


Memory and decision making interact to shape the value of unchosen options

Natalie Biderman ¹✉ & Daphna Shohamy^{1,2}✉

The goal of deliberation is to separate between options so that we can commit to one and leave the other behind. However, deliberation can, paradoxically, also form an association in memory between the chosen and unchosen options. Here, we consider this possibility and examine its consequences for how outcomes affect not only the value of the options we chose, but also, by association, the value of options we did not choose. In five experiments (total $n = 612$), including a preregistered experiment ($n = 235$), we found that the value assigned to unchosen options is inversely related to their chosen counterparts. Moreover, this inverse relationship was associated with participants' memory of the pairs they chose between. Our findings suggest that deciding between options does not end the competition between them. Deliberation binds choice options together in memory such that the learned value of one can affect the inferred value of the other.

¹Department of Psychology and Mortimer B. Zuckerman Mind, Brain, Behavior Institute, Columbia University, New York, NY, USA. ²The Kavli Institute for Brain Science, Columbia University, New York, NY, USA. ✉email: natalie.biderman@columbia.edu; ds2619@columbia.edu

The Latin origin of the verb ‘to decide’ literally means ‘to cut off’ (‘de’ = off, ‘caedere’ = cut). The act of deciding is supposed to cut off the deliberation process to settle on one’s choice. Yet, often the deliberation does not seem to end when a decision is made. Instead, the unchosen option continues to linger in one’s mind. Over the years, research has shown that people continue to think about counterfactuals^{1–6}. Yet, little is known about the mechanism that allows unchosen options to linger in our minds and the consequences of such lingering for the value of unchosen options.

Here we aim to address this gap. The starting point of our inquiry is that thoughts about unchosen alternatives do not come from thin air, but that they are instead tied to the chosen option^{7–9}. For example, consider a decision you had to make, say about where to go on vacation. If the location you ended up choosing did not meet your expectations, this would of course lead you to devalue the choice you made¹⁰. But often in such situations, we also find ourselves automatically thinking back to the unchosen option, which we may now evaluate as a better option than it seemed at the time. Such post-decision experiences suggest that the options we deliberate between remain linked in our minds long after the decision was made. Indeed, deliberation is a comparative process, wherein options are evaluated simultaneously in relation to each other^{11–13}. In this sense, deliberation provides a temporal and conceptual context shared by choice alternatives. Extensive memory research has shown that shared context creates an association between disparate elements, binding them to each other in memory^{14–17}. These findings suggest that rather than ‘cutting off’ the unchosen option, deliberation may, paradoxically, tie the options together.

The possibility of choice options remaining associated in memory could have substantial consequences for how learning shapes value once the outcomes of a choice are revealed. Many studies have investigated the mechanisms by which outcomes of one’s choice drive learning about the chosen option. According to prominent reinforcement learning models, learning the outcomes of chosen options leads to their value being updated based on the difference between the expected value and the actual experienced value¹⁰.

The mechanisms of learning about unchosen options, however, are less clear. Studies of counterfactual learning have focused on how people learn from explicit information about what they could have gained if they had chosen the other option^{18–24}. Yet, in most cases, people are not exposed to the outcomes of their forgone alternatives, leaving open questions about if, and what, people learn about unchosen options afterwards, when they experience outcomes of the choice that they made. In these situations, is the value of the unchosen options updated as well?

Here we hypothesize that the value of unchosen options is updated through their association with the chosen options. Central to this hypothesis is the role of memory. Recent advances suggest that memory associations can facilitate value inference and generalization. The basic idea is that once an association is formed, encountering one item leads to reactivation of the associated items^{25–27}. In a recent study rats were required to make navigation decisions, and when they learned they made the wrong choice and their expected reward was nowhere to be seen, they reactivated the unchosen location⁹. Moreover, studies in humans have shown that if a reward is given to one of two associated items, the value of that reward can spread towards the associated items by reactivating a mnemonic network^{28–30}. This reactivation mechanism has been shown to account for updating of chosen options, however, it remains unknown if it also affects the value of unchosen options.

The current study sought to determine whether the act of deliberation creates a memory association between the

deliberated options and to explore the consequences of this association for later value learning. We hypothesized that learning about the outcomes of the choice leads to reactivation, in memory, of the unchosen option and this in turn leads to a change in the unchosen option’s value. Unlike previous studies showing direct value transfer among associated items^{28–30}, choice options are associated within the deliberation context, which involves a contrast between the options. We, therefore, expect the value to transfer in the opposite direction. That is, if chosen items are explicitly rewarded (or unrewarded), we expect unchosen items to be inferred as unrewarded (or rewarded), a behavioral tendency we refer to as inverse inference of value. Thus, we hypothesize that deliberation may have a somewhat paradoxical role: While it is meant to dissociate choice options, deliberation binds them in our memory. When this bond is reactivated for the purpose of value updating, it continues to serve the deliberation goal of teasing the value of options apart.

Our prediction of an inverse transfer of value between choice options is based upon previous behavioral findings showing that participants continue to separate the value of options even after the decision was terminated. In studies of choice-supportive memory, participants exhibited a bias to better learn and remember the positive aspects of chosen options and the negative aspects of unchosen ones^{31–34}, thus increasing the contrast of value between the options. A similar contrast appears in studies of choice-induced preference change. These studies observed that after participants make a choice—even without any feedback on their choice—the value of chosen options tends to increase, while the value of the unchosen option tends to decrease^{35–38}. Importantly, our study differs from, and therefore extends, choice-induced preference change studies in a crucial aspect. While choice-induced preference change studies examine how preferences are altered as a function of choice itself (regardless of the outcomes of such choices), the current study examines how the inferred value of unchosen options is updated as a function of outcomes.

To test our hypotheses, we devised a multiple-phase behavioral experiment in which participants deliberated between options and then learned the value of their choices (Fig. 1). First, we asked participants to deliberate between pairs of paintings and decide which painting would be more profitable in an upcoming auction (Phase 1, Deliberation; chosen items denoted as S_{chosen} and unchosen items as S_{unchosen}). Importantly, we neither instructed participants to memorize the paintings nor did we specify that there was any dependency between the outcomes of paintings in each pair. Because we were interested in the effects of deliberation and memory on the valuation of unchosen options, we sought to verify that deliberation took place by asking participants to write down the reasons for choosing one option and not the other. Similarly, we sought to verify memory would be robust enough by repeating the deliberation trials several times and giving participants the chance to practice their decisions before committing to their choice.

In the second phase (Phase 2, Outcome Learning), participants learned about the outcome of each of their choices. Because we wanted to assess the effects of associative memory on value updating, we did not provide immediate feedback after each decision. Instead, after all decisions were made we presented only the chosen paintings alongside their auction outcomes.

