Structural Knowledge in Experiential Risky Choice

Kuang-Heng Li (kuangheng.li@warwick.ac.uk)[1], Emmanouil Konstantinidis[1] & Wenjia Joyce Zhao[1]
[1]Department of Psychology, University of Warwick

Abstract

Learning from feedback is a fundamental mechanism through which individuals make decisions under uncertainty. An important open question is how people integrate prior information with repeated feedback to improve their decisions. To address this, we conducted an experiment in which we manipulated whether participants had structural knowledge about the decision environment: half of the participants were told the number of distinct outcomes each option could produce, whereas the other half received no such information and had to infer the outcome structure from feedback alone. Structural knowledge changed behavior under partial feedback, bringing the proportion of risky choice closer to 0.5. We propose a Bayesian model that accounts for participants' initial beliefs about outcome probabilities before receiving any feedback. By capturing choice and confidence patterns across different decision environments, the model demonstrates how structural knowledge shapes learning by altering prior beliefs, thereby linking prior information to experience-based risky choice.

Keywords: risky choice; confidence; decisions-from-experience; structural learning; computational modeling;

Introduction

Risk and uncertainty lie at the core of many real-life decisions. For example, when choosing between different services or products, people often do not know the outcome distribution in advance and must learn it from repeated experience and feedback. To study decision making in such settings, researchers often ask participants to choose between simple monetary gambles (often between a safer option, e.g., £3 with certainty, and a riskier option, £4 with 80% or £0 otherwise) and to learn the outcome distributions of these gambles by making repeated choices and receiving outcome feedback (Barron & Erev, 2003; Hertwig & Erev, 2009). In the standard "decisions-from-experience" (DfE) paradigm, participants receive no prior information about the options and must infer possible outcomes and their associated probabilities from experience alone. In everyday life, however, people often enter decisions with some prior knowledge about the options (e.g., drawn from previous encounters with similar decision problems or from direct descriptive information) which then complements or interacts with the feedback they subsequently receive. An outstanding question and the main focus of the current work is how prior information is integrated with experiential feedback when making decisions under risk (Olschewski et al., 2023; Rakow & Newell, 2010).

Previous research has examined how decisions change when both prior descriptive information about choice options and experiential outcome feedback are available, compared to settings in which only descriptions or only feedback is available. (Erev et al., 2017; Jessup et al., 2008; Lejarraga & Gonzalez, 2011; Weiss-Cohen et al., 2016). For example, Jessup et al. (2008) gave two groups of participants the same risky choice problems and found that the group who were provided with feedback showed higher probability sensitivity than the "no-feedback" group. Subsequent work suggested that, when both descriptions about options and choice feedback are available, participants rely more on repeated experiential feedback than on statistical summaries of the options (i.e., descriptions) (Erev et al., 2017; Lejarraga & Gonzalez, 2011). This tendency appears stronger in more complex decision problems (Lejarraga, 2010; Weiss-Cohen et al., 2018) and when feedback is more reliable .

How do these conclusions generalize to settings in which decision makers have only *partial* knowledge about the options? In such settings, people may possess structural knowledge (information about the form of the environment, such as how many unique outcomes an option can produce), even if they do not know the specific outcomes or their probabilities. Several studies suggest that such structural knowledge can influence risky choice (Glöckner et al., 2016; He & Dai, 2022; Spektor & Wulff, 2024). For example, Glöckner et al. (2016) studied a group of participants who received information about the number of potential outcomes associated with each choice option. Participants then made one-shot decisions following costless sampling of outcomes[1]. This group exhibited over-weighting of small probabilities, contrary to earlier findings showing that experiential feedback typically leads to under-weighting of small probabilities. More broadly, in a comprehensive meta-analysis, Spektor and Wulff (2024) found that participants can obtain structural knowledge from their experience and adjusted their sampling behavior accordingly.

Although prior studies suggest that informing participants about the decision structure influences choice behavior, there is still no computational framework that quantifies how structural knowledge shapes underlying beliefs and decision pro-

[1]Although the free-sampling paradigm differs from the repeated-feedback paradigm, both involve participants receiving sequential feedback and have been shown to produce similar choice behavior (Hertwig & Erev, 2009)

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

cesses. In the current study, we test how structural knowledge influences behavior under repeated feedback and develop a Bayesian model to capture these effects. Because structural knowledge is expected to change decision makers' beliefs (particularly their uncertainty about the options), we also collected trial-by-trial confidence judgments (Lejarraga & Lejarraga, 2020; Olschewski & Scheibehenne, 2024; Salem-Garcia et al., 2023), which provide a direct measure of subjective uncertainty (Fleming, 2024). Together, our behavioral and modeling results offer insight into how prior information is integrated with feedback to guide risky choice, and how representations are updated to support learning.

