

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Risky Decisions from Personal and Observed Experience

Permalink

<https://escholarship.org/uc/item/77v8h56g>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Gutoreva, Alina

Ludvig, Elliot Andrew

Publication Date

2024

Peer reviewed

Risky Decisions from Personal and Observed Experience

Alina Gutoreva (a.gutoreva@kbtu.kz)
Department of IT & Engineering, 59 Tole bi
Almaty, 050000 Kazakhstan

Elliot A. Ludvig (e.a.ludvig@warwick.ac.uk)
Department of Psychology, University Street,
Coventry CV4 7AL, UK

Abstract

People often learn about risks from other people. In the current study, we investigated the impact of social learning on risky decisions from experience by incorporating direct observational learning. Participants were placed in pairs – one participant observed the other participant sampling from different options, and then both made decisions based on this personal/observed experience. Participants tended to underweight rare outcomes less when learning from observed experience, particularly with high-value rare outcomes. This difference was not reliably significant, however, suggesting a subtle effect. The study discusses potential contributing factors such as active hypothesis testing, psychological distance, social environment, competitiveness, and goal alignment to explain the results. Overall, the findings contribute to understanding the dynamics of social learning in risky decision-making.

Keywords: risky decision-making; learning from experience; observational learning; social learning.

Introduction

People live in an environment that consists not only of objects but also of intelligent agents, such as other people. The impact of this social environment on decision-making processes is substantial (Izuma & Adolphs, 2013; Mahmoodi, Bahrami, & Mehring, 2018; Misyak, Noguchi, & Chater, 2016). Social learning by observing other agents' actions and outcomes can be an effective learning strategy especially in risky settings because it detaches the learner from potentially costly outcomes, saves personal energy, and decreases the search space (Hills et al., 2015; Markant & Gureckis, 2013). In decision-making research, two general experimental paradigms are used: Decisions from Description and Decisions from Experience, which represent two general ways people learn about risk (Hertwig & Erev, 2009; Wulff, Mergenthaler-Canseco, & Hertwig, 2018). Less is known, however, about how people make choices from observing others making risky decisions. Here, we incorporate such observational learning into two experiments on risky choice, by having people make decisions after observing the odds and outcomes that other people have experienced.

Personal and social learning differ in at least two ways. First, in personal learning, one actively makes choices and receives outcomes, but in social learning, one passively encounters someone else making choices and outcomes—a

distinction which parallels the distinction between active and passive hypothesis testing (Markant & Gureckis, 2013). Second, in personal learning one reaches an understanding through direct interaction with the environment by living through events (e.g., Hertwig, Hogarth, & Lejarraga, 2018); in social learning, on the other hand, one learns by facilitated interaction with another individual or the products created by them (Heyes, 1994; Rendell et al., 2011). These two learning modes, however, are seldom used in isolation. The complementarity of personal and social learning is especially vivid in risky decision-making because rare or extreme events only happen to a small minority of individuals. In a social context, the nature of rare events can be harnessed to help other individuals, who have not experienced rare or extreme events, to adapt to such potential outcomes (Csibra & Gergely, 2009; Kasperson, Kasperson, Pidgeon, & Slovic, 1988; Wang, 2008). This social aspect of risk, and specifically, social learning, is rarely discussed in decision-making research. But, it might be crucial in understanding risk perception in learning and decision-making (Boyd, Richerson, & Henrich, 2011).

The classical example of social learning is learning through observation (Bandura & Walters, 1977; Heyes, 2016a; Myers, 1970). Observation of experience can be direct and real-time (e.g., watching a competition) or indirect (e.g., watching a recording on a video of the competition) and can be represented as information transfer from one agent to the other (Cloninger, 1981; Hill, Boorman, & Fried, 2016; Olsson, Knapska, & Lindström, 2020). When observing others, people can either receive information about actions that others make by inferring what they try to achieve based on actions (i.e., goal-directed learning) or they can copy the actions of others without knowing which goal they are trying to achieve (i.e., imitation).

In terms of decision-making, observational learning can increase the subjective value of selected choices by others, which are then more likely to be selected by observers (Chung, Christopoulos, King-Casas, Ball, & Chiu, 2015; Suzuki, Jensen, Bossaerts, & O'Doherty, 2016). The value of chosen options is dependent on individual risk preferences: the same risky choice made by others can be perceived as a gentle nudge for a risk-seeking observer or as a strong push for a risk-averse observer (Chung et al., 2015). In general, however, people use information about the choices of others

similarly to information from their own risky decisions (Michael et al., 2020; Suzuki et al., 2012).