In the next phase (Phase 3, Final Decisions) we sought to measure whether the outcome learning led to any value updating for either the chosen or unchosen options. We asked participants to make a new series of decisions between pairs of paintings, choosing the most-valuable painting in each pair. Each trial in this phase presented either two previously chosen stimuli (rewarded and unrewarded in the auction, denoted as

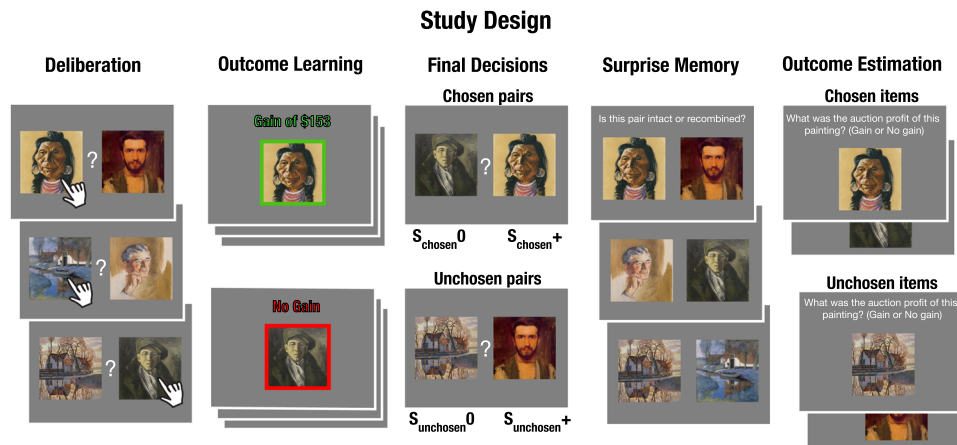


Fig. 1 Study design. In this multi-phase experiment, participants act as art dealers choosing paintings that are later sold in an auction. In the Deliberation phase, participants deliberate and decide which painting they would like to purchase to maximize their profits from the auction. After all decisions are made they learn the auction outcomes in the Outcome Learning phase. Only their selected paintings are presented alongside a colored frame and the profit they made in the auction: either gain (green frame) or no gain (red frame). Then, in the Final Decisions phase, participants are given a choice between two paintings and are asked to choose the most valuable. Unbeknownst to them, the decision pairs are either two previously chosen paintings (rewarded and unrewarded in the auction, denoted as $S_{\text{chosen}+}$ and $S_{\text{chosen}0}$, respectively) or two previously unchosen paintings (initially presented with $S_{\text{chosen}+}$ and $S_{\text{chosen}0}$, denoted as $S_{\text{unchosen}+}$ and $S_{\text{unchosen}0}$, respectively). The two pair types are randomly intermixed. In the Surprise Memory phase, participants are tested for associative memory of the deliberation pairs. They are presented with either the exact pairs from the Deliberation phase (intact) or pairs that include chosen and unchosen paintings that were not previously presented together (recombined). Experiment 1 included an Outcome Estimation phase, wherein participants are asked to estimate the auction profit of all paintings, including those they did not choose. Chosen and unchosen paintings are randomly intermixed. The study additionally included a pre-task and a post-task liking ratings of the paintings, detailed in the “Method” section. Stimuli were art images depicting people, objects, and scenes, randomly intermixed across conditions⁷⁷. A subset of the stimuli are presented here. The art images in the figure include detailed images of the artworks PH-672 (1923), PP-241 (1936), and PH-269 (1941) by Clyfford Still, courtesy the Clyfford Still Museum, Denver, CO © 2021 City and County of Denver / ARS, NY, and the artworks Self Portrait (1900), A Farmbuilding (1900-1901), Farm Near Duivendrecht (1916) by Piet Mondrian, courtesy of Mondrian/Holtzman Trust © 2021. See Supplementary Fig. 1 for the full images.

$S_{\text{chosen}+}$ and $S_{\text{chosen}0}$, respectively), or two previously unchosen stimuli (presented earlier with $S_{\text{chosen}+}$ and $S_{\text{chosen}0}$, denoted as $S_{\text{unchosen}+}$ and $S_{\text{unchosen}0}$, respectively, Fig. 1). Participants were incentivized to choose the more valuable painting in order to earn extra bonus money based upon their performance.

To assess how memory of the deliberation impacts choices, after the last decision phase, participants were presented with a Surprise Memory Test (Phase 4) for the deliberation pairs. In Experiment 1, to assess explicit value inference of unchosen stimuli, at the end of the experiment we told participants that all paintings went on auction, including their previously unchosen paintings, and asked them to estimate which paintings were rewarded in the auction (Last Phase, Outcome Estimation). In addition, to control for participants’ intrinsic preferences for any specific painting, before starting the main experiment we asked participants to rate each painting individually so that we could select items that were relatively neutral in their subjective value.

We were primarily interested in behavior on the Final Decisions phase and its relation to memory. Specifically, it is expected that if participants successfully learned the new values of the chosen items, then they should select $S_{\text{chosen}+}$ over $S_{\text{chosen}0}$. However, the critical question relates to their behavior on trials with unchosen items, for which they never received any direct feedback. As mentioned previously, we predicted an inverse inference of value for unchosen items. In the Final Decisions phase, where participants are asked to select the most-profitable items, this would result in the tendency to select $S_{\text{unchosen}0}$ over $S_{\text{unchosen}+}$. Finally, we hypothesized that this inverse decision bias should be larger when participants better remember which pairs of options appeared together.

Our main experiment (Experiment 1) was pre-registered on Open Science Framework (<https://osf.io/chsvw>) where we elaborated these hypotheses prior to data collection. To preview our

results, across five distinct data sets (total $n = 612$), we found that participants were biased to select unchosen paintings previously paired with unrewarded paintings. This inverse decision bias was strongly related to their memory for the association between the options during deliberation.

Results

All results were analyzed with Bayesian generalized linear models (see details in the “Analysis” section). In each model, we estimated a posterior distribution for regression coefficients and reported the median and 95% highest density interval (HDI) for coefficients of interest. If the 95% HDI of a coefficient excluded zero we concluded that the corresponding variable reliably predicted the outcome. Here, we refer to the value of chosen options after outcome learning as “learned value”, since it results from participants’ experience with the outcomes of their choices. In contrast, for the unchosen options that received no explicit feedback for learning, we use the term “inferred value”.

The value of chosen and unchosen options is inversely related.

We hypothesized that feedback about the chosen option would modulate the value of its associated unchosen option. To examine this, we focused on performance in the Final Decisions phase, allowing us to determine both direct outcome learning for chosen options (which were reinforced in the Outcome Learning phase) as well as inferences about the value of the unchosen options (which were not reinforced at any point in the experiment). We additionally examined an explicit outcome estimation report that participants gave for each painting at the end of the experiment.