Experiment

To investigate the influence of structural knowledge on choice behavior, we asked participants to make repeated choices between two options that generated probabilistic outcomes and manipulated the prior information available to participants. Specifically, half of the participants were informed of the number of potential outcomes associated with each option before receiving any feedback, whereas the other half received no such information. In the absence of structural knowledge, some participants might treat the first observed outcome from each option as fully diagnostic, effectively anchoring the option's value on that initial sample. We hypothesized that providing structural knowledge would change how participants represent option value, potentially leading to choice patterns more in line with equal weighting of the outcomes, as previous research has suggested (He & Dai, 2022). Further, we hypothesized that providing structural knowledge would reduce participants' uncertainty of the decision environment and as a result increase their confidence.

Method

Participants We recruited 137 participants from the psychology participant pool at the University of Warwick. Participants received course credit in addition to a variable performance-based bonus. Three participants were excluded because they did not choose each option at least 5 percent of the time, leaving us with 134 participants (116 female, mean age = 18.82, *SD* = 0.65).

Design In this task, participants made repeated choices between two options with unknown outcome distributions. The experiment included two between-subject factors. The first was type of feedback, which determined whether participants observed outcomes from both chosen and unchosen options (full feedback) or only from the option they chose (partial feedback). The second factor was presence of structural knowledge: In the "no-hint" condition, participants received no information prior to the start of the experiment. In the "number" condition, the number of potential outcome(s) present in the environment was shown on every trial screen[2]. This manipulation is similar to the ones used by Glöckner et al. (2016) and He and Dai (2022).

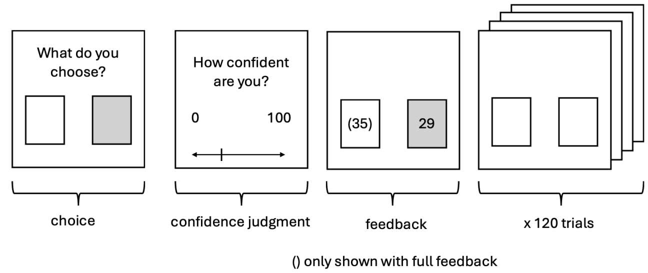

Figure 1: Illustration of the main choice task for the Experiment. Participants made repeated choices between two options for 30 trials in each decision environment. Each choice was followed by a confidence rating on a 0-100 scale that was set to 50 at the start of each trial. After submitting their choice and confidence rating, participants received feedback on their payoffs. In the full feedback condition, payoff from not chosen option (number in the parenthesis) was also shown.

The design also included two within-subject factors, which varied factorially across the four between-subject decision environments. We included these factors to test if the effect of structural knowledge would depend on certain features of the decision environment. The first factor was value of the rare outcome, which determined whether the rare outcome was a high-value "treasure" or a low-value "disaster". The second factor was option composition, which determined the types of options available in the environment: a safe–risky (SR) environment contained one safe option and one risky option, whereas a risky–risky (RR) environment contained two risky options. Each participant completed all four within-subject environments in a randomized order. The payoff distributions for each environment are reported in Table 1.

Table 1: Decision environments used in the experiment

Environment	Safe(r) Option	Risky(er) Option
Treasure-SR	35, 1	29, 0.8; 48
Disaster-SR	21, 1	27, 0.8; 8
Treasure-RR	35, 0.6; 50	30, 0.8; 70
Disaster-RR	45, 0.6; 30	50, 0.8; 10

Note: For each option, the table shows the most likely outcome, the probability of obtaining this outcome and the alternative outcome (if any). SR = safe–risky. RR = risky–risky.

Procedure In each environment, participants were asked to repeatedly choose between two options with the goal of maximizing their payoffs (points; see Figure 1). Each environment consisted of 30 trials. In each trial, participants first made their choice and then rated their confidence on a scale between 0

[2] In the safe vs risky problems, participants were told one option gives one potential outcome and the other gives two potential outcomes. They were not told which option gives one or two potential outcomes. In the risky vs risky problems, they were told both options give two potential outcomes.