In risky choices, people often decide differently depending on whether they learn about the odds and outcomes through explicit description or from personal experience (Hertwig & Erev, 2009). In particular, they tend to act as if they overweight rare events in description yet underweight them in experience. This phenomenon is known as the Description-Experience Gap, which we examined here from a social perspective. One possible reason for this gap arises from the fact that descriptions are social creations, in that they necessarily come from another individual. Thus, we hypothesized that introducing a social element into experience might make those decisions more similar to description (i.e., less underweighting). In a pair of experiments, we incorporated social learning into a decision-from-experience design with rare events. In this novel variation of the task, in some rounds, participants observed samples drawn by another individual and then decided based on this observed experience. The main hypothesis of the conducted experiments was that when people learn from others via observation, the underweighting of rare events would be reduced compared to learning from experience.

Both experiments also tested the following competing hypotheses: On one hand, the social component may create a competitive environment, whereby people seek to secure a higher return than others, and thus may be more risk-seeking overall when observing others. Alternatively, people are often more risk-averse for others than themselves (Charness & Jackson, 2009; Olschewski, Dietsch, & Ludvig, 2019). Thus, the social environment might lead people to be more risk-averse when learning from others. Finally, in line with the usual findings in decisions from experience, people should act as if they underweight rare events when learning from personal experience, acting risk averse. All the above hypotheses were preregistered.

Methods

The raw data files, code for experiments, methods, hypotheses, exclusion criteria, and planned analyses were preregistered on the Open Science Framework (OSF) at <https://osf.io/ucemz> for Experiment 1 and <https://osf.io/68ewd> for Experiment 2.

Participants

In both experiments, participants were recruited in groups of 2-8 via the SONA system from a paid participant pool at the University of Warwick. The number of participants was determined using a power analysis with 80% power to find a small-medium effect size using GPower 3.1 (Mayr, Buchner, Erlder, & Faul, 2007). In Exp 1, the minimum effect size of interest of $d = 0.4$ at the 5% significance level was used. In Exp 2, the effect size of key results of Exp 1 was used, which was $d = 0.27$. In Exp 1, 102 participants took part, and, in Exp 2, 145 participants took part. In both experiments, data

collection was stopped following the necessary number of participants fulfilling the exclusion criterion (see below).

Participants were randomly divided into pairs. From the collected sample, six participants in Exp 1 and nine in Exp 2, were excluded from the analysis as they had no partner and performed a substitute computer task. In Exp 1, from the



Figure 1: The computer screen during sampling choice problems. Participants repeatedly chose between two doors and received feedback about the selected door.

participants who were successfully paired, six were excluded because they did not complete the task due to a software error and 10 were excluded based on the pre-registered exclusion criteria, leaving 80 participants ($M_{\text{age}} = 23.3 \pm 2.3$, 56 women). In Exp 2, from the participants who were successfully paired, 34 were excluded – either they did not follow the instructions or they or their partner failed to pick the dominant option on the catch trials; a further four participants experienced a software error, leaving 98 participants ($M_{\text{age}} = 21.7 \pm 0.7$, 46 women). In both experiments, participants were paid £3 as a show-up fee with a chance to make an additional bonus of up to £22. The mean bonus in Exp 1 was £1.86, and the mean bonus in Exp 2 was £2.97. The research was approved by the relevant University of Warwick research ethics committee.

Design

Both experiments had a within-participants element and a matched-design element. First, within participants, we compared the risky choices based on a participant's sampling from the different options against choices made by that same participant based on observing another participant's sampling. Given the stochasticity in the sampling from a random distribution and differences in choices made during the sampling phase, this comparison is likely made about situations with slightly different experienced outcomes and probabilities. As a result, this analysis was supplemented with a matched-design comparison. Here, we compared the choices of the observed with the observer – in this case, both partners' experiences were identical, but the experience was obtained from different conditions.

Materials

Both experiments consisted of choice problems between a safe option providing a Medium payoff (M) with certainty and a risky option providing a High payoff (H) with probability P and a zero payoff otherwise. Table 1 details the five choice problems in Exp 1 and the six choice problems in Exp 2. Each choice problem was between one risky and one safe option. The risky option had one rarely occurring outcome with 5-10% probability (see Table 1), and the safe option always provided the same medium outcome.

In Exp 1, in Problems 1 and 2, the rare outcome was the high outcome, and the common outcome was the zero outcome. In Problems 3 and 4, for the risky option, the rare outcome was the zero outcome and the common outcome was a high outcome. Thus, the high outcomes could be either a rare or a common event, ensuring no confounding of risk-seeking and underweighting of rare events. The expected values were the same for risky and safe options in Problems 1-4. In Choice Problem 5, the safe option always yielded an outcome which was higher than either possible outcome on the risky option; this was a trick problem, serving as a manipulation check with a dominant safe option to ensure that participants were properly incentivised and paid attention to the task.

