



OPEN

Quantifying the immediate computational effects of preceding outcomes on subsequent risky choices

Hayley R. Brooks & Peter Sokol-Hessner

Forty years ago, prospect theory introduced the notion that risky options are evaluated relative to their recent context, causing a significant shift in the study of risky monetary decision-making in psychology, economics, and neuroscience. Despite the central role of past experiences, it remains unclear whether, how, and how much past experiences quantitatively influence risky monetary choices moment-to-moment in a nominally learning-free setting. We analyzed a large dataset of risky monetary choices with trial-by-trial feedback to quantify how past experiences, or recent events, influence risky choice behavior and the underlying processes. We found larger recent outcomes both negatively influence subsequent risk-taking and positively influence the weight put on potential losses. Using a hierarchical Bayesian framework to fit a modified version of prospect theory, we demonstrated that the same risks will be evaluated differently given different past experiences. The computations underlying risky decision-making are fundamentally dynamic, even if the environment is not.

Prospect theory is conceivably the most successful recent theory in risky decision-making¹. The theory is founded on the idea that people evaluate options relative to “the past and present context of experience” (p. 277)¹. While this central insight led to two of the most successful components of prospect theory, the reflection effect in risk aversion (diminishing marginal utility of values expressed as concavity for gains and convexity for losses) and loss aversion (overweighting losses compared to gains of the same magnitude), it is unclear how recent events *quantitatively* influence risky monetary decision-making, moment-to-moment. We sought to quantify how recent events (e.g. previous outcomes) influence choices by analyzing a large dataset of risky choices in the presence of feedback compiled from four studies^{2–5}.

Laboratory and field studies have suggested that recent events shape the subjective value of risky choice options^{6–9} and actions¹⁰. The directionality and mechanism supporting these effects are unclear as some studies of risky decision-making in the presence of feedback have found that previous gains, relative to losses, lead to less risk-taking^{6,10} or a mixture of less and more risk-taking^{7–9,11}, discrepancies that may be related to how and when outcomes are realized, or paid¹². Theories that have sought to specify the directionality of the effect of feedback on risky choice have not described the size of the effect, its characteristics (e.g. how far it extends in time) or the precise decision processes that would be altered by feedback¹³. Interpreting these previous findings is complicated by the fact that almost all of these studies included design features intended to elicit temporal context effects (e.g. explicitly-signaled contexts, displays with cumulative elements) and thus may have strengthened, or even created, dynamic context effects. While such an approach identifies an upper bound on the possible presence and magnitude of context effects, such effects may not exist in less context-focused environments. Many more studies with trial-by-trial feedback have not analyzed or discussed the effects of feedback on risk-taking^{2–5,14–24}. We thus sought to examine whether, and to what extent, context affects risky decision-making in a setting with no intentional temporal structure. If found, this would establish that the effect of recent events on risky choices is not merely elicited when designs encourage it, but rather that it is part of the fundamental computations underlying risky decision-making.

“Risky decision-making” here refers to decisions with choice options characterized by known values and probabilities²⁵. In such a setting, all choice attributes are explicit and lack temporal dependency (i.e. outcomes or choice attributes on previous trials are unrelated to subsequent trials). Indeed, in the studies compiled for the

Department of Psychology, University of Denver, Denver, CO, USA. ✉e-mail: Peter.Sokol-Hessner@du.edu

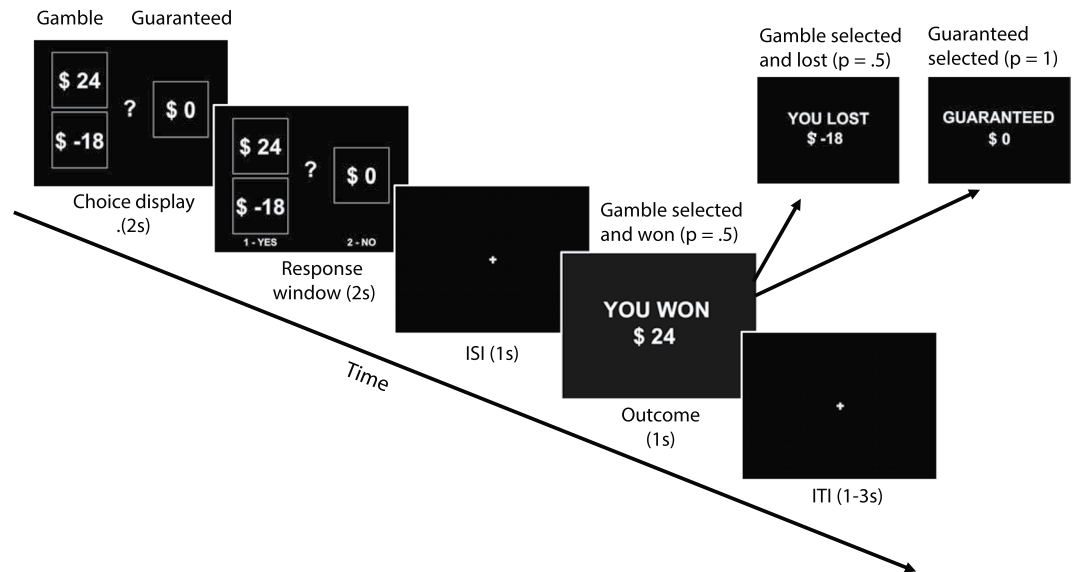


Figure 1. Risky monetary decision-making task structure. Across all four studies^{2–5} compiled for the current analysis, participants made a series of choices between risky and safe options. Every trial featured a risky option with two possible outcomes, each received with 50% probability if that option was selected, and a safe, or guaranteed option received with 100% probability if selected. All choices were followed by outcomes prior to the next choice appearing on the screen. See Supplementary Material and Table S2 for details on the minor variations across studies.

current analysis, all participants passed comprehension quizzes confirming their knowledge of the probabilities and values in the experiments. That there are no probabilities or values to learn or update means that if participants try to do so, they cannot improve and may instead impair their performance. If a dependency is observed, it would establish a fundamental, robust, and unappreciated role for context in affecting subjective value and/or action under risk.

Here, we analyzed a large dataset of risky choice behavior collated across four studies. These studies used an identical risky decision-making task, lacking in temporal dependency, in which participants chose between a risky gamble and a smaller, guaranteed alternative on each trial and received feedback following each trial (see Fig. 1 and Methods for task details)^{2–5}. Here, we use “gambling” and “risk-taking” interchangeably to refer to when participants choose the risky option over the guaranteed alternative. First, we fit mixed effects logistic regression models to the data to examine which previous and current choice attributes and events influence risk-taking. We then used a hierarchical Bayesian framework to fit a modified prospect theory model, simultaneously capturing individual differences and group-level effects across a large sample of participants to identify whether and how temporal context influences the processes underlying risky decision-making. This model captures several qualitatively distinct and dissociable processes that contribute to risky choice behavior: loss aversion, risk aversion, choice consistency, and decision bias (see below for more details on each). For each individual, we estimated baseline values for these four valuation and decision processes, and the extent to which each of these processes changed as a function of recent events. See Methods section and Supplementary Material for more details.

Results

First, we fit mixed effects logistic regression models regressing binary choice (choice of the gamble = 1, guaranteed alternative = 0) on a variety of variables, and second, we fit a hierarchical Bayesian model of the underlying nonlinear processes.

Mixed effects logistic regression models. We fit mixed effects logistic regression models to the risky choice data using the “lme4” package in R (R version 3.5.0; “lme4” version 1.1–21)²⁶. For this analysis, we discuss three types of recent events: previous outcome (the value in dollars of the outcome on the previous trial), previous decision (gamble or guaranteed, coded as 1 for gamble and 0 for guaranteed), and mean expected value (EV) of the previous choice options (the mean of the safe and risky options on the previous trial; see Methods). All mixed effects logistic regression models included binary choice as the binomial outcome variable and both a constant and current choice options (the dollar values of the risky gain, risky loss, and guaranteed alternative) as predictor variables (e.g. $\text{choice}_t \sim \beta_0 + \beta_1 \times \text{risky gain}_t + \beta_2 \times \text{risky loss}_t + \beta_3 \times \text{guaranteed}_t + \beta_4 \times \text{outcome}_{t-1}$). The constant was modeled as a random effect and the choice options were modeled as fixed effects. Each model varied only by the type of recent events included as additional predictor variables (for reliable convergence all were modeled as fixed effects). See Supplementary Material and Table S3 for lme4 code for all models.

Across all mixed effects models discussed in this analysis, current choice options had consistent, significant effects on binary choice, such that larger potential gains led to more risk-taking and larger potential losses or guaranteed alternatives led to less risk-taking. Hereafter, we only discuss the additional effects of recent events on

risk-taking, controlling for effects of current options. See Supplementary Material and Table S3 for full regression results, including the effects of current choice options.