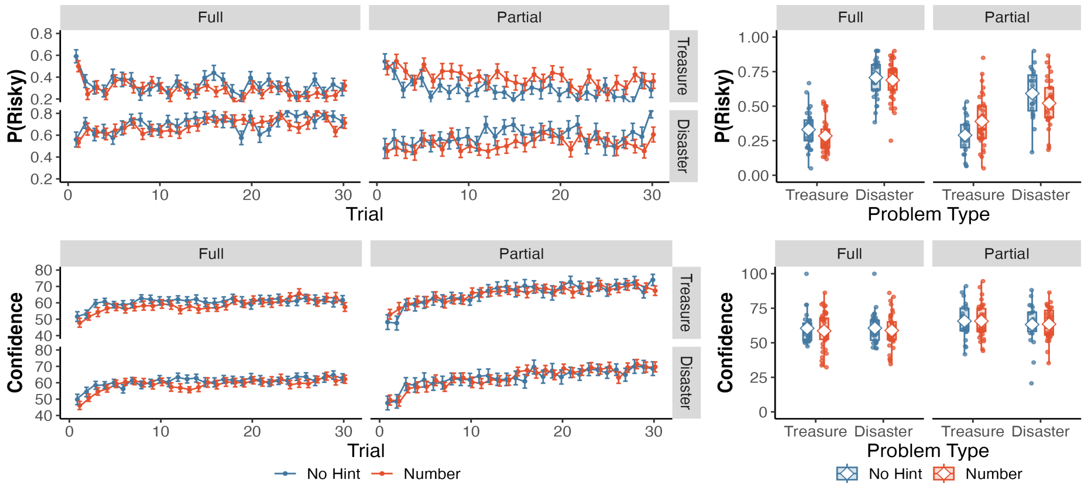

Figure 2: Proportion of risky choice (top) and confidence (bottom). Blue denotes the no-hint condition and orange denotes the number condition. The left panels show how averaged risky choice proportions and confidence vary across trials; dots and error bars represent means and standard errors of the mean. The right panels show participant-level averages of risky choice proportions and confidence. Each point represents a participant, and the white diamonds indicate mean values across participants.

(not at all confident) and 100 (totally confident) that they had made the best choice. After providing confidence, participants received feedback on their choice. A constant tally of the accumulated points for each environment was shown after they completed all trials within that environment, and their total points for the experiment were shown at the end of the experiment. Choice in this task was incentivized, with 1000 points equal to £0.15 of bonus payment. Confidence judgment was not incentivized.

Behavioral Results

To understand the effect of structural knowledge, we fit linear mixed effect models using the "mixed" function in the R package "afex" (Singmann & Kellen, 2019). In two separate models, we regressed the proportion of risky choice and confidence on trial number, structural knowledge (no-hint vs number), feedback type (full feedback vs partial feedback), rare outcome (treasure vs disaster), and option composition (SR vs RR). Both models included a random intercept for participants. This specification allowed us to test not only the main effect of structural knowledge, but also how it interacts with other experimental factors. We estimated all main effects and interactions, and used Kenward-Roger approximation for significance testing. Categorical variables were effect-coded and trial numbers were mean-centered. In the Results section, we report the estimated main effects and the statistically significant interactions involving structural knowledge. Estimates for terms that did not involve structural knowledge fell outside the scope of the present study and therefore not reported for brevity.

Figure 2 summarizes the behavioral results. For risky choice, the main effect of structural knowledge was not significant ($\beta = 0.004 \pm 0.007$, $p = .56$), with the overall proportion of risky choice not differing significantly between the no-hint (M = 0.49, SD = 0.23) and the number conditions (M = 0.47, SD = 0.22). However, there was a significant two-way interaction between rare outcome and structural knowledge (β = -0.018 ± 0.004, $p < .001$), and a significant three-way interaction between rare outcome, structural knowledge, and feedback type (β = 0.025 ± 0.004, $p < .001$). Specifically, in the rare treasure environment under partial feedback, participants chose the riskier option on 39% of trials (SD = 0.19) in the number condition, compared with 29% (SD = 0.11) in the no-hint condition (post-hoc test: β =-0.051 ± 0.020, $p <$ .05). In the rare disaster environments under partial feedback, participants chose the riskier option on 52% of trials (SD = 0.17) in the number condition, compared with 59% (SD = 0.18) in the no-hint condition (post-hoc test: β =0.073 ± 0.020, $p < .001$). For confidence, the main effect of structural knowledge was not significant (β = 0.419 ± 1.024, $p = .68$), and no interactions involving structural knowledge reached significant.[3]