Final Decisions. In Experiment 1 ($n = 235$, see preregistration on Open Science Framework: <https://osf.io/chsvw>), we found that when participants had to decide between two previously chosen items, they preferred $S_{\text{chosen}+}$ over $S_{\text{chosen}0}$ (probability to select

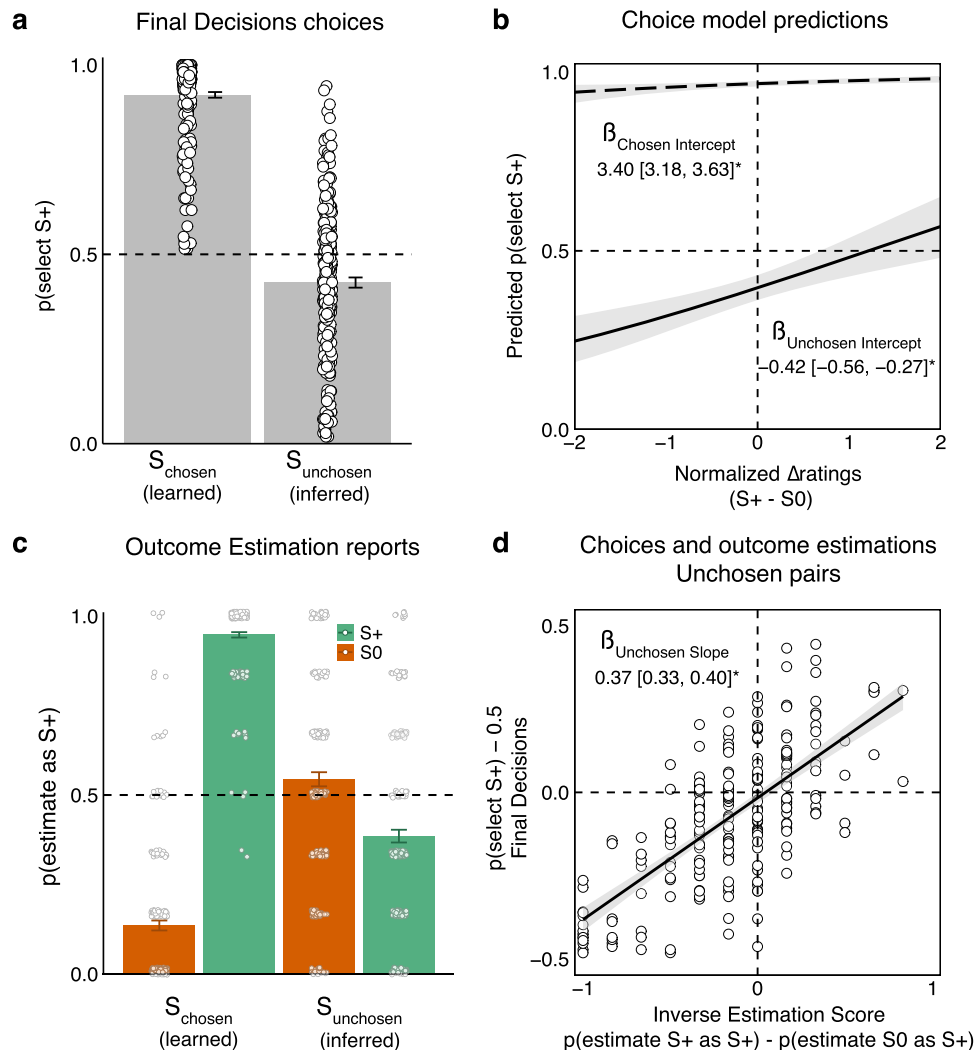


Fig. 2 The inferred value of unchosen options is inversely related to the learned value of chosen options (Experiment 1, $n = 235$). **a** When faced with a choice between two previously chosen stimuli, for which value was explicitly shown [S_{chosen} (learned)], participants tended to select the rewarded option ($S_{\text{chosen}+}$), suggesting they successfully learned their values, while for pairs of previously unchosen options, which were never directly associated with any reward [S_{unchosen} (inferred)], participants tended to select the option previously associated with an unrewarded item ($S_{\text{unchosen}0}$), demonstrating an inverse decision bias. **b** This inverse decision bias was observed even when controlling for initial subjective valuations of the choice options, in a Bayesian logistic regression predicting the probability to choose a rewarded item as a function of pair type and the difference in liking ratings. After rearranging the model coefficients, we can derive separate intercept terms for chosen and unchosen pairs. The intercept coefficient denotes the tendency to choose a rewarded item when there is no difference in liking ratings between the two choice options. For chosen pairs, the intercept is reliably positive, whereas for unchosen pairs it is reliably negative. **c** The inverse inference of value extends beyond the decision phase to explicit estimation of value. When asked to estimate the auction outcomes of each painting, participants correctly remembered the outcomes of the chosen paintings, yet showed the opposite pattern for the unchosen ones. **d** The tendency to select unchosen items previously paired with unrewarded items ($S_{\text{unchosen}0}$ over $S_{\text{unchosen}+}$) in the Final Decisions phase was correlated with an inverse estimation of value, i.e., the tendency to estimate $S_{\text{unchosen}0}$ as rewarded and $S_{\text{unchosen}+}$ as unrewarded in the Outcome Estimation phase. This relationship was assessed in a Bayesian linear regression predicting the mean probability to select rewarded items as a function of inverse estimation of value for chosen and unchosen pairs separately. In panels (**a**) and (**c**), error bars denote the standard error of the mean and points denote trial-averaged data of individual participants. In panels (**b**) and (**d**), the beta coefficients and model fits denote median and 95% highest density interval of the posterior distribution. In panel (**c**), green bars depict rewarded stimuli (S+) and orange bars depict unrewarded stimuli (S0) (for unchosen stimuli, this is the outcome of their chosen counterpart). Source data are provided as a Source Data file.

$S+$: 0.92 ± 0.01 [mean \pm standard error], Fig. 2a), suggesting that they correctly learned the outcomes of their chosen items.

Critically, the opposite pattern was observed for unchosen items: participants were biased to select $S_{\text{unchosen}0}$ over $S_{\text{unchosen}+}$ (probability to select $S+$: 0.43 ± 0.01 , Fig. 2a; for unchosen pairs, $S+$ items are paintings previously paired with a rewarded painting, and $S0$ items are paintings previously paired with an unrewarded painting). Recall that the participants never received explicit feedback on the unchosen items. We quantified these

opposing patterns with a multilevel Bayesian logistic regression model that predicts the probability of selecting $S+$ items in the Final Decisions phase, conditional on the item being chosen or unchosen, and while accounting for the initial subjective likings of each painting. We computed a separate intercept term for chosen and unchosen items that reflects the tendency to select an $S+$ item when there is no liking difference between the choice items (see full model specification in the “Analysis” section, and coefficients in Supplementary Tables 2 and 3). A positive

intercept denotes a tendency to select S+ items and a negative intercept denotes a tendency to select S0 items. As expected, the estimated intercept for chosen pairs was positive ($\beta = 3.40$ [3.18, 3.63]), while for unchosen pairs, it was negative ($\beta = -0.42$ [-0.56, -0.27]; Fig. 2b).

