., & Paraschiv, C. (2011). Experienced vs. described uncertainty: Do we need two prospect theory specifications?. *Management Science*, *57*(10), 1879-1895.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine [Rational man's behavior in face of risk: Critique of the American School's postulates and axioms].*Econometrica, 21*, 503–546.

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.

Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88-97.

Cohen, D., Plonsky, O., & Erev, I. (2020). On the impact of experience on probability weighting in decisions under risk. *Decision*, *7*(2), 153.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*(5), 523-538.

de Palma A, Abdellaoui M, Attanasi G, Ben-Akiva M, Erev I, Fehr-Duda H, Fok D, Fox CR, Hertwig R, Picard N, Wakker PP. Beware of black swans (2014): Taking stock of the description–experience gap in decision under uncertainty. *Marketing Letters*, 25(3), 269-280.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics, 75*, 643–669. http://dx.doi.org/10.2307/1884324

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*(2), 153-177.

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological review*, *124*(4), 369.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and underconfidence: The role of error in judgment processes. *Psychological review*, *101*(3), 519.

Fox, C. R., & Hadar, L. (2006). " Decisions from experience"= sampling error+ prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, *1*(2), 159.

Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General*, *145*(4), 486.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534-539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in cognitive sciences*, *13*(12), 517-523.

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Ed.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 209–235).

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263-292.Lejarraga & Gonzalez, 2011)

Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, *121*(4), 1133-1165.

Marchiori, D., Di Guida, S., & Erev, I. (2015). Noisy retrieval models of over-and undersensitivity to rare events. *Decision*, *2*(2), 82.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press.

Phillips, N. D., Hertwig, R., Kareev, Y., & Avrahami, J. (2014). Rivals in the dark: How competition influences search in decisions under uncertainty. *Cognition*, *133*(1), 104-119.

Plonsky, O., & Erev, I. (2017). Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, *93*, 18-43.

Plonsky, O., & Teodorescu, K. (2020). Perceived patterns in decisions from experience and their influence on choice variability and policy diversification: A response to Ashby, Konstantinidis, & Yechiam, 2017. *Acta psychologica*, *202*, 102953.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*(4), 621-47.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, *109*(24), 96599664.

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton, NJ: Princeton university press.

Wakker, P. P. (1990). Characterizing optimism and pessimism directly through comonotonicity. *Journal of economic theory*, *52*(2), 453-463.

Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511779329

Wakker, P. P., Erev, I., & Weber, E. U. (1994). Comonotonic independence: The critical test between classical and rank-dependent utility theories. *Journal of Risk and Uncertainty*, *9*(3), 195-230.

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). How short-and long-run aspirations impact search and choice in decisions from experience. *Cognition*, *144*, 29-37.

Wulff, D. U., Mergenthaler Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, *144*(2), 140-176.

Yechiam, E., Barron, G., & Erev, I. (2005). The role of personal experience in contributing to different patterns of response to rare terrorist attacks. *Journal of Conflict Resolution*, *49*(3), 430-439.