In Exp 2, in all the choice problems, the rare outcome was always the high outcome. Two choice problems from Exp 1 were used (Problems 1 and 2), plus three novel choice problems and one novel catch choice problem (see Table 1). The expected values were the same for risky and safe options in Problems 1-5. Problem 6 served as a catch problem. In Exp 1, for each option, outcomes were randomly drawn without replacement from a list of 40 possible outcomes, with 2 (5%) or 4 (10%) rare events (i.e., shuffled outcomes). In Exp 2, the outcomes occurred in a random order, based on the predefined probability.

Table 1: Choice problems used in the experiments.

| Problem | Risky (H) | P(H) | Safe |
|--------------|-----------|------|------|
| Exp 1 | | | |
| 1 | 22 | 0.1 | 2.2 |
| 2 | 20 | 0.05 | 1 |
| 3 | 5 | 0.9 | 4.5 |
| 4 | 2 | 0.95 | 1.9 |
| 5 (trick) | 3 | 0.9 | 5 |
| Exp 2 | | | |
| 1 | 22 | 0.1 | 2.2 |
| 2 | 20 | 0.05 | 1 |
| 3 | 16.7 | 0.06 | 1 |
| 4 | 12.5 | 0.08 | 1 |
| 5 | 20 | 0.1 | 2 |
| 6 (trick) | 2 | 0.1 | 4 |

Note: H = higher outcome for risky option; Low outcome in risky choice was always zero (not shown).

Procedure

The experiments were performed on Windows 10 computers using PsychoPy software version 1.90.3 (Peirce, 2009). Participants were first presented with a summary of the study and provided informed consent. Participants initially sat back-to-back facing their computers. The experiment began with instructions on the computer screen of each participant's computer. As can be seen in Figure 2, participants started each round at their computers and were then informed whether, in the upcoming round, they were to sample for themselves as the *experiencer*, or their partner would sample

and they would observe, as the *observer*. When a participant sampled for themselves, they stayed at their computer and waited for their partner to sit next to them to observe. When participants observed their partner, they moved to sit next to their partner's computer. As illustrated in Figure 1, on each trial, participants were presented with pictures of two doors, and they indicated which door they wished to sample by using the keyboard arrow key (left or right arrow) for the corresponding door. The procedure was analogous to the Ludvig & Spetch (2011) Doors Task. Selections were immediately followed by feedback for 1.2 s, which showed the points corresponding to that door only, as in the sampling paradigm (Hertwig & Erev, 2009). One participant (the *experiencer*) sampled 40 times from the two options using the arrow keys, while the *observer* observed the selections; both participants observed the feedback after each selection as shown on the right of Figure 1. After the round finished, the *observer* moved back to their computer and indicated their preferred option based on the 40 viewed samples, as a single final decision (Figure 2). Participants did not observe each other's final decisions. During the task, participants were asked not to talk nor directly intervene in any way in their partner's selections.

Each participant experienced all choice problems twice, once by sampling from the options themselves while their partner was observing them (i.e., standard Decisions from Experience, DfE), and once by observing their partner's sampling from the same choice problems (i.e., Decisions from Observation, DfO). The order of rounds was randomly shuffled for each pair. Each round had a different pair of door images representing the two options. This shift was necessary to limit the possibility that participants recognized that the choice problems were repeated in the observed and




Figure 2: Schematic of the experimental set up. Participants played the choice problems in pairs. In the sampling phase, one of the participants in a pair sampled the doors using the arrow keys, while the other participant observed them. For the final choice, the participant who observed went back to their own computer, so that each participant made their final choice without observing what the other chose.

experienced conditions and only differed in terms of the sampling differences between the participants. Nevertheless, in the qualitative feedback after the experiments, about a third of the participants reported that they recognized that they played the same choice problems twice. The pay-out for participants depended on the final choices: out of the ten choices they made, one of their chosen options was randomly selected and then played out using the generative odds as given in Table 1. This outcome served as a bonus in addition

to the show-up fee, making all final choices incentive-compatible. The experiments lasted less than 30 minutes.

Data Analysis

The dependent measures were the proportion of risky choices and the underweighting of rare events in Exp 1 and the proportion of risky choices only in Exp 2. The proportion of risky choices was defined as the ratio of the number of times the risky option was chosen in the final decision and calculated across all (non-trick) choice problems in each condition (see Table 1). In Exp 1, the degree of underweighting was defined as the sum of safe choices made in Problems 1-2 (where the rare event was a large win) and risky choices made in Problems 3-4 (where the rare event was a zero outcome) divided by four, the overall number of choice problems. The proportion of risky choices and degree of underweighting were calculated for each participant.

All statistical comparisons were made twice: once within participants and once with matched participants. The within-participant comparison pitted the final choices in Experience against those in Observation or the same participant. The matched-participant comparison pitted the final choices for the observed in their Experience and the observer in their Observation.