Interestingly, we also observed that the tendency to choose unchosen paintings previously paired with unrewarded paintings ($S_{\text{unchosen}0}$) was accompanied with faster responses compared to choices of paintings previously paired with rewarded paintings ($S_{\text{unchosen}+}$, see Supplementary Text 1 and Supplementary Fig. 2b). Considering the research showing that choices for rewards are faster than choices to avoid loss^{39,40}, this pattern in reaction times mirrors the inverse decision bias and suggests that participants viewed $S_{\text{unchosen}0}$ items to be more valuable.

Together, these results demonstrate that there is an inverse inference of value: in chosen pairs, participants tend to select S+ items, but in unchosen pairs they tend to select S0 items and this tendency is accompanied with faster reaction times.

Outcome Estimation. We next asked whether the aforementioned inverse decision bias extends to explicit change in value of choice options. To this end, we analyzed the explicit reports in the Outcome Estimation phase. Indeed, participants' explicit reports revealed that the estimated value of unchosen items mirrors the inverse decision bias (Fig. 2c). Participants estimated the auction outcomes of unchosen paintings to be inversely related to the chosen ones (probability to estimate an item as rewarded in $S_{\text{unchosen}+}$: 0.39 ± 0.02 ; and in $S_{\text{unchosen}0}$: 0.54 ± 0.02 ; multilevel Bayesian logistic regression: $\beta_{\text{unchosen}+} = -0.57$ [-0.75, -0.39], $\beta_{\text{unchosen}0} = 0.24$ [0.03, 0.45]). Here too, we validated that participants correctly learned which of their chosen paintings resulted in reward (probability to estimate an item as rewarded in $S_{\text{chosen}+}$: 0.95 ± 0.01 ; and in $S_{\text{chosen}0}$: 0.14 ± 0.01 ; multilevel Bayesian logistic regression: $\beta_{\text{chosen}+} = 3.74$ [3.30, 4.27], $\beta_{\text{chosen}0} = -2.64$ [-3.04, -2.30], see model specification in the "Analysis" section).

Lastly, we tested whether the inverse decision bias observed in the Final Decisions phase is related to participants' explicit reports in the Outcome Estimation phase. To assess overall performance in the Outcome Estimation phase, we computed an inverse estimation score for chosen and unchosen items separately. This was the difference in the mean probability to estimate an item as S+ between items that were rewarded (S+) and those that were not (S0), the difference between the green and red bars in Fig. 2c; for unchosen items the assigned outcome was that of the chosen item they were paired with). Negative scores will suggest an inverse inference of value (S+ items estimated as S0 and S0 items estimated as S+), and positive scores will suggest a direct inference of value (S+ items estimated as S+ and S0 items estimated as S0). Performance in the Final Decisions phase was computed as a linear transformation of the mean probability to select S+ items (i.e., $p(\text{select S+}) - 0.5$, see the "Analysis" section for details), separately for chosen and unchosen pairs. Negative estimates signify a tendency to choose S0 over S+, and positive estimates signify a tendency to choose S+ over S0. We ran a Bayesian linear regression predicting choices in the Final Decisions phase as a function of choice type (chosen vs. unchosen pairs) and the inverse estimation score from the Outcome Estimation phase. We then rearranged the coefficients to get separate slope estimates of the effect of inverse estimation score on decisions for chosen and unchosen pairs. In both pair types, the slope term was substantial ($\beta_{\text{chosen}} = 0.32$ [0.26, 0.37], $\beta_{\text{unchosen}} = 0.37$ [0.33, 0.40], see Fig. 2d), suggesting that decision bias in the Final Decisions phase was correlated with the explicit estimation reports.

Decision bias is related to associative memory. We hypothesized that the change in the value of the unchosen option would

be related to the strength of the associative memory linking the two options. That is, to transfer value between a pair of options, participants need to have associated in memory the competing options included in each pair. To evaluate the relationship between decision bias and memory, we defined a summary statistic for each. We defined a pair memory score as the mean accuracy in the Surprise Memory phase, where participants were asked whether pairs of paintings appeared together in the Deliberation phase (see Supplementary Text 2 for analysis of the Surprise Memory phase). A decision bias score for each participant was calculated as the difference in the probability of choosing S+ paintings for chosen and unchosen pairs in the Final Decisions phase. Higher decision bias scores depict larger contrast in decision patterns between the two pair types, reflecting a stronger inverse decision bias. Bayesian linear regression between these two measures revealed that memory was associated with decision bias ($\beta_{\text{intercept}} = 0.06$ [-0.08, 0.20]; $\beta_{\text{memory}} = 0.67$ [0.47, 0.88]), such that better memory was related to stronger inverse decision bias across participants (Fig. 3a).

The correlation between memory and decision bias was also evident within participants. To assess within-participant variability, we used the outcome estimation reports of each participant to split their deliberation pairs into two kinds, per participant: (1) "direct transfer" pairs and (2) "inverse transfer" pairs. In the Outcome Estimation phase, for both chosen and unchosen items, we asked participants to judge whether an item results in a gain or no-gain. We defined direct transfer pairs as those for which participants judged both members of the pair in the same manner (judged both as gain or both as no-gain), alluding to the interpretation that the explicitly learned value of the chosen option was directly transferred to the unchosen option. We define the remaining pairs, i.e., those receiving opposed judgements, as inverse transfer pairs. We then compared pair memory performance between the two kinds of deliberation pairs. We found that pairs memory was more accurate for inverse compared to intact transfer pairs (mean \pm SEM memory accuracy for inverse transfer: 0.75 ± 0.01 ; intact transfer: 0.68 ± 0.02 ; multilevel logistic regression predicting memory accuracy: $\beta_{\text{intercept}} = 1.03$ [0.91, 1.16]; $\beta_{\text{pair-type}} = 0.15$ [0.05, 0.25], Fig. 3b). Together, these findings suggest that memory is related to an inverse decision bias, both between- and within-participants.

Critically, our findings do not seem to be related to any explicit assumptions about the task structure. We examined the possibility that participants might have assumed that for every decision they make in the Deliberation phase, one painting would result in a gain and the other would not. While our instructions did not indicate any such relationship, we asked participants at the end of the study about their strategy in deciding between unchosen pairs. 22 out of 235 participants (9%) stated that they assumed that if their chosen painting did not gain in the auction, its associated unchosen painting probably did. The remaining participants stated they either guessed, chose according to their liking, or tried to look for similarities with the gaining chosen paintings. Importantly, the inverse decision bias remained substantial even when we excluded the 22 participants who explicitly used the inverse heuristic (probability to select $S_{\text{unchosen}+}$: 0.45 ± 0.01 , unchosen intercept: $\beta = -0.30$ [-0.44, -0.16], probability to select $S_{\text{chosen}+}$: 0.92 ± 0.01 , chosen intercept: $\beta = 3.31$ [3.09, 3.57]). Moreover, the inverse decision bias was still associated with outcome estimation score (unchosen pairs: $\beta = 0.35$ [0.31, 0.40]) and memory of the initial deliberation pairs ($\beta = 0.56$ [0.33, 0.78]).