To first examine which types of recent events had the strongest effects on subsequent risk-taking, we performed three separate regressions (models 1–3), each featuring one of the following: the previous outcome (model 1), the previous decision (model 2), or the mean EV of the previous choice option to capture the average magnitude of the previous trial (model 3). Because the number of parameters in each of these models was identical, to compare models, we used log-likelihood values, or the degree to which each model produced choice likelihoods that reflected the actual choices on each trial. The best-fitting model had the highest (least negative) log-likelihood value.

We found that previous outcomes (model 1) had a significant, negative influence on binary choice, such that as outcomes on the previous trial increased in magnitude, risk-taking on the current trial decreased ($\beta = -0.03$ (0.003), $p < 2 \times 10^{-16}$). This model (model 1, log-likelihood = $-10,638.5$) outperformed models that, instead of the previous outcome amount, featured previous decisions (model 2, log-likelihood = $-10,682.7$) or the mean EV of the previous choice options (model 3, log-likelihood = $-10,669.2$). When all three regressors (previous outcomes, decisions, and mean EV) were pit directly against each other by including them in the same model, model 4, previous outcomes predicted risk-taking on the current trial beyond previous decision and mean EV of the previous choice options (previous outcome: $\beta = -0.03$ (0.004), $p = 4.08 \times 10^{-15}$; previous decision: $\beta = -0.04$ (0.02), $p = 0.06$; mean EV of previous choice options: $\beta = -0.002$ (0.006), $p = 0.78$).

To illustrate the effect size of past outcomes on a given choice, assuming that the participant was indifferent between the current choice options (i.e. the value of the gamble and guaranteed alternative currently under consideration were equal), the estimates from model 1 indicate that the probability of choosing the risky option following a gain of +\$20 was 35%, while that following a loss of -\$20 was 65% (Fig. 2a), a 30% difference in the probability of risk-taking.

Having established that the effect of previous events on current binary choices was most clearly accounted for by previous outcomes, we subsequently sought to establish the characteristics of that effect over time, valence, and outcome type in models 5–8. First, we examined how far the effect of previous outcome extended over time. In model 5, we regressed binary choice onto the outcome one trial back (t-1), two trials back (t-2), and three trials back (t-3). We found that the negative effect of previous outcomes declined over time, such that it was significant at one and two trials back, but not three trials back (outcome_{t-1}: $\beta = -0.03$ (0.003), $p < 2 \times 10^{-16}$; outcome_{t-2}: $\beta = -0.008$ (0.003), $p = 0.01$; outcome_{t-3}: $\beta = -0.004$ (0.003), $p = 0.19$; Fig. 2b).

Next, in model 6, we tested whether the previous outcome effect was due to the outcome amount or simply the outcome valence, as previous outcomes included gain, zero, and loss amounts. Regressing binary choice onto regressors for previous outcome amount and for previous outcome valence (modeled as +1 for gains, 0 for zero, and -1 for losses) in model 6, we found that the amount of the previous outcome predicted risk-taking on the current trial beyond the valence of the previous outcome (previous outcome amount: $\beta = -0.04$ (0.005), $p < 5.95 \times 10^{-14}$; previous outcome valence: $\beta = 0.08$ (0.04), $p = 0.08$).

Finally, we examined whether the previous outcome effect differed as a function of outcome type (gain, loss, or guaranteed). Differences on the basis of outcome type could arise from differences between gains (risky and safe) and losses due to loss aversion^{27,28}, or differences between risky outcomes (gain and loss) and guaranteed outcomes due to possible expectation-based learning processes (for example, processes reflecting prediction errors^{29–31}) as well as more complex interactions between valence and uncertainty. In model 7, we thus regressed binary choice onto three separate regressors, one each for previous risky gain outcome amount, previous risky loss outcome amount, and previous guaranteed outcome amount. All three regressors were unsurprisingly significant, replicating the overall effect of previous outcomes on binary choice (previous gain: $\beta = -0.03$ (0.004), $p < 2 \times 10^{-16}$; previous loss: $\beta = -0.03$ (0.008), $p = 0.0003$; previous guaranteed: $\beta = -0.02$ (0.009), $p = 0.005$). Furthermore, these three coefficients were not significantly different from one another as tested in two ways. First, model 7 (with separate coefficients for previous risky gain, risky loss, and guaranteed outcomes) did not perform significantly better than model 1 (with a single previous outcome amount regressor; because model 1 is a nested version of model 7, we used the likelihood ratio test; LRT statistic = 1.17, $df = 2$, $p = 0.56$). Second, estimates for the effects of previous gain, loss, and guaranteed outcomes in model 7 were not significantly different from one another in any of the three possible pairwise comparisons (Wald tests, all three p 's > 0.31 ; Fig. 2c). Together, these findings demonstrate that the effect of previous outcomes on the current binary choice did not differ by outcome type.

For an alternative approach to the analysis of outcome type, in model 8, we regressed binary choice on two regressors that represented outcome amount separately by valence only (with regressors for previous gain amount collapsing across risky and guaranteed outcome amounts, and previous loss amount), and included a separate term for previous risky versus guaranteed choices (coded +1/−1, respectively, as in the other models) which was interacted with previous gain amount. Model 8 similarly identified significant effects of previous gain and loss amounts (previous gain amount: $\beta = -0.03$ (0.005), $p = 3.22 \times 10^{-8}$; previous loss amount: $\beta = -0.04$ (0.009), $p = 5.49 \times 10^{-5}$) that were not significantly different from each other (Wald test, $p = 0.44$). Model 8 also replicated model 4's finding of a weak effect of previous choices ($\beta = -0.05$ (0.02), $p = 0.06$) when accounting for previous outcomes. Finally, the interaction of previous gain outcome amount with previous risky versus guaranteed choices was not significant ($\beta = -0.0009$ (0.005), $p = 0.86$). Model 8's results are consistent with models 6 and 7 in identifying that the effect of previous outcomes does not significantly vary by the valence of the outcome or by the type of previous outcome, but is instead best described as an effect of the previous outcome amount, whatever its type or valence.

Thus far, we established that the effect of previous outcome on the current binary choice was short-lasting (Model 5), driven by outcome amount (Model 6), and did not differ by outcome type or valence (Models 7 and 8).