The data from all non-excluded participants was tested for normality using a Kolmogorov-Smirnov test and tested for equality of sample variances with Levene's test. In Experiment 1, a one-sample Kolmogorov-Smirnov test showed that both variables—proportion of risky choices and underweighting of rare events—were not normally distributed ($D(79) = 0.5, p < .01$ for both tests). Levene's test for equality of variances showed that the variability in the two conditions for both proportion of risky choices ($F = 1.65, p = .20$) and underweighting of rare events ($F = 0.18, p = .67$) were not reliably different. Because the data did not meet the requirements for a parametric test (i.e., the data was not normally distributed), Wilcoxon tests for within-subject comparison were performed. As a robustness check, standard t-tests were also performed and yielded qualitatively similar results. Bayes Factors were also calculated using the *BayesFactor* library in R and are reported below.

In Experiment 2, Kolmogorov-Smirnov and Levene's tests were performed to assess differences in the distribution and variance of samples in each population. Both the proportion of risky choices and the underweighting of rare events was not normally distributed ($D(98) = 0.5, p < .001$ and $D(98) = 0.5, p < .001$ respectively). The distribution between the two conditions was not different ($D(98) = 0.112, p = .568$), and the variances between the conditions were also not reliably different ($F(1,194) = 3.01, p = .085$). Thus, for Experiment 2 data, paired-sample and matched-sample t-tests for within- and between-subject comparisons were performed accordingly.

All data analysis was conducted in RStudio (Version 1.2.5033). Effect sizes were calculated as Cohen's d , and mean differences are presented with 95% confidence

intervals. The data analyses followed the pre-registered plan, except for the exploratory analysis which is acknowledged.

Results

The main hypothesis in the experiment was that, given the proposed mapping between personal and social learning, DfO would differ from DfE (i.e., less underweighting of rare events). As shown in Figure 3A, participants tended to underweight rare events in DfO ($54.0 \pm 3.1\%$) 5.7% less than in DfE ($59.7 \pm 3.0\%$), although this trend was non-significant in a one-tailed test, $W(n=80) = 926, d = 0.23, p = .076; BF = 0.439$. In Observation, participants did not reliably act as if they underweighted rare events, but rather acted as if they had no preference for the safe or risky option ($z = 1.34, d = 0.14, p = .18, BF = 0.228$). As predicted, however, when making choices from experience, participants did underweight rare events, $z = 3.06, d = 0.36, p = .002, BF = 9.870$. The direction and the mild effect are in line with the hypothesis that learning from observation is different from experience, but the evidence was limited.

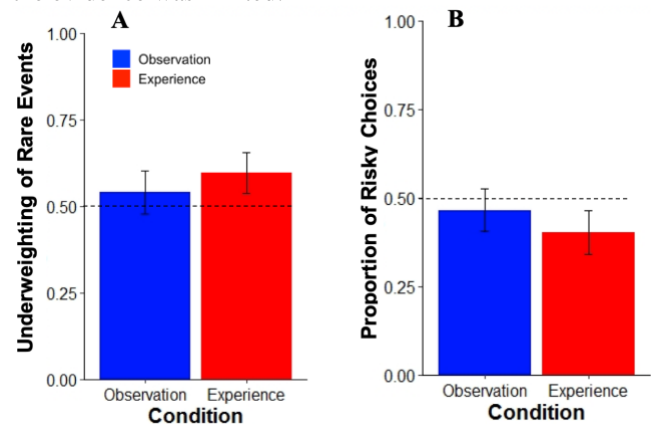


Figure 3: Experiment 1. Mean proportion of (A) underweighting of rare events and (B) risky choices across all choice problems for Observation and Experience conditions. The dashed line represents chance level. Error bars represent 95% CI.

We also hypothesised that participants may be more risk-seeking when observing others, or they may be more risk-averse when learning from others as compared to learning from personal experience. Figure 3B shows that people were non-significantly more risk-seeking in observation as compared to experience, $46.6 \pm 3.0\%$ vs. $40.3 \pm 3.1\%$ respectively; $z = 1.77, d = .23, p = .076, BF = 0.439$. Deciding from observation led participants to be slightly more risk-seeking compared to deciding from personal experience, thereby decreasing the usual bias observed in the decisions from experience literature. However, in experience, participants also acted as would be expected from the literature, by underweighting rare events. As hypothesized, when comparing the proportion of risky choices in Experience against chance level, participants were risk-