Replication of the main results across independent data sets. We ran a series of follow-up experiments to test the robustness of

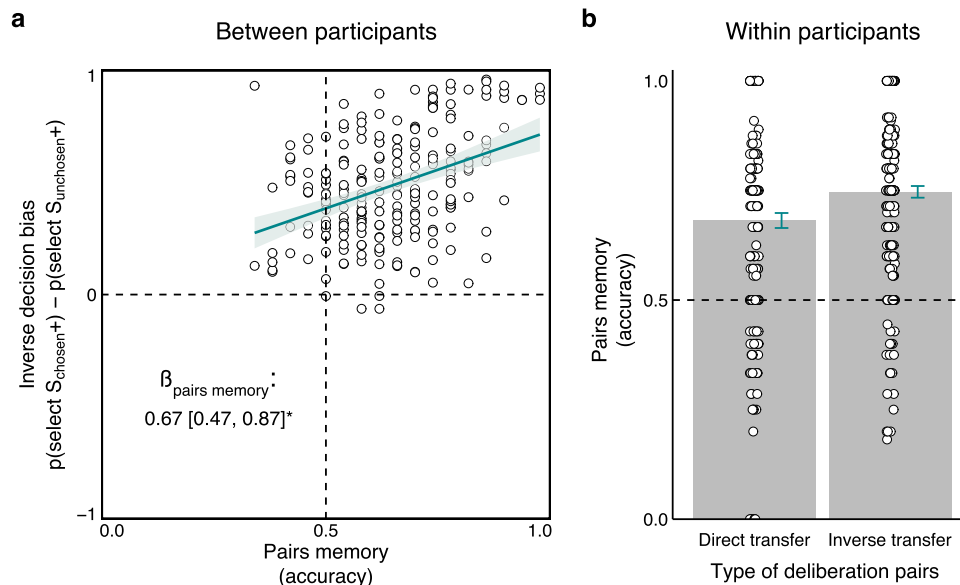


Fig. 3 Inverse inference of value is related to associative memory (Experiment 1, $n = 235$). **a** The group-level inverse decision bias (difference in mean probability to choose rewarded items for chosen and unchosen pairs) is related to memory for the deliberation pairs. Points denote summarized observations of participants and the turquoise line denotes the fit of a Bayesian linear regression predicting inverse decision bias as a function of memory accuracy. Model fit and memory coefficient (beta) depict median and 95% highest density interval estimates. **b** The deliberation pairs were separated into two types for each participant based on outcome estimations (direct transfer: both the chosen and unchosen items within a pair were estimated with the same outcome, inverse transfer: the chosen and unchosen items were estimated to be in opposition). Pairs memory was better for inverse compared to direct transfer within participants. Points denote trial-averaged data of individual participants and error bars denote standard error of the mean. S+ denotes a rewarded stimulus (for unchosen stimuli, this is the outcome of their chosen counterpart). Source data are provided as a Source Data file.

our effects—namely, the inverse decision bias and its relation to memory—and their generalizability across variants of the experimental design. We ran three variants of the task with slight modifications in the Deliberation phase (see details in Supplementary Text 3 and Supplementary Fig 3), while keeping the general structure the same. We varied the number of times each pair was presented during the Deliberation phase (Experiment 2, $n = 96$), the amount of gain that could be won for choosing the rewarded option (Experiment 3, $n = 95$), and the high-level association between the pairs of paintings by informing participants that some of the pairs were painted by the same painter (Experiment 4, $n = 93$). In all experiments, we replicated the inverse decision bias as well as its correlation with memory (see Fig. 4, Supplementary Text 3 and Supplementary Tables 1, 4 and 5 for details).

Discussion

In five experiments (total $n = 612$), we studied how exposure to outcomes of chosen options modulates the value of their unchosen counterparts. We discovered that after outcomes are revealed, the value assigned to unchosen options is inversely related to the learned outcomes of chosen options. This inverse bias manifested both in choice behavior and in explicit value estimation and it was associated with participants' memory of the pairs of options that they had deliberated about to begin with.

Our results have important implications for theories of value updating. Consistent with reinforcement learning models, we show that outcomes change the value assigned to chosen options. With respect to unchosen options, previous studies showed that direct feedback about their hypothetical outcomes leads to changes in value and subsequent behavior^{18,20–22,24,41–43}. In life outside the laboratory, however, we are usually exposed only to the outcomes of our chosen option. Here we show that such ecological exposure is sufficient to facilitate value inference even

for unchosen options, and that this value inference is related to memory.

We found that stronger memory for the deliberated options is related to a stronger discrepancy between the value assigned to the chosen and unchosen options. This result suggests that choosing between options leaves a memory trace. By definition, deliberation is meant to tease apart the value of competing options in the service of making the decision; our findings suggest that deliberation and choice also bind pairs of choice options in memory. Consequently, unchosen options do not vanish from memory after a decision is made, but rather they continue to linger through their link to the chosen options.

We show that participants use the association between choice options to infer the value of unchosen options. This finding complements and extends previous studies reporting transfer of value between associated items in the same direction, which allows agents to generalize reward value across associated exemplars. For example, in the sensory preconditioning task, pairs of neutral items are associated by virtue of appearing in temporal proximity. Subsequently, just one item gains feedback—it is either rewarded or not. When probed to choose between items that did not receive feedback, participants tend to select those previously paired with rewarded items^{28–30,44}. In contrast, our participants tended to avoid the items whose counterpart was previously rewarded. Put in learning terms, when the chosen option proved to be successful, participants' choices in our task reflected avoidance of, rather than approach to, the unchosen option. One important difference between our task and the sensory preconditioning task is the manner in which the association is formed. In both tasks a pair of items appears in close temporal proximity, yet in our task participants are also asked to decide between these items and the act of deliberation seems to result in an inverse association between the deliberated options.

Another interesting difference between the current findings and previous findings of value inference is related to whether the