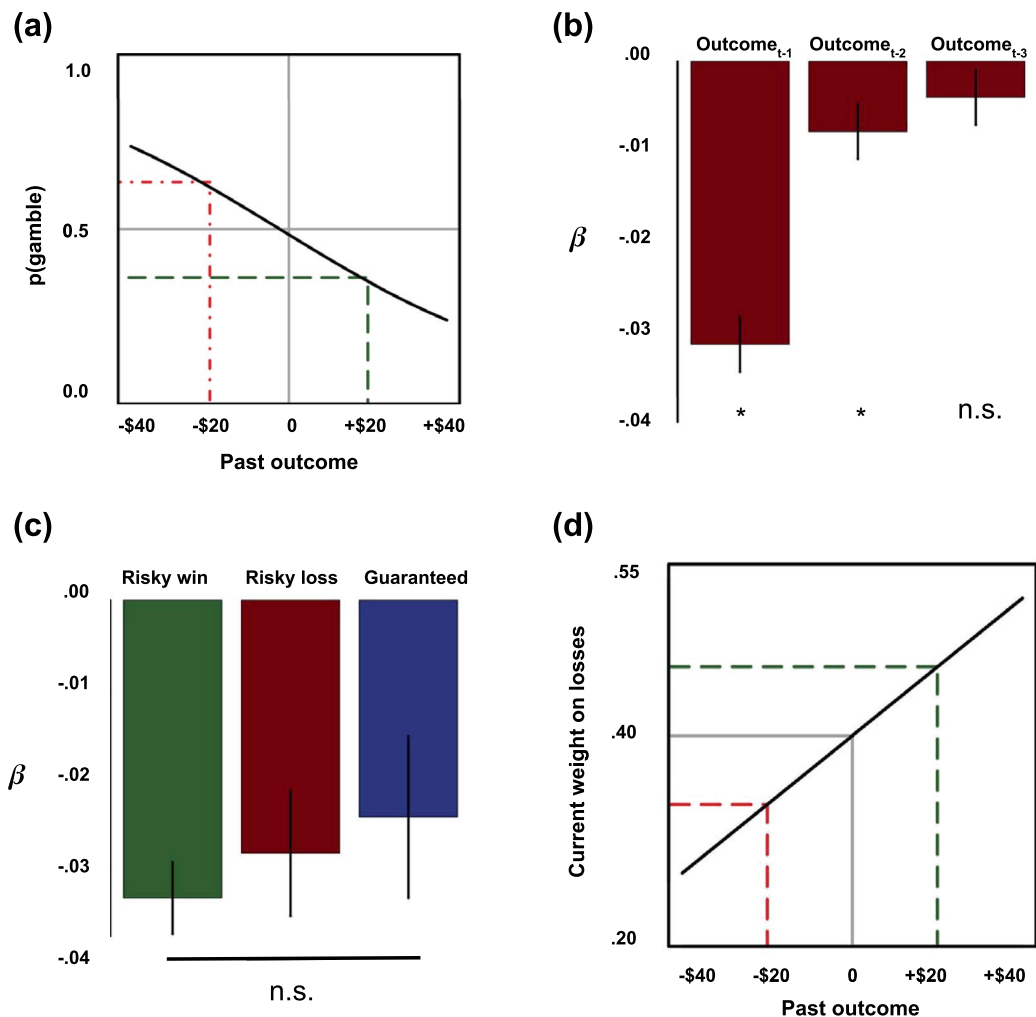


Figure 2. Effects of previous outcomes on current risky choice. **(a)** Visualizing the effect size of previous outcomes of -\$20/+ \$20 on the current choice, assuming indifference on the current trial (model 1). **(b)** Estimates of the effect of past outcomes up to three trials back ($t-1$, $t-2$, $t-3$) on risk-taking on the current trial in model 5 (errors bars represent standard errors). **(c)** Testing whether the effect of previous outcomes varies by type of previous outcome (risky gain, risky loss, or guaranteed alternative) in model 7 (errors bars represent standard errors). No pairwise comparisons were significant (Wald's Test, all p 's > 0.31), and model 7 was not significantly better than model 1, which collapsed across outcome types (likelihood ratio test, $p = 0.56$). **(d)** Visualizing the effect size of previous outcomes of -\$20/+ \$20 on the current weight of losses (model 9).

All of the above effects are value-independent in that recent events directly shift the probability of choosing the risky option. However, it is unclear whether recent events may additionally influence valuation of the current choice options. To test whether previous outcomes influence how individuals assess the value of the options on the current trial, in model 9 we regressed binary choice onto previous outcome and included three interaction terms between previous outcome and each of the current choice options (risky gain, risky loss, and guaranteed). While the value-independent effect of previous outcomes on the current binary choice remained ($\beta = -0.46(0.21)$, $p = 0.03$), we additionally found that the outcome on the previous trial increased the weight put on the potential losses (outcome_{t-1} × risky loss; $\beta = 0.1(0.03)$, $p = 0.002$), but not potential gains (outcome_{t-1} × risky gain; $\beta = 0.01(0.04)$, $p = 0.8$) or potential guaranteed alternatives (outcome_{t-1} × guaranteed; $\beta = -0.06(0.08)$, $p = 0.5$). To illustrate the effect size of past outcomes on the valuation of subsequent potential losses, the weight placed on potential losses after an outcome of + \$20 was 0.46, or 142% of the weight on losses after an outcome of -\$20, which was 0.33 (Fig. 2d).

Hierarchical Bayesian estimation. Mixed effects logistic regressions allow us to detect individual- and group-level differences for linear processes, but we know that some risky decision-making processes such as risk aversion are not linear^{2-5,21,32}.

To address this, for the second analysis we used hierarchical Bayesian estimation, an approach that allowed the fitting of all of the decision-making data at once while simultaneously estimating both individual- and group-level parameters of nonlinear models (like prospect theory) previously shown to fit these data well^{2-5,10,21,32-34}. We used




Figure 3. Representative moment-to-moment changes in the loss aversion parameter (λ) for a single hypothetical individual as a function of three different sets of outcome experiences ((A) yellow, (B) crimson, and (C) dark red lines), illustrating that the same person will have different preferences, and therefore make different choices, if they have had different recent experiences. The initial loss aversion parameter value (1.57) and the value of δ^λ (0.01) used for this effect size illustration are the group mean values. Note that any given participant's initial loss aversion or updating parameter values may have been higher or lower than these (see Results for confidence intervals). Outcome histories used here are real and were experienced by three different participants².

Markov-Chain Monte Carlo (MCMC) sampling techniques in a hierarchical Bayesian framework (using “rstan” version 2.17.3)³⁵ to fit a modified version of a 4-parameter prospect theory-inspired model (Prospect Theory Plus, PT +)¹⁰. PT + captured four distinct decision-making processes: risk aversion (ρ), loss aversion (λ), choice consistency (μ), and decision bias (db), assuming a linear probability weighting function. Using this Bayesian framework, we estimated both the group-level distribution and individual values for each of the four PT + parameters. We also modeled the change in each of these four parameters over time by altering PT + to include four additional updating parameters (four δ^θ terms, i.e. δ^ρ , δ^λ , δ^μ , δ^{db}). The updating parameters controlled how much previous decision outcomes shifted each of the four original PT + parameters on the subsequent trials. Positive (negative) values of δ^θ would indicate that θ increased (decreased) as previous outcomes increased, while a zero value of δ^θ would indicate no net adjustment of θ by previous outcomes. Because PT + consists of value-dependent parameters (ρ , λ), a value-independent parameter (db), and an intermediate parameter linking value and action (μ), we were able to simultaneously detect the effects of context on linear and non-linear action- and valuation-related decision-making processes. See Methods for the complete modeling procedure.

MCMC estimation procedures produce “chains” of sampled parameter values, in proportion to their likelihood. Using Stan (“rstan” version 2.17.3)³⁵, we ran twenty chains of 10,000 samples each, discarding the first 5,000 samples of each chain as a burn-in period, resulting in 100,000 samples. Priors were selected to be as uninformative as possible and were normal, uniform or Cauchy distributions, described in more detail in Table S1. Each of the twenty chains converged on similar distributions of parameter values (mean Rhat for group-level mean parameters = 1.002, range = 1.0003–1.0034; ideal = 1). The total number of effective samples for each of the group-level mean parameters were $\rho = 10,782$, $\lambda = 100,000$, $\mu = 100,000$, $db = 12,783$, $\delta^\rho = 6,292$, $\delta^\lambda = 13,003$, $\delta^\mu = 12,594$, and $\delta^{db} = 6,592$.

First, we examined mean values and 95% confidence intervals for each of the baseline group-level parameter estimates for risk aversion, loss aversion, choice consistency and decision bias. In our sample, participants were risk averse for gains and risk seeking for losses ($\rho = 0.65$, 95% CI = [0.58 0.73]), were mildly loss averse ($\lambda = 1.57$, 95% CI = [1.44 1.71]), were consistent in their choices for the risky option ($\mu = 22.2$, 95% CI = [18.9 26.0]), and had a bias to gamble ($db = -0.58$, 95% CI = [-0.68 -0.47]), consistent with others' findings^{2–5,10,21,32}.

Next, we tested whether each of the four parameters updated as a function of previous outcomes. Examining the mean values and 95% CIs for each of the four group-level mean updating parameters, we found that large positive previous outcomes increased loss aversion ($\delta^\lambda = 0.013$, 95% CI = [0.003 0.02]), consistent with our finding from the linear mixed effects logistic regression analyses that large previous outcomes increased the weight put on potential losses (see above, model 9). We also found that large positive previous outcomes increased consistency across binary choices ($\delta^\mu = 0.08$, 95% CI = [0.05 0.10]) and reduced the bias to gamble ($\delta^{db} = 0.03$, 95% CI = [0.02 0.05]). We found no effect of previous outcomes on risk aversion ($\delta^\rho = 0.005$, 95% CI = [-0.01 0.02]). The variation identified in these confidence intervals means that for identical previous outcomes, different people will react differently. More importantly, the finding of any effect of recent outcomes also means that if otherwise initially-similar people have different experiences, they will subsequently make systematically different risky choices. Figure 3 illustrates this effect by plotting loss aversion over time for one individual experiencing three different histories of outcomes. Risky choice behavior is not just about the individual decision-maker and their choice preferences, but also their past experiences.

Discussion

Here, we have quantified the consistent and strong effects of recent outcomes on subsequent risky choices. We have demonstrated that the size of the monetary amounts received in previous outcomes has a negative relationship with current risky choices, the effect is short-lasting and linear across outcome types, and recent events appear to have value-independent and value-dependent effects on current actions. These findings are especially notable given the absence of explicit manipulations or design features encouraging integration of previous events with current choices, and the explicit, learning-free setting featuring fully-known values and probabilities. This differs from studies (e.g. on the decision by sampling theory^{36–40} or others^{7–11}) that are designed around contextual influence. That risk-taking changes as a function of contextual manipulations in those studies likely reflects both the intrinsic contextual sensitivity of risk-taking as well as the success of the experimental manipulations, but it is unclear how much of the overall effect is aligned with which. We provide strong evidence here that risky monetary decision-making is in fact fundamentally dynamic, even when the environment is decidedly not.