Consistent with our hypothesis, the presence of structural knowledge brought the proportion of risky choice closer to 0.50 under partial feedback (He & Dai, 2022). Specifically, in rare-treasure environments structural knowledge increased risky choice, whereas in rare-disaster environments it decreased risky choice. One potential explanation is that, treating the two outcomes as more equally likely would increase the subjective value of the riskier option in treasure environments (by placing more weight on the rare high-value outcome) but decrease its value in disaster environments (by placing more

[3]We also ran the analysis separately on data obtained in safe-risky environments and in risk-risky environments and found there was no effect of structural knowledge on risky choice in the these environments

weight on the rare low-value outcome). The fact that this pattern emerged only under partial feedback suggests that participants relied more on structural knowledge when experiential feedback was limited. Contrary to our hypothesis, structural knowledge did not influence confidence judgments. In the next section, we introduce a computational model to examine more directly how structural knowledge changes participants' beliefs and, in turn, influenced experience-based risky choice.

Bayesian Model

Overview

Our goal was to develop a model that can predict both choice and confidence across experimental conditions and capture the effect of structural knowledge on choice proportions. We propose a model with separate learning and response components. The learning component assumes that participants have prior expectations about outcome probabilities and update their belief in a Bayesian way as outcomes are observed. Our modeling framework was motivated by prior work showing that participants learn the set of possible outcomes and their associated probabilities for each option (Mason et al., 2025; Szollosi et al., 2019), and by accounts that attribute the differences between description-based and experience-based risky choice to Bayesian belief-updating mechanisms (Aydogan, 2021). The Bayesian framework also provides a natural way to quantify the uncertainty participants have about each option, which can predict choice behavior in both stationary (Gershman, 2018) and changing environments (Speekenbrink & Konstantinidis, 2015) and thus allow us to better model confidence judgments.

Learning

We assume that before trial t, participants hold a subjective outcome distribution for option j, such that it yields outcome $R_1^{(j)}$ with probability $\theta^{(j,t)}$ and outcome $R_2^{(j)}$ with probability $1-\theta^{(j,t)}$. Here, for notation convenience, we define $R_1^{(j)}$ as the first outcome value that a participant observes from option j and $R_2^{(j)}$ as the other possible outcome value. Before $R_2^{(j)}$ is observed for the first time, it is assumed to be 100, which is an arbitrary value set to facilitate initial exploration (Gonzalez & Dutt, 2011). We model uncertainty about $\theta^{(j,t)}$ with a Beta distribution,

$$\theta^{(j,t)} \sim Beta(b_1^{(j,t)}, b_2^{(j,t)})$$

before trial 1 this becomes,

$$\theta^{(j,1)} \sim Beta(b_1^{(j,1)}, b_2^{(j,1)})$$

where $b_1^{(j,1)}$ and $b_2^{(j,1)}$ are pseudo-count parameters that capture prior evidence in favor of receiving $R_1^{(j)}$ and $R_2^{(j)}$ respectively. Under this parametrization, the prior expectation for receiving $R_1^{(j)}$ when participants first enter a given decision environment is

$$\mathbb{E}[\theta^{(j,1)}] = \frac{b_1^{(j,1)}}{b_1^{(j,1)} + b_2^{(j,1)}}$$

We assume that this expected value is the same for both options and denote it δ. We refer to δ as the prior bias parameter, which captures participants' prior belief about how common outcome $R_1^{(j)}$ is relative to outcome $R_2^{(j)}$. For parsimony, we assume that all participants share the same total amount of prior evidence (i.e., the same total number of pseudo-counts) by imposing the constraint $b_1^{(j,1)} + b_2^{(j,1)} = 2$, and estimate only how this evidence is allocated between the two outcomes across participants and conditions. When $\delta = 0.5$, $R_1^{(j)}$ is given equal weight as $R_2^{(j)}$. As δ increases, the prior assigns greater weight to $R_1^{(j)}$.