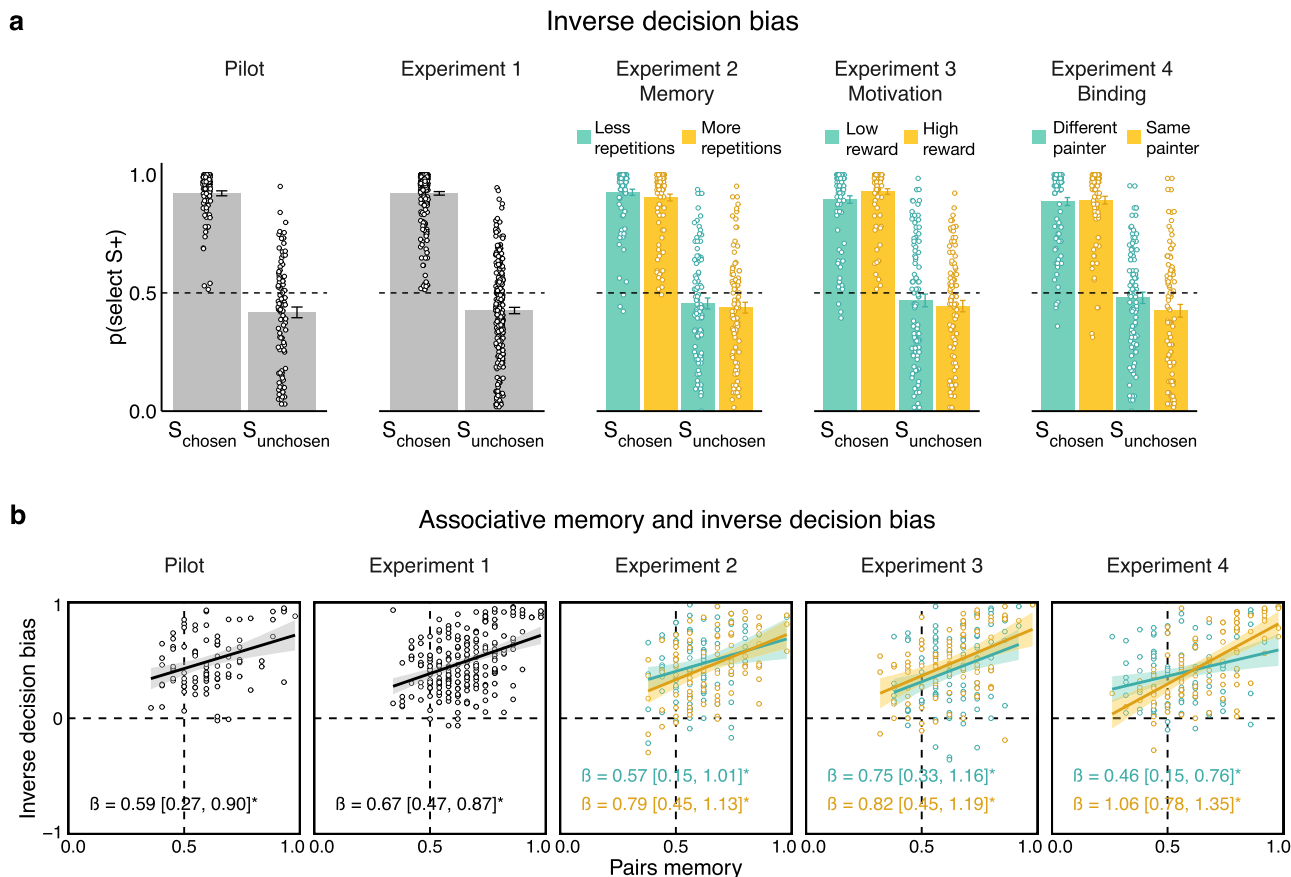


Fig. 4 Replicating the main effects across five distinct data sets. Experiment 1 ($n = 235$) was pre-registered on Open Science Framework (<https://osf.io/chsvw>) following a Pilot study ($n = 93$). Experiments 2 to 4 include different experimental manipulations of the Deliberation phase (depicted by yellow and cyan colors). We manipulated the number of times each deliberation pair was presented (Experiment 2, memory manipulation, $n = 96$), the magnitude of potential gains (Experiment 3, motivation manipulation, $n = 95$), and the high-level association of pairs by assigning some of them as belonging to the same painter (Experiment 4, binding manipulation, $n = 93$). **a** The tendency to select rewarded paintings for chosen pairs ($S_{\text{chosen}+}$), but unrewarded ones for unchosen pairs ($S_{\text{unchosen}0}$) was replicated across conditions. **b** Memory of the deliberation pairs was related to inverse decision bias (difference in decision tendency of chosen and unchosen pairs) in all experiments and conditions. Points denote trial-averaged data of individual participants. In panel (a), error bars denote standard error of the mean. In panel (b), betas represent the memory slope coefficients in a Bayesian linear regression predicting inverse decision bias as a function of memory accuracy, and for both betas and model fits we show median and 95% highest-density intervals of the corresponding posterior distribution. S_+ denotes a rewarded stimulus (for unchosen stimuli, this is the outcome of their chosen counterpart). Source data are provided as a Source Data file.

memory associations driving changes in value are implicit or explicit. In sensory preconditioning and similar tasks (such as acquired equivalence²⁵), memory’s influence on value-based decisions appears to happen outside of conscious awareness and explicit memory for associations is orthogonal to their influence on choice. Here we find, repeatedly, that explicit memory for the association between choice options is related to the extent of the inverse bias in decisions. An important challenge for future work will be to develop a principled understanding of the factors that determine when this relationship is implicit vs. explicit and to what extent it may be related to whether the initial encoding of associations is incidental (as in sensory preconditioning) or whether it is a deeper encoding that itself happens in the context of a decision (as in the current task).

In our data, we observed large individual differences in value transfer from chosen to unchosen options (see Fig. 2a). While the majority of participants transferred value in the opposite direction, some exhibited a direct value transfer. This tendency was evident even between deliberation pairs of the same participant. Notably, individual differences in value transfer is not uncommon in the sensory preconditioning task^{28–30,44}, but future studies could use additional physiological and neuroimaging methods to

find correlates that predict these differences. An additional possible source of variation could be related to retroactive or proactive interference, which may lead to the chosen and unchosen options becoming incorrectly associated. This cannot account for the full effects we observe and, in particular, does not account for the observed relationship between memory and decision bias. Nonetheless, in the future it will be interesting to assess how much of the variance we observe in decision bias is related to forgetting and proactive/retroactive interference.

The inverse inference of value resonates with studies showing that making a decision leads to selective memory for information supporting the decision. The classical finding is that participants remember their chosen option more favorably and remember their unchosen option less favorably^{31,32,34}, and they also exhibit higher learning rates for the positive outcomes of their chosen options and the negative outcomes of their unchosen options^{33,41}. These tasks differ from our task in that participants are exposed directly to the outcomes of their unchosen options and exhibit biases in the memory of these outcomes. Here we show that even without any exposure to the unchosen option’s outcome, participants infer that its value is opposed to that of the chosen option. Importantly, the inverse inference of value does not always

support choices: unchosen options could be inferred as rewarding if the chosen option was not rewarding.

Our findings suggest that the act of deciding can modify an option's value, consistent with studies of choice-induced preference change^{35,36,38,45–47}. These studies show that for two equally valued options, the mere act of deciding induces a tendency to overvalue the chosen option and undervalue the unchosen one. Both options' values are modified right after the decision is made, as a by-product of it. Our study differs from, and therefore extends, choice-induced preference change studies in an important way. In choice-induced preference change, the change in value occurs as a function of a single choice and a choice alone, not as a function of learning from subsequent outcomes. The latter is what our task targets. While our study did not evaluate changes in subjective preferences (e.g., using ratings), but rather changes in the value people assign to choice options after learning about their outcomes, our results echo choice-induced preference-change effects. As with choice-induced preference-change, we find that people tend to separate the value of their choice options after the decision was terminated. Yet, importantly, we find that people use the outcomes they received for their chosen options as a reference point to infer the value of their unchosen options. They do not only devalue these options, but in some cases they might even assume they were better options than their chosen ones. Finally, our findings are also consistent with recent studies showing that choice-induced preference change effects were correlated with a better memory of the initial choice^{35,48–51}.