While many dynamic behaviors are well-explained by learning, it is not a compelling explanation here for two reasons: the task, and participants' reports. First, all monetary values in the task were explicitly shown on every trial and detailed instructions explicitly communicated that probabilities were stable (and simple, being either 100% or 50% in all cases) across all trials. This means that it is unclear what could have been learned in the first place as no elements of the task were hidden, implicit, or otherwise unknown at any time. Second, all participants completed quizzes prior to data collection directly assessing comprehension of the explicit, non-dynamic nature of all task elements, making it unlikely that any of them went on to try to learn about any of these elements. Thus, the major remaining possible explanations for the observed dynamic effects are shifts in the processes underlying valuation and action.

We reported effects of previous outcomes on three of the four decision processes that we modeled (loss aversion, choice consistency, and decision bias). It is important to note that there are additional decision-making processes not included in our model that previous outcomes may influence on a trial-by-trial basis (e.g. probability weighting)⁴¹. Future research should address additional decision-making processes to more completely identify which processes are affected by previous events and the magnitude and characteristics of that influence.

The strong effects we identified establishes that the influence of temporal context on risk is so fundamental and robust that it is visible even if existing risky decision-making studies and models are not designed to capture that influence. Conceptually similar recent studies have examined the implications of dynamic context effects for static models of value-based decision-making^{38,42}, while others have begun to develop inherently contextual models of cognition and decision-making using frameworks like normalization^{43–45}, sampling^{36,46–49}, and others^{10,50}. However, none of these models yet account for the trial-by-trial effects of previous outcomes on subsequent valuation and choices observed here. The development of sophisticated, contextually-sensitive models will be critical if future studies are to optimally tune their designs to further characterize and quantify the effects of context on valuation and action.

The creation of such models will also be indispensable to efforts to identify the underlying affective and neural correlates. While affect is clearly related to risky decision-making^{2,4,19,51,52}, it is still unclear how affect may be related to the effects of context in the same domain. Theories have suggested a variety of possible ways affect could interact with or represent context, including that positive affect in a given moment guides future behavior to maintain that affect¹³; that affective moods may represent recent experiences in reinforcement learning⁵³; or that affect may have an even broader and more fundamental role in representing a decision-maker's context^{54,55}. Given the recent finding that the central neural systems of valuation and choice may represent some contextual variables¹⁰, it seems likely that regions like the ventral striatum, amygdala, and ventromedial prefrontal cortex, which are deeply involved in decision-making^{51,56,57}, are primary candidates for the neural mechanisms of context effects on risky decision-making.

Together, these findings suggest that risky monetary decision-making with moment-to-moment feedback is inherently dynamic and that a context-dependent, time-sensitive model is necessary for understanding temporal context effects underlying valuation and risky choice behavior. This critical insight will not only advance the study of risky monetary decision-making but may have real-world implications in settings where people may have different experiences (e.g. in aging, or across socioeconomic statuses), or think and feel differently about those experiences (e.g. in depression). Risky choices must be considered not only as a product of the individual decision-maker, but of what they have experienced.

Method

Data. We analyzed data collated across four previously published studies of risky monetary decision-making^{2–5}. Each study had between 30 and 120 participants completing between 140 and 180 trials of a structurally identical risky monetary decision-making task (see Supplementary Material and Table S2 for a more in-depth discussion of the minor methodological variations across the four studies) in a fully-instructed laboratory setting (including detailed and guided task instruction, comprehension quizzes, and practice trials), producing high-quality incentive-compatible individual-level choice data. The complete dataset comprised a total of 64,953 binary choices made by a total of 234 participants (128 females, mean age = 23.4 (4.8), median age = 22; 207 missed trials were removed from analysis). To minimize confounding influences of the different manipulations and structure present in the four studies, portions of the data were removed from this analysis so that all data used reflected participants' first experience with the task, in the absence of any external manipulation. To accomplish this, we removed all choices made on Day 2^{2,4,5}, made with an emotion regulation-like strategy², made under the influence of a medication⁴, or made after a stress manipulation⁵. The remaining data consisted of 23,373 choices made by a total of 151 participants (88 females), mean (sd) age of 23.4 (4.7) years old (median age = 22). All four studies were IRB approved and featured standard consent procedures. For more detailed information, see the

individual studies²⁻⁵. The University of Denver's Institutional Review Board determined that the current project did not require IRB review (January 18, 2019).

Task. All studies had the same basic instruction and trial-level structure. Each began with detailed task instructions that were reviewed with the experimenter, a brief quiz to assess task comprehension, and a series of practice trials to ensure participants understood the task and were fully informed. In the main task, participants made a series of binary choices between risky gambles and guaranteed alternatives (see Fig. 1). The gamble always offered two potential outcomes with equal (50%) probability, while the guaranteed alternative delivered a single outcome with 100% probability (if chosen). Choice options were displayed for 2–4 s, followed by a response window of 2 s, then an interstimulus interval of 1 s before the choice outcome was realized for 1 s. Intertrial intervals varied between 1–3 s. The majority of trials were gain-loss trials in which potential risky gamble outcomes featured a gain amount (\$2–\$12) or a corresponding loss amount that was a negative multiple ($\times 0.25$ – $\times 2$) of the gain amount. The remaining trial types were gain-only in which potential risky gamble outcomes featured a gain amount (\$2–\$30) or an outcome of \$0, and a smaller guaranteed alternative (\$1–\$12). Every participant completed both gain-loss and gain-only trial types intermixed. At the end of the task, participants were paid a portion (10%) of trial outcomes relative to an initial endowment of \$30, making all four studies incentive-compatible. The four studies varied slightly in total number of trials and blocks and the number of gain-loss and gain-only trials. See Supplementary Material and Table S2 for details on the minor variations across studies.

Analysis: logistic regression models. To examine whether recent events influence subsequent risky choice behavior, we fit mixed effects logistic regression models to the risky binary choice data using the package “lme4” in R²⁶. All generalized mixed effects models reported here are logistic regressions that use a linear combination of predictors and a binomial link function. We regressed binary choice on current choice options (including current risky gain, risky loss, and guaranteed alternative amounts), and one or more elements of previous events, including previous outcomes, previous decisions, and/or mean expected value of the previous choice options. When applicable, mean expected value on a given trial was calculated as $[(0.5 \times \text{risky gain} + 0.5 \times \text{risky loss}) + (1 \times \text{guaranteed})]/2$. In general, regressions included an intercept, current choice options, and previous choice outcomes/decisions/options as fixed effects and a random intercept for each participant. See Table S3 for lme4 code corresponding to each regression (models 1–9) in the Results.

We tested the output of model 7 in two ways: testing the entire model against model 1 with a likelihood ratio test (because model 1 is nested within model 7, that is, model 1 is a special case of model 7, in which the parameter estimates for previous risky gains, risky losses, and safe outcomes are equal) and testing the equality of specific pairs of coefficients with Wald Tests. The model comparison likelihood ratio test statistic was chi-square distributed with two degrees of freedom for the two extra regressors. The three Wald tests compared the coefficients estimated in model 7 for the effects of previous gains vs losses, previous gains vs guaranteed alternatives, and previous losses vs guaranteed alternatives. Tests were conducted using the “linearHypothesis” function in the package “car” in R⁵⁸, and were chi-square distributed with 1 degree of freedom. A Wald test was also used to test the effects of previous gain outcomes against previous loss outcomes in model 8 in an identical procedure.