After receiving feedback on trial t, participants update their belief about outcome distributions. If the observed reward from option j is $R_1^{(j)}$, then $b_1^{(j,t+1)} = b_1^{(j,t)} + 1$ and $b_2^{(j,t+1)}$ is the same as $b_2^{(j,t)}$. In the full feedback condition, participants observe outcomes from both options, but previous research has shown that they tend to pay less attention to forgone payoffs (Palminteri et al., 2015; Yechiam & Busemeyer, 2005). To capture this asymmetry, we introduce an attention parameter $\omega \in [0, 1]$ to control how much participants learn from observed but foregone outcomes. In summary, b parameters are updated on each trial based on

$$b_1^{(j,t+1)} = \begin{cases} b_1^{(j,t)} + 1, & \text{if } j \text{ is chosen and } R^{(j,t)} = R_1^{(j)} \\ b_1^{(j,t)} + \omega, & \text{if } j \text{ is not chosen and } R^{(j,t)} = R_1^{(j)} \\ b_1^{(j,t)}, & \text{otherwise} \end{cases}$$

The updating from $b_2^{(j,t)}$ to $b_2^{(j,t+1)}$ is defined analogously.

Response

The learning component of the model captures participants' beliefs about outcome distributions on each trial, and can be used to predict both choice and confidence. Given $\theta^{(j,t)}$, the subjective probability of observing outcome $R_1^{(j)}$ for option j at trial t, we can quantify both the magnitude of expected rewards and the uncertainty associated with getting these rewards. The former is computed via the expected value,

$$\mathbb{E}[R^{(j,t)}] = \mathbb{E}[\theta^{(j,t)}]R_1^{(j)} + (1 - \mathbb{E}[\theta^{(j,t)}])R_2^{(j)},$$

and the latter is given by uncertainty in $\theta^{(j,t)}$, quantified as the standard deviation of the Beta distribution,

$$UC^{(j,t)} = \sqrt{\frac{b_1^{(j,t)} b_2^{(j,t)}}{(b_1^{(j,t)} + b_2^{(j,t)})^2 (b_1^{(j,t)} + b_2^{(j,t)} + 1)}}.$$

We assume that both factors influence choice and confidence.

For choice, we modeled the probability of choosing option j on trial t using a softmax function,

$$P(Choice^{(t)} = j) = \frac{e^{V^{(j,t)}}}{\sum_{k=1}^{2} e^{V^{(k,t)}}}.$$

Option values were computed using a linear combination of expected value and uncertainty, i.e.,

$$V^{(1,t)} = -\alpha_0 + \alpha_1 \mathbb{E}[R^{(1,t)}] + \alpha_2 UC^{(1,t)},$$

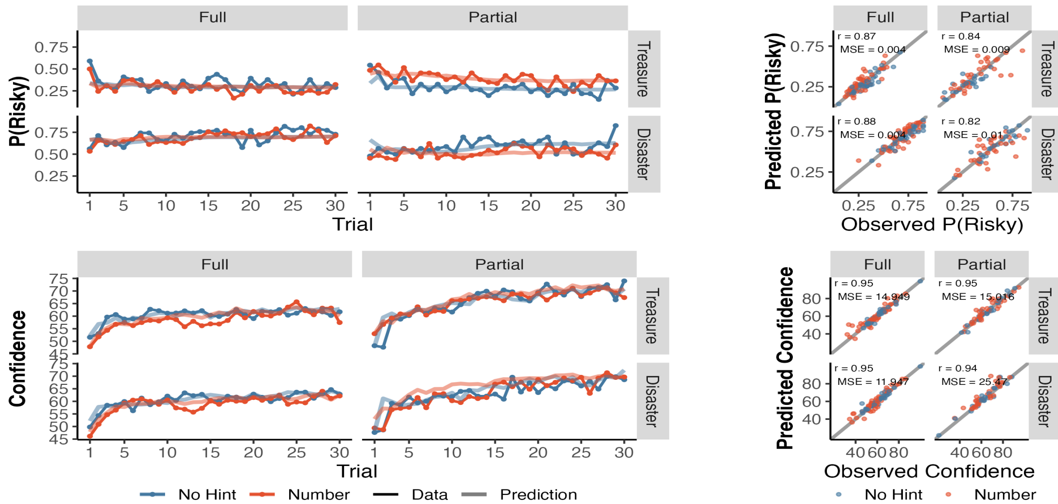

Figure 3: Model predictions for proportion of risky choice (top) and confidence (bottom). Blue denotes the no-hint condition and orange the number condition. Columns correspond to type of feedback (Full vs. Partial), and rows correspond to value of rare outcome (Treasure vs. Disaster). Left panels show trial-wise averages. Solid lines show observed data and shaded lines show model predictions. Right panels show individual-level averages. Predicted values are plotted against observed values. The diagonal indicates perfect prediction. r is the Pearson correlation coefficient and MSE is the mean square error.

$$V^{(2,t)} = \alpha_0 + \alpha_1 \mathbb{E}[R^{(2,t)}] + \alpha_2 UC^{(2,t)}.$$

Here α_0 captures a response bias toward choosing option 2 versus option 1. The addition of α_0 to the model has been shown to improve parameter recoverability (Wilson & Collins, 2019).