Our findings add to recent advances showing that episodic memory guides value-based decisions (for a review see ref. 52). Recent studies have found that people consult specific past events when evaluating their next choice^{53,54}. Decision makers are biased to sample a past decision when given an incidental reminder of the decision context⁵⁵, if the option resulted in a high reward and they remember the association between the option and the reward⁵⁶, and when the option seems familiar⁵⁷. Furthermore, recent studies have shown that the simple act of making a decision, rather than having someone else make it for you, increases the episodic memory of that decision^{50,58,59}, thereby making it more accessible for later retrieval. The current study both builds upon and extends these findings by suggesting that memory mechanisms also contribute to the valuation of unchosen options, potentially by retrieving the deliberation context in which the choice options were associated.

An important future direction will be to test whether memory and deliberation are necessary for inverse bias to occur. The present study was designed to assess the effects of deliberation, and memory of the deliberation, on the valuation of unchosen options. To facilitate this intent, we designed the task to increase the chances that participants would actually deliberate and would be able to learn the deliberation pairings as well as the outcomes of the chosen options. Following this logic (and prior studies²⁸), we repeated the pairings and outcomes several times to assure they were learned sufficiently, and asked participants to write down the reasons for choosing one option and not the other to promote deliberation. Furthermore, we did not provide immediate feedback after every decision but rather delayed the feedback and presented it only for the chosen options. This design feature was meant to facilitate the use of long-term memory associations formed during deliberation to transfer value between choice options, rather than doing so while the two options are kept simultaneously in working memory when the feedback is immediate.

As a next step, having found that sufficient deliberation and memory reliably manifest in inverse inference of value, it will be important to further determine the boundary conditions

necessary for this effect. For example, when deliberation is not as extensive or when outcomes are provided immediately after the decision was made (as in standard reinforcement learning tasks), rather than later on. In Experiment 2, we manipulated the memory strength of the deliberation pairs by changing the number of times they were repeated (two vs. six times). Pairs that were repeated more often were remembered better than those repeated less often. Critically, however, there was no difference in decision bias for unchosen pairs between these conditions, nor was there any difference in the correlation of decision bias with memory (see Supplementary Text 3). Future studies could assess the parameters that modulate the inverse inference of value.

In addition, while the decisions in our study involved deliberating and choosing between two options, an open question is what would happen in non-binary decisions. For example, when the decision involves more than two options⁶⁰, when the decision is between engaging with an option or continuing to search for better alternatives (foraging decisions)^{43,61–63}, or when the decision requires to first select the options to consider amongst many alternatives⁶⁴, among other types of decisions. Future studies could assess whether learning about chosen options in non-binary decisions modulates the value of all unchosen alternatives considered in a specific context, only the one dominant unchosen alternative, or none at all. These different possible outcomes would provide insight on how people reason about the interplay between the value of chosen and unchosen options across different contexts.

Our study sheds light on value transfer between chosen and unchosen options and raises questions about when the value transfer takes place. We eliminated the possibility of value transfer during deliberation (or immediately following it) by withholding any feedback at the initial deliberation stage. For the value of unchosen options to change in relation to the chosen ones, participants had to associate the two options to each other in memory. We, therefore, suggest that deliberation creates a memory association between choice options, and this association is later reactivated to facilitate value transfer. Previous studies have suggested that value can transfer during encoding^{28,29} (here, when participants learn the value of chosen options during Outcome Learning phase), or right before new decisions are faced^{44,65} (here, in the Final Decisions phase). While future studies are needed to directly disambiguate these options, our findings provide initial evidence supporting the hypothesis of value transfer during encoding. Specifically, we found that responses were faster when choosing unchosen items previously paired with unrewarded items ($S_{\text{unchosen}0}$), suggesting that these items were perceived as more valuable^{39,40}. These findings allude to the possibility that the positive value of $S_{\text{unchosen}0}$ items was most likely already constructed before the participants faced the new choice in the Final Decisions phase; if the value was constructed on the spot during the decision phase, we would have expected slower, not faster, reaction times.

An important open question is whether the pattern of inverse bias is specifically modulated by counterfactual thinking. That is, when participants learn about the outcomes of their decisions, do they engage in thinking about what would have happened had they chosen the other, unchosen, option, and consequently infer that the value of that option is in opposition to their chosen option. Studies of counterfactual thinking suggest that thoughts about counterfactuals may arise when encountering outcomes of chosen options⁷. Furthermore, recent studies have shown that counterfactual thinking activates similar brain regions as episodic memory^{66,67}, and counterfactual construction can lead to memory distortions⁶⁸.

Relatedly, spontaneous counterfactual thoughts are triggered more frequently after experiencing negative outcomes compared

to positive ones^{7,8,69,70}. Our behavioral test (Final Decisions phase) had participants choose between rewarded and unrewarded stimuli in chosen and unchosen pairs of stimuli, so we could not discern whether the tendency to select $S_{\text{unchosen}0}$ over $S_{\text{unchosen}+}$ stemmed from the inferred positive value of $S_{\text{unchosen}0}$, the inferred negative value of $S_{\text{unchosen}+}$, or both. When participants were asked to explicitly report the potential outcomes of unchosen stimuli (Outcome Estimation phase), there was a greater tendency to report an inverse outcome for $S_{\text{unchosen}+}$ compared to $S_{\text{unchosen}0}$ (see Fig. 2c). That is, unchosen stimuli that were paired with rewarded chosen stimuli were categorized more frequently as unrewarded. Yet this tendency might have originated from a weaker signal for unrewarded chosen stimuli compared to rewarded ones, because we did not introduce losses but only zero gains. Future studies could introduce losses rather than no gains and test whether the asymmetry following negative versus positive outcomes in counterfactual thinking studies also translates to asymmetry in the modulation of the value of unchosen options.

To conclude, our findings suggest that deciding between competing options does not end the competition between them. Deliberation binds choice options together in memory, such that the learned value of one can affect the inferred value of the other.

Methods

Preregistration and sample size determination. Experiment 1 was preregistered on Open Science Framework (<https://osf.io/chsvw>). Prior to data collection we ran a Pilot study ($n = 93$) which confirmed our main hypotheses and guided later sample size decisions. Specifically, we found that participants exhibited an inverse decision bias (selected rewarded items in chosen pairs and unrewarded items in unchosen pairs), which was correlated with pairs memory (see Fig. 4 and Supplementary Tables 1–5). The goal of Experiment 1 was to replicate these findings using a substantially larger sample size. Following⁷¹, we defined the sample size of Experiment 1 to be 2.5 times the sample size of the Pilot study, which is 235 participants (we calculated 2.5 times the original sample size of 93 participants and rounded up by a few). A sample size of 235 participants gave us above 99% power to detect the effect size of the Pilot study, computed based on a simplified version of our main logistic regression model (see the “Analysis” section). Namely, we fitted a logistic function for each participant using “lme4” package⁷² and then tested the unchosen intercept coefficients against zero ($t(92) = -3.71$; $p < 0.001$; Cohen’s $d = 0.39$; power analysis was performed using “pwr” package⁷³ and Cohen’s d computation was performed using “lsr” package⁷⁴).