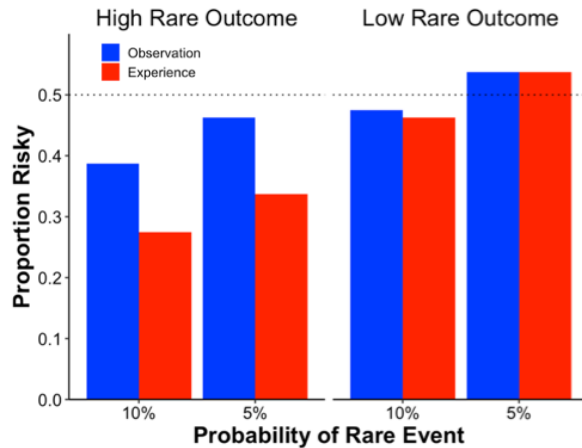


Figure 4: Exp 1. Mean proportion of risky choices in Observation and Experience for each choice problem. The dashed line represents the chance level.

averse, choosing the risky option significantly less often than chance level, $z = 3.07$, $d = 0.35$, $p = .001$, $BF = 1.15$. Overall, people were acting as if they were underweighting rare events in Experience, but there was no reliable difference between the Observation and Experience conditions. The matched-sample analysis, controlling for the variability in the encountered sampling, showed a similar picture as in the paired comparison but with slightly statistically more robust results. In Experiment 1, people underweighted rare outcomes less in Observation as compared to the Experience condition ($W(n=80) = 558.5$, $d = 0.18$, $p = .051$, $BF = 0.228$; Fig.3A) and also chose the risky option with similar frequency in the observed as compared to the experienced condition ($W(n=80) = 847$, $d = 0.17$, $p = .076$, $BF = 0.174$, Fig. 3B).

As part of an exploratory analysis, the data from Exp 1 was split based on the choice problem type: high-value (£22 and £20) or low-value (£0) outcome choice problems (see Table 1). As shown in Figure 4, participants chose the risky option significantly less often in choice problems with high-

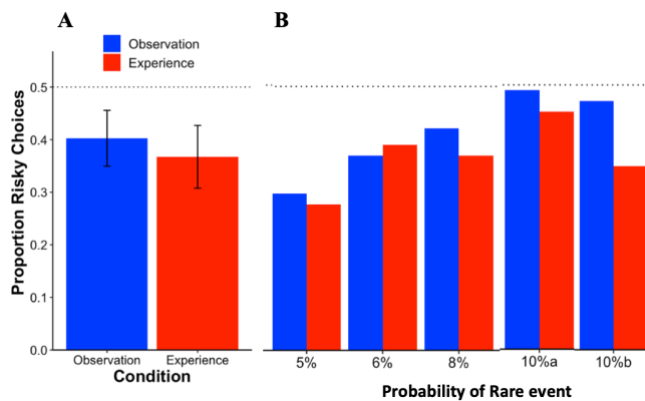


Figure 5: Experiment 2. (A) Mean ($\pm 95\%$ CI) proportion risky choice in Observation and Experience. (B) Proportions of participants who picked the risky option for each choice problem comparing Observation and Experience conditions. Dashed line represents chance.

value rare outcomes than in those with low-value rare outcomes, $W(n=80) = 621$, $d = 0.3$, $p = .007$, $BF = 1.834$. More notably, however, with the high-value rare outcomes, people were slightly more risk-seeking with observation, thereby acting as if they underweighted the rare events less, as hypothesized. This result motivated Exp 2, which only used high-value rare outcomes.

The main hypothesis of Exp 2, as in Exp 1, was that Observation would exhibit a pattern more similar to that generally seen with Description. Specifically, based on the previous results, we predicted that participants would overweight the high-value rare outcomes more and thus select more riskily in Observation than in Experience. As Figure 5 shows, as expected, people were significantly risk-averse in DfE, acting as though they underweighted rare events, $t(97) = -4.51$, $d = 0.46$, $p < .001$, $BF = 903.7$. People were also risk-averse in DfO, choosing the risky option significantly less often than chance level, $t(97) = -3.36$, $d = 0.34$, $p = .001$; $BF = 20.1$. People again chose the risky option slightly more often in DfO ($40.3 \pm 5.3\%$) than in DfE ($36.7 \pm 5.9\%$) condition, but the difference was not significant in a one-tailed paired t-test, $t(97) = 1.43$, $d = 0.14$, $p = .078$; $BF = 0.301$. This result was supplemented by a matched-sample analysis, which controls for the sampling the participants encountered, and the picture was the same for the paired comparison, with no difference between the conditions, $t(97) = 1.25$, $d = 0.12$, $p = .106$, $BF = 0.206$.

Discussion

This study examined risky decision-making in individuals learning from personal experience compared to social experience. The results of Exp 1 confirmed the participants acted as if they underweighted rare events, acting risk-averse whenever the rare events were big wins by the decision-making literature, even whilst being observed by another person (Ruggeri et al., 2020; Wulff & Hertwig, 2018) The hypothesis that participants may be more risk-seeking simply because they observe others did not hold as well as the hypothesis stating that people might be more risk-averse when learning from others. In observation, this underweighting bias decreased, but participants still underweighted rare events as they usually do in DfE. Thus, learning from social experience, specifically from observing a partner making risky choices, though an abstraction from personal experience, is still similar to personal experience.