Analysis: non-linear prospect theory plus model. We fit a modified version of a 4-parameter prospect theory-inspired model of non-linear processes underlying valuation and choice (Prospect Theory Plus, PT+) ¹⁰. Prospect Theory and related models, like PT+, are a class of nonlinear models that are both commonly used on, and acknowledged to provide a good fit to risky decision-making data^{1,27,28,33,34}. In brief, the PT+ model transformed objective values to subjective values in utility functions, which were then used to compute the likelihood of making a risky choice using a softmax function. The model contains four parameters capturing four distinct decision-making processes in three equations. The subjective values of risky gains, risky losses, and guaranteed alternatives were calculated with Eqs. (1) & (2):

$$u(x^+) = p(x^+)(x^+)^{\rho} \quad (1)$$

$$u(x^-) = p(x^-) \times -\lambda \times (-x^-)^{\rho} \quad (2)$$

In these equations, $p(x)$ is the probability of receiving the value, x . Consider an example choice between a gamble with a risky gain of +\$12 and a risky loss of -\$10, and a guaranteed alternative of \$0. The utility calculations for the risky gain, risky loss, and guaranteed alternative would be as follows for someone with parameters $\rho = 0.8$ and $\lambda = 1.5$:

$$\begin{aligned} u(\text{risky gain}) &= 0.5 \times (+\$12)^{0.8} = +3.65 \\ u(\text{risky loss}) &= 0.5 \times -1.5 \times (-(-\$10))^{0.8} = -4.73 \\ u(\text{guaranteed alternative}) &= 1 \times (\$0)^{0.8} = 0 \end{aligned}$$

The parameter ρ , present in Eqs. (1) and (2), captures diminishing sensitivity to changes in value as the absolute value increases, represented as curvature of the utility function. For choices exclusively in the gain domain, when $\rho < 1$, an individual is risk averse, $\rho = 1$ indicates risk neutrality, while $\rho > 1$ indicates risk seeking. In the loss domain, the effect of ρ on risk-taking is reversed: $\rho < 1$ indicates risk seeking for losses, $\rho = 1$ indicates risk neutrality, and $\rho > 1$ indicates risk aversion for losses. This reversal is often called the reflection effect¹.

The parameter λ , present in Eq. (2), captures loss aversion as a multiplicative weight on losses as compared to gains. When $\lambda = 1$, gains and losses are valued equally (“gain-loss neutral”), while $\lambda > 1$ indicates the

overvaluation of losses compared to gains (“loss averse”), and $\lambda < 1$ indicates the undervaluation of losses compared to gains (“gain seeking”).

The third equation is the logit, or softmax function which integrates the utility of the gamble (i.e. the sum of the utility of the risky gain and risky loss) and the guaranteed alternative, and a value-independent decision bias (db) to calculate the probability (between 0 and 1) of accepting the gamble:

$$p(\text{accept gamble}) = (1 + e^{-\mu \times (u(\text{gamble}) - u(\text{guaranteed}) - db)})^{-1} \quad (3)$$

Equation (3) contains two parameters. First, μ (μ) quantifies the extent to which an individual’s binary choices respond to changes across trials in the value of the gamble as compared to the guaranteed alternative. A μ of 0 indicates complete randomness across choices, whereas increasing values of μ indicate greater consistency across choices. Second, the parameter db represents a value-independent decision bias toward or away from risk-taking, captured as a leftward or rightward shift in the softmax function. $db < 0$ indicates a bias to gamble (softmax function shifts left on x -axis) and $db > 0$ indicates a bias to not gamble (softmax function shifts right on x -axis). For more on decision bias, see reference¹⁰.

The parameters ρ and λ are value-dependent parameters, as they affect the transformation of objective value to subjective value. The parameter db is value-independent as it simply affects a propensity for action (risk-taking) without regard to the values present on a given trial. μ is an intermediate kind of parameter, as it links values to action.

Analysis: hierarchical bayesian model-fitting. To estimate the parameter values of PT+, we used Markov-Chain Monte Carlo (MCMC) sampling techniques in a hierarchical Bayesian framework. This approach explicitly fits individual-level parameters (e.g. participant 1’s decision bias, participant 2’s decision bias, etc.) and group-level parameters (e.g. the group-level mean and standard deviation describing the distribution of decision bias parameter values across individuals), thereby capturing both individual variability and the similarity across individuals to maximal statistical effect while simultaneously reducing the impact of outliers. This kind of approach has been used successfully to model risky decision-making in both our own⁵ and others^{33,34} research.

To model each of the four main parameters of PT+ over time, we modeled their initial value and an update parameter that controlled how much previous decision outcomes shifted that parameter on subsequent trials. To do this, we altered PT+ to include four additional updating parameters (four δ^θ terms, i.e. δ^ρ , δ^λ , δ^μ , δ^{db}). Thus, each of the four main PT+ parameters were modeled for a given participant as follows:

$$\theta_1 = \theta_{\text{initial}} \quad (4)$$

$$\theta_t = \theta_{t-1} + \text{outcome}_{t-1} \times \delta^\theta \quad (5)$$

Equation (4) represents the baseline value for each parameter (λ , ρ , μ and db any of which is represented by θ) on the first trial for a given participant. Equation (5) describes the value of each parameter on a given trial t (θ_t , for $t > 1$) as a function of the previous value of that parameter (θ_{t-1}), the outcome on the previous trial (outcome $_{t-1}$, linearly scaled to be between 0 and 1), and the update parameter (δ^θ).

All participants’ individual-level parameters (each of ρ , λ , μ , db , δ^ρ , δ^λ , δ^μ , and δ^{db}) were normally-distributed around a group-level mean and standard deviation for that parameter. When examining model output, we therefore focused analysis on estimates of the group-level mean parameters (e.g. λ_{mean}).

For each parameter, we calculated 95% confidence intervals (CIs) for the posterior distribution of samples for that parameter. In the case of the update parameters, we examined whether each 95% CI contained zero. If the CI did not contain zero, we could conclude, with 95% confidence, that the true value of that update parameter was not zero, suggesting that previous outcomes had indeed influenced the value of that parameter of PT+ over time.

See Supplementary Materials for more information, including priors for the group-level parameters and detailed explanation of the mechanics of model estimation and the use of mathematical transformations to gently implement parameter bounds. The full Stan model code is also available on the Open Science Framework: <https://osf.io/npd54/>.

Data availability

The complete datasets analyzed in the current study, along with the Stan code for the hierarchical Bayesian model of PT+ are available on the Open Science Framework at <https://osf.io/qxa9h/>.

Received: 5 November 2019; Accepted: 18 May 2020;

Published online: 18 June 2020

References

1. Kahneman, D. & Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–292 (1979).
2. Sokol-Hessner, P. *et al.* Thinking like a trader selectively reduces individuals’ loss aversion. *PNAS* **31**, (2009).
3. Sokol-Hessner, P., Hartley, C. A., Hamilton, J. R. & Phelps, E. A. Interoceptive ability predicts aversion to losses. *Cogn. Emot.* **29**, 695–701 (2015).
4. Sokol-Hessner, P. *et al.* Determinants of Propranolol’s Selective Effect on Loss Aversion. *Psychol. Sci.* **26**, 1123–1130 (2015).
5. Sokol-Hessner, P., Raio, C. M., Gottesman, S. P., Lackovic, S. F. & Phelps, E. A. Acute stress does not affect risky monetary decision-making. *Neurobiol. Stress* **5**, 19–25 (2016).
6. Brevers, D., He, Q., Xue, G. & Bechara, A. Neural correlates of the impact of prior outcomes on subsequent monetary decision-making in frequent poker players. *Biol. Psychol.* **124**, 30–38 (2017).
7. Hytönen, K. *et al.* Path dependence in risky choice: Affective and deliberative processes in brain and behavior. *J. Econ. Behav. Organ.* **107**, 566–581 (2014).