For confidence ratings, we modeled them using an ordered logit formulation with fixed thresholds between scale points. Conditional on the participant's choice of option j on trial t, the probability of reporting a confidence lower than $c \in \{1, 2, ..., 100\}$ is

$$P(Conf^{(t)} \leq c - 1) | Choice^{(t)} = j) = \frac{1}{1 + e^{-D_c^{(j,t)}}},$$

where

$$D_c^{(j,t)} = \tau_c - [\gamma_0 + \gamma_1 P^{(j,t)}(Better) + \gamma_2 UC^{(j,t)}].$$

Here τ_c is the c-th cut point; $\tau_c = 1, 2, ..., 100$, where $\tau_1 = 0$, $\tau_{c+1} = \tau_c + 0.1$. The cut points were set a priori for convenience and had minimal impact on model fit. γ_0 is an intercept; $P^{(j,t)}(Better)$ is the probability that the chosen option yields a better outcome than the other option, derived from the subjective outcome distributions; γ_1 captures how this probability influences confidence ratings; $UC^{(j,t)}$ is the standard deviation of the Beta distribution as defined above; γ_2 is the uncertainty sensitivity parameter.

Modeling Results

The model was fitted separately to data collected from each participant, with all conditions combined, using cmdstanr (Gabry et al., 2025) . The prior distributions for α and γ parameters were $N(0, 10)$. The prior distributions for δ and ω were $Beta(1, 1)$. Model fit was assessed using the leave-one-out crossed-validation information criterion (LOOIC) (Vehtari et al., 2017).

Figure 3 shows the full model's aggregate predictions of risky choice proportions and confidence judgments on the left, and individual-level predictions on the right. The model was able to capture the main aggregate patterns in the data as well as individual differences across conditions. It also captured the key behavioral pattern discussed in the Behavioral Results section: the proportion of risky choice was closer to 0.5 in the number condition than in the no-hint condition when participants received only limited, partial feedback.

Figure 4 shows the posterior means of the model parameters. For choice, the values of α_1 were significantly higher than 0 in both the no-hint condition (M = 0.11, SD = 0.11, V = 1548, p < .001) and the number condition (M = 0.10, SD = 0.12, V = 2971, p < .001), indicating that participants' choices were sensitive to expected value difference between the options. The values of α_2 were larger than 0 in the no-hint condition (M = 1.70, SD = 7.13, V = 1062, p < .05) and did not differ significantly from 0 in the number condition (M = -0.84, SD = 9.76, V = 1481, p = .77). This pattern is consistent with reduced uncertainty-driven exploration when structural knowledge was provided. There were no significant differences in parameter values between the no-hint condition and the number condition, except for α_2 which were marginally different (W = 2452, p = .23).

For confidence, the values of γ_1 were significantly higher than 0 in both the no-hint condition (M = 2.39, SD = 2.58, V = 1548, p < .001) and the number condition (M = 2.36, SD = 2.72, V = 2889, p < .001), indicating that participants' confidence judgments were sensitive to value difference in both