In both experiments, participants decided from social experience similarly to personal experience, with only a marginal difference. The difference manifested in a smaller underweighting bias in decisions from experience. In Exp 1, two choice problem types were used: (i) rare high value (£22 and £20) or (ii) rare low value (£0). In Observation, participants were more risk-seeking compared to learning from personal experience, with a decreased underweighting bias. However, the results were mild and appeared in the rare high-value choice problems only. In the low-value rare outcome choice problems, however, people chose almost

identically when learning about choice problems by personal experience or observing others' experiences. In Exp 2, learning from observation and personal experience was further explored using high-value rare-outcome choice problems only (aka "rare treasures"). With only rare high-value choice problems, the effect was expected to be higher, but the difference between learning from personal and social experience was again not statistically different, even though the results of Exp 2 confirmed the direction of Exp 1.

One possible explanation for the slight difference in conditions is the greater emotional distance in observation. In both experiments, the sampling phase was not directly incentivized, but participants could still experience the ups and downs of the sampling experience — more so if they were experiencing it directly. Even though there was a similar trend in Exp 2 as compared to Exp 1, when only looking at high-value rare-outcome choice problems, the effect was marginal. Potentially, the low-value rare outcomes impacted both the Observer and the Experiencer, but a rare high outcome only impacted the experiencer due to the extremity of the outcome and the potential emotional charge (Konstantinidis et al., 2017a; Madan et al., 2014; Zaki, et al, 2016). Furthermore, this effect is likely to be relative, because in Exp 2, only the very highest outcomes produced a greater difference between the decisions based on different learning modes. This suggests that given similar goals, risk preferences based on observational learning likely depend on more basic elements of the choice, such as the proportion of risky samples and the magnitude and frequency of rare events. Indeed, attention to learning from experience is often directed towards the outcome rather than the probability of these events (Fiedler & Glöckner, 2012).

Observation can be viewed as a passive hypothesis testing that lacks personal agency to select options (Denrell & Le Mens, 2007; Markant & Gureckis, 2013). The Experiencer's choice was responsible for what both the Experiencer and Observer encountered in sampling, and thus, influenced both partners' final choice. In hypothesis testing by reception, a learner is in a passive mode of inference where they need to make sense of the information that is received to which they have no or only partial control (Bruner, 1961). In a risky-choice setting, passive learning can be particularly advantageous — there is no need to risk personal resources by engaging in actions and outcomes. Passive learning helps to limit affective sampling, hot-stove effects, and recency bias, increasing performance in some tasks (Chi, 2009; Gureckis & Markant, 2012; Markant & Gureckis, 2013).

The similarity of risky choices in the personal and social experience conditions can be explained through partner-aligned goals. Goal alignment happens when making decisions based on other's choices, which maximise personal utility or when there is small or no difference between personal and others' risk preferences manifested in choice (Baker, Saxe, & Tenenbaum, 2009; Michael et al., 2020; Shamay-Tsoory, 2019). Given that the partner's goals were identical — to receive as many points as possible in the final

choice — both partners' risk preference was learnt in sampling, with the only difference in who was actively choosing (Baker et al., 2009). The observer's subjective values were thus dependent on the Experiencer's selected options, which impacted individual risk preferences (Chung et al., 2015; Suzuki et al., 2016). This preference could also play a role in identifying bias in partners' sampling. This suggests that even in the current experiments where there was little difference between the Experiencer and the Observer, passive sampling might create a necessary abstraction from choice to adjust bias. An alternative explanation is that Observers might have had lower engagement with the task during observation because they were not directly involved in sampling. To account for this, a future experiment could potentially make the sampling phase consequential: Partner A would do the sampling, but Partner B would receive the consequences.

Limitations

One limitation of the current experiments is that they did not contain losses. In Exp 1, the value of the rare outcome (high or low) influenced whether there was a difference between the decision from experience or observation. One possibility would be to introduce problems with rare, large losses ("rare disasters"). Similar to Exp 1, underweighting large losses leads to risk aversion, which is opposite to what happens with underweighting big wins. This would also allow a clear distinction between the effects on risk preference and the effects on rare-event weighting which was not addressed in Exp 2. One additional limitation of the current study was that in these experiments, the repetition of the choice problems across observation and experience was only camouflaged by shuffling the problems across trials between the participants in a pair, even while the outcomes were similar across different problems (see Table 1).