8. Post, T., Van Den Assem, M. J., Baltussen, G. & Thaler, R. H. Deal or no deal? decision making under risk in a large-payoff game show. *Am. Econ. Rev.* **98**, 38–71 (2008).
9. Suhonen, N. & Saastamoinen, J. How Do Prior Gains and Losses Affect Subsequent Risk Taking? New Evidence from Individual-Level Horse Race Bets. *Manage. Sci.* **64**, 2797–2808 (2017).
10. Juechems, K., Balaguer, J., Ruz, M. & Summerfield, C. Ventromedial Prefrontal Cortex Encodes a Latent Estimate of Cumulative Reward. *Neuron* **93**, 705–714.e4 (2017).
11. Thaler, R. H. & Johnson, E. J. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Manage. Sci.* **36**, 643–660 (1990).
12. Imas, A. The realization effect: Risk-taking after realized versus paper losses. *Am. Econ. Rev.* **106**, 2086–2109 (2016).
13. Isen, A. M. & Patrick, R. The Effect of Positive Feelings on Risk Taking: When the Chips Are Down. *Organizational Behavior and Human Performance* **31**, 194–202 (1983).
14. Pabst, S., Brand, M. & Wolf, O. T. Stress effects on framed decisions: There are differences for gains and losses. *Front. Behav. Neurosci.* (2013). <https://doi.org/10.3389/fnbeh.2013.00142>
15. Porcelli, A. J. & Delgado, M. R. Acute stress modulates risk taking in financial decision making. *Psychol. Sci.* (2009). <https://doi.org/10.1111/j.1467-9280.2009.02288.x>
16. Klueen, L. M., Agorastos, A., Wiedemann, K. & Schwabe, L. Cortisol boosts risky decision-making behavior in men but not in women. *Psychoneuroendocrinology* (2017). <https://doi.org/10.1016/j.psyneuen.2017.07.240>
17. Solway, A., Lohrenz, T. & Montague, P. R. Loss Aversion Correlates With the Propensity to Deploy Model-Based Control. *Front. Neurosci.* **13**, (2019).
18. Ortiz-Teran, E., Ortiz, T., Turrero, A. & Lopez-Pascual, J. Neural implications of investment banking experience in decision-making under risk and ambiguity. *J. Neurosci. Psychol. Econ.* (2019). <https://doi.org/10.1037/npe0000100>
19. Sokol-Hessner, P., Camerer, C. F. & Phelps, E. A. Emotion regulation reduces loss aversion and decreases amygdala responses to losses. *Soc. Cogn. Affect. Neurosci.* **8**, 341–350 (2013).
20. Gardner, M. & Steinberg, L. Peer influence on risk taking, risk preference, and risky decision making in adolescence and adulthood: An experimental study. *Dev. Psychol.* (2005). <https://doi.org/10.1037/0012-1649.41.4.625>
21. Chumbley, J. R. et al. Endogenous Cortisol and Loss Aversion. *Psychol. Sci.* **25**, 2102–2105 (2014).
22. Starcke, K., Wolf, O. T., Markowitsch, H. J. & Brand, M. Anticipatory Stress Influences Decision Making Under Explicit Risk Conditions. *Behav. Neurosci.* (2008). <https://doi.org/10.1037/a0013281>
23. Pabst, S., Brand, M. & Wolf, O. T. Stress and decision making: A few minutes make all the difference. *Behav. Brain Res.* (2013). <https://doi.org/10.1016/j.bbr.2013.04.046>
24. Pabst, S., Schoofs, D., Pawlikowski, M., Brand, M. & Wolf, O. T. Paradoxical effects of stress and an executive task on decisions under risk. *Behav. Neurosci.* (2013). <https://doi.org/10.1037/a0032334>
25. Knight, F. H. Risk, uncertainty and profit. (Houghton Mifflin Company, 1921).
26. Bates, D., Maechler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
27. Sokol-Hessner, P. & Rutledge, R. B. The Psychological and Neural Basis of Loss Aversion. *Curr. Dir. Psychol. Sci.* **28**, 20–27 (2019).
28. Camerer, C. Prospect theory in the wild: Evidence from the field. in *Choices, values, and frames* (eds. D., Kahneman & A., Tversky) 288–300 (Cambridge University Press, 2000).
29. Glimcher, P. W. Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proc. Natl. Acad. Sci. USA* **108**, 15647–15654 (2011).
30. Gläscher, J., Daw, N., Dayan, P. & O’Doherty, J. P. States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* **66**, 585–595 (2010).
31. Schultz, W. Dopamine reward prediction-error signalling: A two-component response. *Nature Reviews Neuroscience* **17**, 183–195 (2016).
32. Charpentier, C. J., Aylward, J., Roiser, J. P. & Robinson, O. J. Enhanced Risk Aversion, But Not Loss Aversion, in Unmedicated Pathological Anxiety. *Biol. Psychiatry* (2017). <https://doi.org/10.1016/j.biopsych.2016.12.010>
33. Nilsson, H., Rieskamp, J. & Wagenmakers, E.-J. Hierarchical Bayesian parameter estimation for cumulative prospect theory. *J. Math. Psychol.* **55**, 84–93 (2011).
34. Scheibehenne, B. & Pachur, T. Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychon. Bull. Rev.* **22**, 391–407 (2015).
35. Stan Development Team. RStan: the R interface to Stan. (2018).
36. Stewart, N., Chater, N. & Brown, G. D. A. Decision by sampling. *Cogn. Psychol.* **53**, 1–26 (2006).
37. Walasek, L. & Stewart, N. How to Make Loss Aversion Disappear and Reverse: Tests of the Decision by Sampling Origin of Loss Aversion. *J. Exp. Psychol. Gen.* **144**, 7–11 (2015).
38. Walasek, L. & Stewart, N. Context-dependent sensitivity to losses: Range and skew manipulations. *J. Exp. Psychol. Learn. Mem. Cogn.* **45**, 957–968 (2019).
39. Rigoli, F., Rutledge, R. B., Dayan, P. & Dolan, R. J. The influence of contextual reward statistics on risk preference. *Neuroimage* **128**, 74–84 (2016).
40. Rigoli, F., Friston, K. J. & Dolan, R. J. Neural processes mediating contextual influences on human choice behaviour. *Nat. Commun.* **7**, (2016).
41. Jessup, R. K., Bishara, A. J. & Busemeyer, J. R. Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science* **19**, (2008).
42. Agranov, M. & Ortoleva, P. Stochastic choice and preferences for randomization. *J. Polit. Econ.* **125**, 40–68 (2017).
43. Louie, K., Khaw, M. W. & Glimcher, P. W. Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci. USA* **110**, 6139–44 (2013).
44. Tymula, A. A. & Glimcher, P. W. Expected Subjective Value Theory (ESVT): A Representation of Decision Under Risk and Certainty. *SSRN Electron. J.* (2016). <https://doi.org/10.2139/ssrn.2783638>
45. Khaw, M. W., Glimcher, P. W. & Louie, K. Normalized value coding explains dynamic adaptation in the human valuation process. *Proc. Natl. Acad. Sci.* **114**, 12696–12701 (2017).
46. Bornstein, A. M. & Norman, K. A. Reinstated episodic context guides sampling-based decisions for reward. *Nat. Neurosci.* (2017). <https://doi.org/10.1038/nn.4573>
47. Bornstein, A. M., Khaw, M. W., Shohamy, D. & Daw, N. D. Reminders of past choices bias decisions for reward in humans. *Nat. Commun.* (2017). <https://doi.org/10.1038/ncomms15958>
48. Barron, H. C., Dolan, R. J. & Behrens, T. E. J. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat. Neurosci.* (2013). <https://doi.org/10.1038/nn.3515>
49. Bhui, R. & Gershman, S. J. Decision by sampling implements efficient coding of psychoeconomic functions. *Psychol. Rev.* **125**, 985–1001 (2018).
50. Rigoli, F., Mathys, C., Friston, K. J. & Dolan, R. J. A unifying Bayesian account of contextual effects in value-based choice. *PLoS Comput. Biol.* **13**, 1–28 (2017).
51. Phelps, E. A., Lempert, K. M. & Sokol-Hessner, P. Emotion and Decision Making: Multiple Modulatory Neural Circuits. *Annu. Rev. Neurosci.* **37**, 263–287 (2014).

52. Suter, R., Pachur, T. & Hertwig, R. How affect shapes risky choice: Distorted probability weighting versus probability neglect. *J. Behav. Decis. Mak.* **29**, 437–449 (2016).
53. Eldar, E., Rutledge, R. B., Dolan, R. J. & Niv, Y. Mood as Representation of Momentum. *Trends in Cognitive Sciences* (2016). <https://doi.org/10.1016/j.tics.2015.07.010>
54. Loewenstein, G. F., Hsee, C. K., Weber, E. U. & Welch, N. Risk as Feelings. *Psychol. Bull.* (2001). <https://doi.org/10.1037/0033-2909.127.2.267>
55. Lerner, J. S., Li, Y., Valdesolo, P. & Kassam, K. S. Emotion and Decision Making. *Annu. Rev. Psychol.* (2015). <https://doi.org/10.1146/annurev-psych-010213-115043>
56. Clithero, J. A. & Rangel, A. Informatic parcellation of the network involved in the computation of subjective value. *Soc. Cogn. Affect. Neurosci.* (2013). <https://doi.org/10.1093/scan/nst106>
57. Bartra, O., McGuire, J. T. & Kable, J. W. The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage* (2013). <https://doi.org/10.1016/j.neuroimage.2013.02.063>
58. Fox, J. & Weisberg, S. An {R} Companion to Applied Regression. (Sage, 2019).

Acknowledgements

The authors are grateful to Elizabeth Phelps, Paul Glimcher, and Agnieszka Tymula for thoughtful conversations.

Author contributions

P.S.H. and H.R.B. analyzed the data and wrote the main manuscript. H.R.B. prepared Figs. 1–3. Both authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-66502-y>.