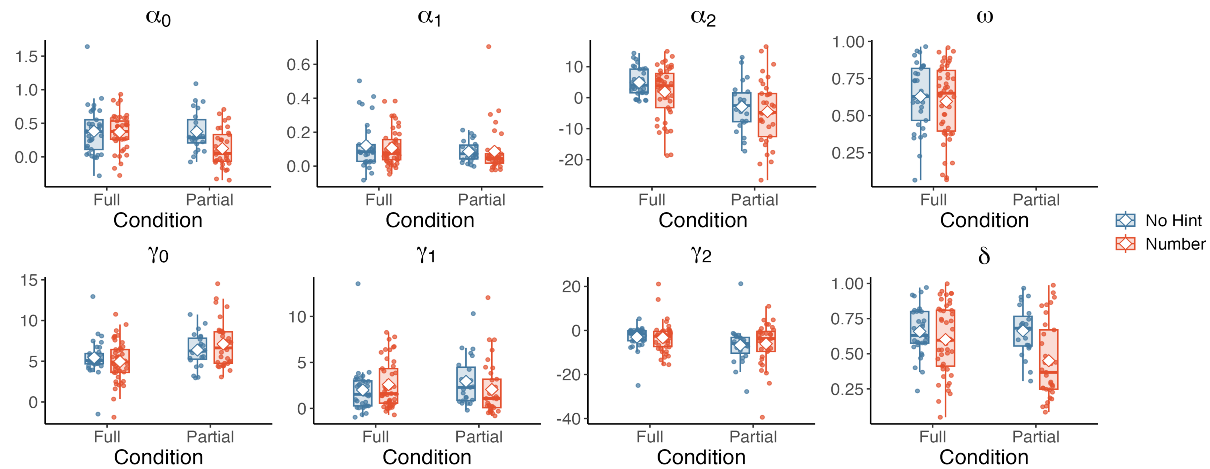

Figure 4: Distributions of posterior means of model parameters. Blue denotes the no-hint condition and orange the number condition. Each point shows parameter from one participant. Mean values are marked by white diamonds. ω was not estimated in the partial feedback condition therefore no values are shown.

conditions. The values of γ_2 were significantly lower than 0 in both the no-hint condition (M = -4.53, SD = 6.98), V = 132, p < .001), and the number condition (M = -4.43, SD = 8.17, V = 515, p < .001), suggesting that higher uncertainty of the chosen options reduced confidence in both conditions. There were no significant differences in parameter values between the no-hint condition and the number condition.

Given that the focus of our study was to investigate the effect of structural knowledge, we examined the values of δ, the prior bias parameter that captures participants' initial belief about outcome probabilities before any feedback is observed. While δ differed widely across individuals, we found evidence that biases were significantly higher (i.e., closer to 1 and further away from 0.5) in the no-hint condition (M = 0.66, SD = 0.17) than in the number condition (M = 0.54, SD = 0.27), W = 2775, p < .01. The direction of the difference is consistent with our expectation, in that informing participants about the potential number of outcomes reduced the tendency to anchor on the first observed outcome, resulting in a prior closer to equal likelihood for the two outcomes. We also assessed the importance of including the δ parameter in our model by comparing the full model to a constrained model with fixed prior bias (δ = 0.5). Allowing the prior-bias parameter to be freely estimated produced a non-significant improvement in model fit in the no-hint condition (ΔMean LOOIC = 5.44, V = 587, p = .08) and a smaller improvement in the number condition (ΔMean LOOIC = 1.91, V = 1320, p = .27).

To explain why structural knowledge brings the proportion of risky choice closer to 0.5 under partial feedback, we focused on participants in the partial-feedback condition and tested whether individual differences in fitted parameters accounted for their behavior. Specifically, we regressed each participant's proportion of risky choice in the rare-treasure and rare-disaster environments (analyzed separately) on the fitted parameters α_1, α_2, and δ. Consistent with the hypothesis that prior beliefs about the outcome distribution contribute to changes in risky choice, we found evidence that δ was associated with a higher proportion of risky choice in the rare treasure environments (β = -0.239 ± 0.093, p < .05) but not in the rare disaster environments (β = 0.014 ± 0.108, p = .89).

To demonstrate that the model can support reliable inference, we conducted a parameter recovery analysis. For each participant, we generated a synthetic dataset using the fitted model and the posterior mean of each estimated parameters. We then refitted the same model to these synthetic datasets. Parameter recovery was assessed by computing, for each parameter, the Pearson correlation between the generating values and the recovered estimates. All parameters except ω showed satisfactory recovery (Pearson's r > 0.6).

Concluding Remarks

In this study we investigated the effect of structural knowledge on risky experiential choice. We proposed a Bayesian model that explicitly accounts for prior beliefs about outcome distributions and allows these beliefs to be updated as experiential feedback is received, in accordance with Bayes' rule. We found that the tendency for risky choice proportions to approach 0.5 in the number condition could be partially accounted for by differences in participants' beliefs about outcome probabilities before receiving any feedback. Our model provides a general framework for incorporating different forms of structural knowledge in experiential risky choice. By precisely defining different types of structural knowledge and quantitatively estimating how they influence learning, choice, and confidence, our framework offers a step toward a more complete understanding of how people integrate prior information with feedback to update representations of choice options, and how these representations shape decision making under risk.