Conclusion

The experiments provide a new angle to the study of risky decisions from experience by introducing a social element, through decisions from observation (DfO). The two experiments suggest that learning from social experience by direct observation is similar to learning from experience in this novel experimental setting. Even though the results were not significant, people showed a consistent trend in the hypothesized direction, whereby social learning decreased the underweighting bias found in learning from personal experience (DfE). This difference, however, was mild with a particular effect of "rare treasures" as compared to rare low-value outcomes. The decrease of underweighting bias in observation can potentially be explained by passive hypothesis testing given an aligned goal with experience when learning about risk. In conclusion, despite these intriguing trends, directly observing experience was not sufficient to reliably reduce the underweighting of rare event bias in risky decision-making.

References

- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349.
<https://doi.org/10.1016/j.cognition.2009.07.005>
- Bandura, A., & Walters, R. H. (1977). Social learning theory. In *Englewood Cliffs, NJ: Prentice-hall*. (1st ed.).
<https://doi.org/10.1111/j.1460-2466.1978.tb01621.x>
- Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(2), 10918–10925.
<https://doi.org/10.1073/pnas.1100290108>
- Charness, G. B., & Jackson, M. O. (2009). The Role of Responsibility in Strategic Risk-Taking. *Journal of Economic Behavior & Organization*, *69*(3), 241–247. Retrieved from <http://www.econ.ucsb.edu/~charness/>,
- Chi, M. T. H. (2009). Active-Constructive-Interactive: A Conceptual Framework for Differentiating Learning Activities. *Topics in Cognitive Science*, *1*(1), 73–105.
<https://doi.org/10.1111/j.1756-8765.2008.01005.x>
- Chung, D., Christopoulos, G. I., King-Casas, B., Ball, S. B., & Chiu, P. H. (2015). Social signals of safety and risk confer utility and have asymmetric effects on observers' choices. *Nature Neuroscience*, *18*(6), 912–916. <https://doi.org/10.1038/nn.4022>
- Cloninger, C. R. (1981). The Dynamics of Social Learning. *Science*, *213*(4510), 858–859.
<https://doi.org/10.1126/science.213.4510.858>
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13*(4), 148–153.
<https://doi.org/10.1016/j.tics.2009.01.005>
- Denrell, J., & Le Mens, G. (2007). Interdependent sampling and social influence. *Psychological Review*, *114*(2), 398–422. <https://doi.org/10.1037/0033-295X.114.2.398>
- Fiedler, S., & Glöckner, A. (2012). The dynamics of decision making in risky choice: An eye-tracking analysis. *Frontiers in Psychology*, *3*(OCT), 25643.
<https://doi.org/10.3389/fpsyg.2012.00335>
- Gureckis, T. M., & Markant, D. B. (2012). Self-Directed Learning: A Cognitive and Computational Perspective. *Perspectives on Psychological Science*, *7*(5), 464–481.
<https://doi.org/10.1177/1745691612454304>
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.
<https://doi.org/10.1016/j.tics.2009.09.004>
- Hertwig, R., Hogarth, R. M., & Lejarraga, T. (2018). Experience and Description: Exploring Two Paths to Knowledge. *Current Directions in Psychological Science*, *27*(2), 123–128.
<https://doi.org/10.1177/0963721417740645>
- Heyes, C. M. (1994). Social Learning in Animals: Categories and Mechanisms. *Biological Reviews*, *69*(2), 207–231.
<https://doi.org/10.1111/j.1469-185X.1994.tb01506.x>
- Heyes, Cecilia M. (2016). Born Pupils? Natural Pedagogy and Cultural Pedagogy. *Perspectives on Psychological Science*, *11*(2), 280–295.
- Hill, M. R., Boorman, E. D., & Fried, I. (2016). Observational learning computations in neurons of the human anterior cingulate cortex. *Nature Communications*, *7*(1), 12722.
<https://doi.org/10.1038/ncomms12722>
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Bateson, M., ... Wolfe, J. W. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, *19*(1), 46–54.
<https://doi.org/10.1016/j.tics.2014.10.004>
- Izuma, K., & Adolphs, R. (2013). Social manipulation of preference in the human brain. *Neuron*, *78*(3), 563–573.
- Kasperson, J. X., Kasperson, R. E., Pidgeon, N., & Slovic, P. (1988). The Social Amplification of Risk A Conceptual Framework. In *The Social Amplification of Risk* (pp. 13–46).
<https://doi.org/http://dx.doi.org/10.1017/CBO9780511550461.002>
- Konstantinidis, E., Taylor, R., & Newell, D. (2017). Absolute not relative extremeness of outcomes sway risky decisions from experience Manuscript. *Psychonomic Bulletin & Review*, *25*(5), 1925–1933.
- Ludvig, E. A., Madan, C. R., McMillan, N., Xu, Y., & Spetch, M. L. (2018). Living near the edge: How extreme outcomes and their neighbours drive risky