Correspondence and requests for materials should be addressed to P.S.-H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020

Quantifying the immediate computational effects of preceding outcomes on subsequent risky choices

Hayley R. Brooks, Peter Sokol-Hessner

SUPPLEMENTARY MATERIAL

Study Details

The four studies included in this analysis had the same basic instructions and task structure including the initial endowment, the timing of decision and feedback events on each trial, the possible monetary amounts, and the task-related payment. The studies varied by participation fee (\$15-\$20), number of trials (140-180), number of days (1-2), and external manipulation. For this analysis, we excluded choices made under any experimental manipulation and on day 2 to capture participants' first experience with the risky monetary decision-making task. In Study 1, 30 participants switched between two cognitive reappraisal strategies ("regulate" and "attend") during the risky monetary choice task¹. In the "attend", or baseline condition, participants were instructed to think about each choice independent of the other choices. In the "regulate" condition, participants were instructed to think about their choices as a portfolio ("you win some, you lose some"). Study 2 (N=37) involved no experimental manipulations during the risky choice task². Study 3 was a two-day, double-blind, placebo-controlled, within-subjects study in which 47 participants were given the beta-blocker propranolol or a placebo prior to the risky monetary choice task³. In Study 4, a two-day, within-subjects study, 120 participants completed a cold-pressor task or a control manipulation (a warm-water bath) with equal probability prior to the risky monetary decision-making task on each day⁴. In total, there was 64,953 choices across 234 participants. For this analysis, we excluded the choices made in the "regulate" condition (4,200 trials from Study 1), in the propranolol condition (3,432 trials from Study 3), in the cold-pressor condition (8,968 trials from Study 4), and on day 2 (an additional 7,013 trials from Study 3 and 17,967 additional trials from Study 4), leaving a total of 23,373 trials across 151 participants. See Table S1 for summaries of the demographic and methodological differences across the four studies.

Additional MCMC Estimation Details

The Stan model code is available on the Open Science Framework: <https://osf.io/npd54/>

Priors

Sampling priors were selected to be uninformative as possible and were normal (mean, standard deviation), uniform (lower limit, upper limit), or cauchy (location, scale) distributions (see Table S1 below). Parameters were sampled in a different space than they were applied. A transformation was used to convert sampled values to applied values (see section below, *Transformation*)

parameter	Group Mean prior	Group SD prior
λ	Normal(0,30)	Cauchy(0,2.5)
ρ	Normal(0,30)	Cauchy(0,2.5)
μ	Uniform(0,30)	Cauchy (0,2.5)
db	Normal(0,30)	Cauchy(0,2.5)
δ^λ	Normal(0,10)	Cauchy(0,2.5)
δ^ρ	Normal(0,10)	Cauchy(0,2.5)
δ^μ	Normal(0,10)	Cauchy(0,2.5)
δ^{db}	Normal(0,10)	Cauchy(0,2.5)

Table S1: Priors for the group-level mean and standard deviation for each of the main parameters in the prospect theory plus model.

Transformations

For stability of estimation, the MCMC model transformed the sampled baseline parameters λ , ρ , and μ before applying them to the data in manner identical to that used previously⁴ and similar to approaches used by others^{5,6}, in effect implementing a lognormal structure for those parameters. These transformations served to gracefully implement the minimum number of practical bounds on parameter values, without which models would experience numerical faults (overflow; impossible values) that would prevent successful estimation. See *Transformation Rationale* below for more discussion. In all cases, the values discussed in the text and in plots reflect the applied values of these parameters (i.e. after the transformation).

First baseline values for the PT+ parameters λ , ρ , and μ , all of which have theoretical lower bounds of 0, were sampled in unbounded ‘sampling’ space (bounds of [-infinity, +infinity]). To transform from sampling space to ‘applied’ space, these unbounded values were passed through an exponential (e.g. if the sampled value was R, $\exp(R)$ gave the applied value, ρ). An exponential transform produces strictly positive values of ρ for all real values of R (that is, ρ is bounded [0, +infinity]), thereby meeting the basic requirement that values of λ , ρ , and μ , be above 0. All effect sizes and plots reflect the applied (transformed, bounded) values of the parameters (that is, ρ not R).

The exponential transform was not applied to the decision bias parameter, which is theoretically unbounded.

Second, the update parameters (i.e. the δ^θ parameters), were transformed using individual softmax-based functions to gracefully constrain parameters between lower and upper bounds

symmetric around zero, while allowing sampling to occur smoothly in unbounded space. The softmax equations were built so adjustment terms in ‘applied’ space thus had lower/upper bounds of $[-1, +1]$ (db and λ), and $[-0.25, +0.25]$ (ρ and μ). All effect sizes reflect values in these applied (transformed, bounded) spaces.

In all cases, the final 95% confidence intervals for all parameters did not approach their respective bounds, suggesting that these bounds, while effective in enabling model estimation, did not interfere with identification of the most likely values of these parameters. While one cannot definitively state that there are no values of the parameters outside these bounds more likely than those we sampled, it is also not possible to identify them, as model estimation is not stable without these reasonable, psychologically plausible, and commonly-used bounds. The model was coded in Stan (see Methods) – for complete model code, including all aspects of parameter transformations, see <https://osf.io/npd54/>.

Transformation rationale

We have previously published this hierarchical Bayesian implementation of prospect theory⁴, which is structurally similar to that used by others^{5,6}. In essence, in this approach the ‘sampled’ space of the parameters for rho, lambda, and mu are unbounded (i.e. with bounds of $[-\infty, +\infty]$), but the final parameter values applied to data (what we call “applied space”; after the use of the exponential, see above) are bounded $[0, +\infty]$. We use the exponential to implement a lower bound of 0 for three of our four parameters (rho, lambda, and mu), for two main reasons:

1) It smoothly implements a lower bound of 0 (which is required for rho, lambda, and mu), while leaving the parameters unconstrained in the positive direction, as they have no theoretical upper bound (even if their plausible, expected, and psychologically-likely values are generally lower).

2) It gracefully allows the summation of independent terms contributing to the value of a parameter on a given trial in unbounded, ‘sampled’ space (e.g. the effect of previous outcomes on the loss aversion parameter) while preventing those summations from under-flowing.

Interaction between the priors and transformations

As a result of the prior distributions going through an exponential to generate softly-bounded final parameters, the priors favored lower values for each of the parameters. This approach was deemed reasonable for three reasons. First, lower values of these parameters are indeed psychologically most likely (e.g. rho values tend to be reported between .5 and 1.5; lambda values between 0.5 and 4). Second, the priors, after being put through the exponential transformation are still broad and relatively uninformative. Lastly, our MCMC sampling procedure discarded the first 5,000 samples (50%) of each of the twenty chains, effectively eliminating the influence of the selected priors on the final sampled posterior distributions.

1. Sokol-Hessner, P. *et al.* *Thinking like a trader selectively reduces individuals' loss aversion.* *PNAS* **31**, (2009).
2. Sokol-Hessner, P., Hartley, C. A., Hamilton, J. R. & Phelps, E. A. Interoceptive ability predicts aversion to losses. *Cogn. Emot.* **29**, 695–701 (2015).
3. Sokol-Hessner, P. *et al.* Determinants of Propranolol's Selective Effect on Loss Aversion. *Psychol. Sci.* **26**, 1123–1130 (2015).
4. Sokol-Hessner, P., Raio, C. M., Gottesman, S. P., Lackovic, S. F. & Phelps, E. A. Acute stress does not affect risky monetary decision-making. *Neurobiol. Stress* **5**, 19–25 (2016).
5. Nilsson, H., Rieskamp, J. & Wagenmakers, E.-J. Hierarchical Bayesian parameter estimation for cumulative prospect theory. *J. Math. Psychol.* **55**, 84–93 (2011).
6. Scheibehenne, B. & Pachur, T. Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychon. Bull. Rev.* **22**, 391–407 (2015).

Table S2. Characteristics of the previously published studies in the current analysis.