References

Aydogan, I. (2021). Prior beliefs and ambiguity attitudes in decision from experience. *Management Science*, *67*(11), 6934–6945. https://doi.org/10.1287/mnsc.2020.3841

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. https://doi.org/10.1002/bdm.443

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*, 369–409. https://doi.org/10.1037/rev0000062

Fleming, S. M. (2024). Metacognition and confidence: A review and synthesis. *Annual Review of Psychology*, *75*(Volume 75, 2024), 241–268. https://doi.org/10.1146/annurev-psych-022423-032425

Gabry, J., Češnovar, R., Johnson, A., & Bronder, S. (2025). *Cmdstanr: R interface to 'cmdstan'* [R package version 0.9.0, https://discourse.mc-stan.org]. https://mc-stan.org/cmdstanr/

Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, *173*, 34–42. https://doi.org/10.1016/j.cognition.2017.12.014

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), 486–508. https://doi.org/10.1037/a0040103

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*(4), 523–551. https://doi.org/10.1037/a0024558

He, Z., & Dai, J. (2022). Context-dependent outcome expectation contributes to experience-based risky choice. *Judgment and Decision Making*, *17*(1), 70–90. https://doi.org/10.1017/S1930297500009037

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

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, *19*(10), 1015–1022. https://doi.org/10.1111/j.1467-9280.2008.02193.x

Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making*, *23*(1), 100–116. https://doi.org/10.1002/bdm.666

Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, *116*(2), 286–295. https://doi.org/10.1016/j.obhdp.2011.05.001

Lejarraga, T., & Lejarraga, J. (2020). Confidence and the description–experience distinction. *Organizational Behavior and Human Decision Processes*, *161*, 201–212. https://doi.org/10.1016/j.obhdp.2020.07.003

Mason, A., Lindskog, M., Hertwig, R., & Wulff, D. U. (2025). Explicit memory representations in decisions from experience. *bioRxiv*, 2025.10.27.684917. https://doi.org/10.1101/2025.10.27.684917

Olschewski, S., Luckman, A., Mason, A., Ludvig, E. A., & Konstantinidis, E. (2023). The future of decisions from experience: Connecting real-world decision problems to cognitive processes. *Perspectives on Psychological Science*, *19*(1), 82–102. https://doi.org/10.1177/17456916231179138

Olschewski, S., & Scheibehenne, B. (2024). What's in a sample? epistemic uncertainty and metacognitive awareness in risk taking. *Cognitive Psychology*, *149*, 101642. https://doi.org/10.1016/j.cogpsych.2024.101642

Palminteri, S., Khamassi, M., Joffily, M., & Coricelli, G. (2015). Contextual modulation of value signals in reward and punishment learning. *Nature Communications*, *6*(1), 8096. https://doi.org/10.1038/ncomms9096

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*(1), 1–14. https://doi.org/10.1002/bdm.681

Salem-Garcia, N., Palminteri, S., & Lebreton, M. (2023). Linking confidence biases to reinforcement-learning processes. *Psychological Review*, *130*(4), 1017–1043. https://doi.org/10.1037/rev0000424

Singmann, H., & Kellen, D. (2019). An introduction to mixed models for experimental psychology. In *New methods in cognitive psychology* (pp. 4–31). Routledge. https://doi.org/10.4324/9780429318405-2

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, *7*(2), 351–367. https://doi.org/10.1111/tops.12145

Spektor, M. S., & Wulff, D. U. (2024). Predecisional information search adaptively reduces three types of uncertainty. *Proceedings of the National Academy of Sciences*, *121*(47), e2311714121. https://doi.org/10.1073/pnas.2311714121

Szollosi, A., Liang, G., Konstantinidis, E., Donkin, C., & Newell, B. R. (2019). Simultaneous underweighting and overestimation of rare events: Unpacking a paradox. *Journal of Experimental Psychology: General*, *148*(12), 2207–2217. https://doi.org/10.1037/xge0000603

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing*, *27*(5), 1413–1432.

Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, *135*, 55–69. https://doi.org/10.1016/j.obhdp.2016.05.005

Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2018). Task complexity moderates the influence of descriptions in decisions from experience. *Cognition*, *170*, 209–227. https://doi.org/10.1016/j.cognition.2017.10.005

Wilson, R. C., & Collins, A. G. E. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, *8*, e49547. https://doi.org/10.7554/eLife.49547

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*(3), 387–402. https://doi.org/10.3758/BF03193783