- choice. *Journal of Experimental Psychology: General*, 47(12), 1905–1918. <https://doi.org/https://doi.org/10.1037/xge0000414>
- Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS ONE*, 6(6), e20262. <https://doi.org/10.1371/journal.pone.0020262>
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, 21(3), 629–636. <https://doi.org/10.3758/s13423-013-0542-9>
- Mahmoodi, A., Bahrami, B., & Mehring, C. (2018). Reciprocity of social influence. *Nature Communications*, 9(1), 1–9. <https://doi.org/10.1038/s41467-018-04925-y>
- Markant, D. B., & Gureckis, T. M. (2013). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1), 94–122. <https://doi.org/10.1037/a0032108>
- Mayr, S., Buchner, A., Erdfelder, E., & Faul, F. (2007). A short tutorial of GPower. *Tutorials in Quantitative Methods for Psychology*, 3(2), 51–59. <https://doi.org/10.1037/0096-1523.32.4.932>
- Michael, J., Gutoreva, A., Lee, M. H., Tan, P. N., Bruce, E. M., Székely, M., ... Ludvig, E. A. (2020). Decision-makers use social information to update their preferences but choose for others as they do for themselves. *Journal of Behavioral Decision Making*, 33(3), 270–286. <https://doi.org/10.1002/bdm.2163>
- Misyak, J., Noguchi, T., & Chater, N. (2016). Instantaneous Conventions: The Emergence of Flexible Communicative Signals. *Psychological Science*, 27(12), 1550–1561. <https://doi.org/10.1177/0956797616661199>
- Myers, W. A. (1970). Observational learning in monkeys. *Journal of the Experimental Analysis of Behavior*, 14(2), 225–235. <https://doi.org/10.1901/jeab.1970.14-225>
- Olschewski, S., Dietsch, M., & Ludvig, E. A. (2019). Competitive motives explain risk aversion for others in decisions from experience. *Judgment and Decision Making*, 14(1), 58–71.
- Olsson, A., Knapska, E., & Lindström, B. (2020). The neural and computational systems of social learning. *Nature Reviews Neuroscience*, 21(4), 197–212. <https://doi.org/10.1038/s41583-020-0276-4>
- Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2(10). <https://doi.org/10.3389/neuro.11.010.2008>
- Rendell, L., Fogarty, L., Hoppitt, W. J. E., Morgan, T. J. H., Webster, M. M., & Laland, K. N. (2011). Cognitive culture: Theoretical and empirical insights into social learning strategies. *Trends in Cognitive Sciences*, 15(2), 68–76. <https://doi.org/10.1016/j.tics.2010.12.002>
- Ruggeri, K., Alí, S., Berge, M. L., Bertoldo, G., Bjørndal, L. D., Cortijos-Bernabeu, A., ... Folke, T. (2020). Replicating patterns of prospect theory for decision under risk. *Nature Human Behaviour*, 4(6), 622–633. <https://doi.org/10.1038/s41562-020-0886-x>
- Shamay-Tsoory, S. G. (2019). Herding Brains: A Core Neural Mechanism for Social Alignment. *Trends in Cognitive Sciences*, 23(3), 1–25. <https://doi.org/10.1016/j.tics.2019.01.002>
- Suzuki, S., Harasawa, N., Ueno, K., Gardner, J. L., Ichinohe, N., Haruno, M., ... Nakahara, H. (2012). Learning to Simulate Others' Decisions. *Neuron*, 74(6), 1125–1137. <https://doi.org/10.1016/j.neuron.2012.04.030>
- Suzuki, S., Jensen, E. L. S., Bossaerts, P., & O'Doherty, J. P. (2016). Behavioral contagion during learning about another agent's risk-preferences acts on the neural representation of decision-risk. *Proceedings of the National Academy of Sciences*, 113(14), 3755–3760. <https://doi.org/10.1073/pnas.1600092113>
- Wang, X. T. (2008). Risk communication and risky choice in context: Ambiguity and ambivalence hypothesis. *Annals of the New York Academy of Sciences*, 1128(1), 78–89. <https://doi.org/10.1196/annals.1399.009>
- Wulff, D. U., & Hertwig, R. (2018). A Meta-Analytic Review of Two Modes of Learning and the Description–Experience Gap. *Psychological Bulletin*, 144(2).
- 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.
<https://doi.org/10.1037/bul0000115>

Zaki, J., Kallman, S., Wimmer, G. E., Ochsner, K., & Shohamy, D. (2016). Social Cognition as Reinforcement Learning: Feedback Modulates Emotion Inference. *Journal of Cognitive Neuroscience*, 28(9), 1–13. <https://doi.org/10.1162/jocn>