Study #	1	2	3	4
Title	Thinking like a trader selectively reduces individuals' loss aversion.	Interoceptive ability predicts aversion to losses.	Determinants of propranolol's selective effect on loss aversion.	Acute stress does not affect risky monetary decision-making.
Authors	Sokol-Hessner, P., Hsu, M., Curley, N.G., Delgado, M.R., Camerer, C.F., Phelps, E.A.	Sokol-Hessner, P., Hartley, C.A., Hamilton, J.R., Phelps, E.A.	Sokol-Hessner, P., Lackovic, S.F., Tobe, R.H., Camerer, C.F., Leventhal, B.L., Phelps, E.A.	Sokol-Hessner, P., Raio, C.M., Gottesman, S.P., Lackovic, S.F., Phelps, E.A.
Year	2009	2015a	2015b	2016
Journal	PNAS	Cognition and Emotion	Psychological Science	Neurobiology of Stress
N	30	37	47	120
Sex	11 females; 19 males	25 females; 12 males	22 females; 25 males	64 females; 56 males
Age, mean(sd)	22(3) years	24.32(4.6) years	26.6(5.1) years	22.4(4.5) years
Trials per participant	140 (120 gain-loss; 20 gain-only)	180 (150 gain-loss; 30 gain-only)	150 (120 gain-loss; 30 gain-only)	
Days	1		2	
Incentive structure	Subject fee	\$15		\$20/hr
	Endowment	\$30		
	Payout	outcomes on 10% of randomly selected trials		
Timing of task events	4s viewing window, 2s response period, 1s ISI, 1s outcome, 1-3s ITI	2s viewing window, 2s response window, 1s ISI, 1s outcome, 1-3s ITI		
Experimental methodology	One-day, within-subjects study. Participants completed a risky monetary choice task while switching between two cognitive reappraisal (attend vs. regulate) strategies.	One-day, within-subjects study. Participants completed a heartbeat detection task (as a measure interoception) followed by the risky monetary gambling task.	Double-blind, placebo-controlled, within-subjects study. On each day, participants were administered either propranolol or a placebo prior to completing risky monetary choice task.	Fully-crossed, two-day, within-subjects study. Participants completed a cold pressor task or a control manipulation with equal probability prior to making a series of risky monetary choices. Acute stress was measured by salivary cortisol samples.
Exclusion criteria for this analysis	Choices made during "reappraisal" condition.	NA	Choices made during the administration of propranolol and on day 2.	Choices made during the stress condition and on day 2.
Missed trials	0	19	32	42

Table S3. Generalized linear modeling results using the "lme4" package (R version 3.5.0; "lme4" version 1.1-21).

Model 1 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + outcome amount(t-1) + (1 Subject ID), data, family = "binomial")				
Model 1 results				
AIC	BIC	logLik	deviance	df.resid
21289	21337.3	-10638.5	21277	23216
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.067457	0.104245	-0.647	0.518
risky gain amount(t)	0.299218	0.007126	41.99	< 2e-16 ***
risky loss amount (t)	0.397172	0.006532	60.804	< 2e-16 ***
safe amount (t)	-0.590955	0.014532	-40.666	< 2e-16 ***
outcome amount(t-1)	-0.030859	0.003092	-9.979	< 2e-16 ***
Model 2 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + choice(t-1) + (1 Subject ID), data, family = "binomial")				
Model 2 results				
AIC	BIC	logLik	deviance	df.resid
21377.5	21425.8	-10682.7	21365.5	23216
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.134037	0.105756	-1.267	0.205003
risky gain amount(t)	0.299066	0.007109	42.066	< 2e-16 ***
risky loss amount (t)	0.396355	0.006520	60.795	< 2e-16 ***
safe amount (t)	-0.590458	0.014502	-40.717	< 2e-16 ***
choice(t-1)	-0.068814	0.018671	-3.686	0.000228 ***
Model 3 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + mean EV(t-1) + (1 Subject ID), data, family = "binomial")				

Model 3 results				
AIC	BIC	logLik	deviance	df.resid
21350.4	21398.7	-10669.2	21338.4	23216
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.088786	0.103914	-0.854	0.393
risky gain amount(t)	0.299548	0.007121	42.066	< 2e-16 ***
risky loss amount (t)	0.39585	0.006517	60.744	< 2e-16 ***
safe amount (t)	-0.59061	0.014521	-40.674	< 2e-16 ***
mean EV (t-1)	-0.032006	0.00505	-6.337	2.34e-10 ***

Model 4 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + mean EV(t-1) + choice(t-1) + outcome amount(t-1) + (1|Subject ID), data, family = "binomial")

Model 4 results				
AIC	BIC	logLik	deviance	df.resid
21288.7	21353.2	-10636.4	21272.7	23214
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.078854	0.105503	-0.747	0.4548
risky gain amount(t)	0.299485	0.00713	42.004	< 2e-16 ***
risky loss amount (t)	0.397197	0.006534	60.787	< 2e-16 ***
safe amount (t)	-0.591271	0.014536	-40.676	< 2e-16 ***
mean EV(t-1)	-0.001755	0.006312	-0.278	0.7809
choice(t-1)	-0.037183	0.019637	-1.893	0.0583 .
outcome amount(t-1)	-0.029276	0.003728	-7.853	4.08e-15 ***

Model 5 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + outcome amount(t-1) + outcome amount(t-2) + outcome amount(t-3) + (1|Subject ID), data, family = "binomial")

Model 5 results				
AIC	BIC	logLik	deviance	df.resid
20941.7	21006	-10462.8	20925.7	22912

Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.066689	0.105297	-0.633	0.5265
risky gain amount(t)	0.300642	0.00716	41.961	< 2e-16 ***
risky loss amount (t)	0.397693	0.006589	60.361	< 2e-16 ***
safe amount (t)	-0.592879	0.014614	-40.569	< 2e-16 ***
outcome amount(t-1)	-0.031339	0.003114	-10.064	< 2e-16 ***
outcome amount(t-2)	-0.007769	0.003050	-2.547	0.0109 *
outcome amount(t-3)	-0.003964	0.003057	-1.296	0.1948

Model 6 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + outcome amount (t-1) + outcome valence(t-1) + (1|Subject ID), data, family = "binomial")

Model 6 results				
AIC	BIC	logLik	deviance	df.resid
21287.8	21344.2	-10636.9	21273.8	23215

Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.066248	0.104291	-0.635	0.525
risky gain amount(t)	0.299004	0.007126	41.96	< 2e-16 ***
risky loss amount (t)	0.397148	0.006533	60.795	< 2e-16 ***
safe amount (t)	-0.590608	0.014532	-40.642	< 2e-16 ***
outcome amount (t-1)	-0.037835	0.005038	-7.509	5.95e-14 ***
outcome valence (t-1)	0.079089	0.044875	1.762	0.078 .

Model 7 = glmer(choice ~ 1 + risky gain amount(t) + risky loss amount(t) + safe amount(t) + win amount(t-1) + loss amount(t-1) + safe amount(t-1) + (1|Subject ID), data, family = "binomial")

Model 7 results				
AIC	BIC	logLik	deviance	df.resid

21291.8	21356.2	-10637.9	21275.8	23214
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.064496	0.104951	-0.615	0.538861
risky gain amount(t)	0.299152	0.007126	41.979	< 2e-16 ***
risky loss amount (t)	0.397241	0.006533	60.802	< 2e-16 ***
safe amount (t)	-0.590842	0.014532	-40.659	< 2e-16 ***
win amount(t-1)	-0.033083	0.003973	-8.327	< 2e-16 ***
loss amount(t-1)	-0.02797	0.007819	-3.577	0.000348 ***
safe amount(t-1)	-0.023994	0.008558	-2.804	0.005053 **
Model 8 = glmer(choice ~ 1 + risky gain amount(t)+ risky loss amount(t) + safe amount(t) + gain amount(t-1)*choice(t-1) + loss amount(t-1) + (1 Subject ID), data, family = "binomial")				
Model 8 results				
AIC	BIC	logLik	deviance	df.resid
21290.2	21362.7	-10636.1	21272.2	23213
Fixed effects				
	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.0899602	0.106447	-0.845	0.398
risky gain amount(t)	0.2994444	0.00713	42	< 2e-16 ***
risky loss amount (t)	0.3972658	0.006533	60.808	< 2e-16 ***
safe amount (t)	-0.5912799	0.014536	-40.676	< 2e-16 ***
gain amount(t-1)	-0.0272841	0.0049348	-5.529	3.22e-08 ***
choice(t-1)	-0.0458073	0.024232	-1.89	0.0587 .
loss amount(t-1)	-0.0357708	0.008868	-4.034	5.49e-05 ***
gain amount(t-1)*choice(t-1)	-0.0008651	0.004943	-0.175	0.8611

Model 9 = glmer(choice ~ 1 + risky gain amount(t)*outcome amount(t-1) + risky loss amount(t)*outcome amount(t-1) + safe amount(t)*outcome amount(t-1) + (1|Subject ID), data, family = "binomial")

Model 8 results

AIC	BIC	logLik	deviance	df.resid
21278.6	21351.1	-10630.3	21260.6	23213

Fixed effects

	Estimate	Std Error	Z value	Pr(> z)
intercept	-0.08573	0.1045	-0.82	0.4120
risky gain amount(t)	.299348	0.007348	40.738	< 2e-16 ***
risky loss amount (t)	0.394433	0.0066	59.761	< 2e-16 ***
safe amount (t)	-0.590369	0.014989	-39.386	< 2e-16 ***
outcome amount(t-1)	-0.45562	0.205493	-2.217	0.02660 *
risky gain amount(t)* outcome amount(t-1)	0.012088	0.04071	0.297	0.76653
risky loss amount(t)* outcome amount(t-1)	0.1019	0.03271	3.128	0.00176**
safe amount(t)* outcome amount(t-1)	-0.056176	0.082578	-0.68	0.49633