Participants. The research was approved by the Institutional Review Board (IRB) at Columbia University through Columbia IRB Protocol #AANA1488. 235 Mechanical Turk participants took part in Experiment 1 for an average payment of \$6 (mean age: 28.54 ± 4.49 , 132 female, 101 male, 2 other). All participants provided informed consent for their participation in the experiment. 377 participants participated in the Pilot study and Experiments 2 to 4 (see Supplemental Text 3 for demographic information). Additional 50 participants across all experiments were removed from analyses because they met our predefined exclusion criteria for online studies (see below, all exclusion criteria were preregistered on OSF: <https://osf.io/chsvw>). We added a restriction on Mechanical Turk to include only US-based participants within the age range of 18–36, and an approval rate of above 90%.

Exclusion criteria. We applied the following exclusion criteria which were all aimed to ensure that the online participants were attending the task: (1) Below chance performance (probability to choose gain items below 0.5) in the Final Decisions phase for chosen pairs, indicating participants who did not learn the new values of chosen paintings (for Experiments 2 to 4, the exclusion performance is computed across conditions); (2) More than 25 missed responses in either the Outcome Learning phase (where participants had to register the outcomes they observed) or the Final Decisions phase; (3) More than 25 events where participants were browsing a different window in any experimental phase (blur-focus events detected using jsPsych library⁷⁵); (4) More than 10 trials in the Deliberation phase where responses were too fast (below 300 ms; these trials were accompanied with a warning), signifying no actual deliberation; and (5) More than 10 failed attempts to answer a comprehension quiz administered after instructions in any experimental phase. Participants who met at least one of these exclusion criteria were removed from all analyses.

Materials. Stimuli were images of representational paintings that depict people, landscapes, or objects, randomly intermixed across conditions. The stimuli were

collected by Celia Durkin from various online databases, and a subset of these stimuli are published online⁷⁶. The stimuli were converted to 300×300 pixel size and were presented on a gray background (RGB: 128, 128, 128; see Fig. 1).

Procedure. Experiment 1 included the following consecutive phases: (1) Pre-task Ratings, (2) Deliberation, (3) Outcome Learning, (4) Final Decisions, (5) Surprise Memory, (6) Post-task Ratings, and (7) Outcome Estimation. Each phase began with instructions followed by a comprehension quiz. Participants were not informed about the upcoming phases.

Pre-task ratings. Participants were presented with 60 paintings and were asked to rate their liking of each painting using a continuous scale (from “not at all” to “very much”, responses were then scaled from 0 to 100). At the end of the phase, the paintings were sorted by their ratings, and the 24 paintings rated in the middle of the distribution were selected to construct the deliberation pairs. The selected paintings were shuffled and distributed across 12 pairs.

Deliberation. Participants were instructed to act as art dealers who will choose paintings to sell in an auction and their goal is to maximize profits from that auction. During the Deliberation phase, participants saw 12 pairs of paintings and were asked to deliberate and choose one of the paintings in each pair. They were instructed to take their time and were given up to 10 s to make a decision. If they did not respond during this time, they were prompted to make the decision again and to respond more quickly. To ensure deliberation and encoding of the pairs, we told participants they will practice the decisions a couple of times before committing to their final choices. The 12 deliberation trials were repeated three times in three separate blocks, with random order of trials within each block. Overall, participants were consistent with their initial choices (mean and standard error of the number of times participants repeated the same decision: $2.87 [0.01]$). In subsequent phases, we used the choices made in the last Deliberation block. To increase deliberation, participants were asked to write down the reasons for their decisions using a text box. The text boxes appeared once for every deliberation pair across the first two blocks. Participants were told that one of their decisions will be played out for real, and they will receive 1% of their chosen painting’s auction earnings, if indeed it resulted in a gain. At the end of the Deliberation phase, all participants were informed that they received extra \$1.5 bonus money (we made sure one of the chosen items gained \$150).

Outcome learning. Because we were interested in the long-term effects of associative memory on value updating, we did not provide immediate feedback after every decision. We assumed that after making a choice, participants could still maintain both choice options in their working memory and feedback for the chosen option could transfer to the unchosen option just by virtue of being activated simultaneously in working memory. Thus, only after the completion of all decisions, we presented the auction outcomes for chosen paintings alone. Half of the deliberation pairs were randomly assigned as rewarded items (denoted as S+), with earnings centered around \$150 and standard deviation of \$10. The other half were assigned as unrewarded items, with \$0 earnings (denoted as S0). Paintings were presented in the center of the screen alongside a colored outcome and frame (green for S+, red for S0, Fig. 1) for 2 s. To facilitate learning, we repeated each painting six times across three blocks, in a randomized order. To ensure participants’ attention, we asked them to press the space bar to see the outcome auction and then to register the outcome by pressing a corresponding key while it is presented on the screen (up arrow for gain, down arrow for no gain). If they missed an outcome registration they saw a warning asking them to respond faster.

Final decisions. Participants were asked to prepare a portfolio of high-valued paintings and to do so they had to make additional decisions. They were incentivized to select the more valuable painting in each pair by being informed that they could earn extra bonus money based upon their performance in this phase. The final-decisions pairs were constructed from all possible combinations of S+ and S0 paintings, separately for previously chosen and unchosen paintings (for unchosen paintings, S+ and S0 assignment was according to their chosen counterparts). This yielded 36 unique chosen pairs and 36 unique unchosen pairs, for a total of 72 pairs. To increase the number of trials, we repeated the decision trials three times across three separate blocks, each included all 72 pairs, randomly intermixed. In each block, the rewarded painting appeared on the left and the right sides equal number of times. Participants had 2.5 s to make a decision and, if they failed to respond during this time, they were shown a warning asking them to respond faster. Unbeknownst to participants, the potential bonus money was based only on their performance for the chosen pairs and was up to \$2.

Surprise memory. To test for memory of the deliberation pairs, we presented 24 pairs of paintings and asked participants whether each pair was intact or recombined. Intact pairs were pairs of paintings that appeared in the Deliberation phase, and recombined pairs were pairs that included a chosen painting and an unchosen painting that did not previously appear together. If participants responded “intact”, they were also asked to indicate which of the two paintings they previously had chosen in the Deliberation phase.

Post-task ratings. Similarly to the Pre-task Ratings phase, we presented all 60 paintings again and asked participants to rate how much they like them, using a continuous scale.

Outcome estimation. On this final phase, we told participants that all paintings were sent to auction, including those they didn't choose in the beginning of the experiment, and asked them to estimate the auction outcomes. We presented previously chosen and unchosen paintings in a random order, and for each painting we asked participants whether the painting resulted in a gain or not, and then to rate how confident they are in their response from 1 ("completely unsure") to 6 ("completely sure").

At the end of the experiment we asked participants about their decision strategies throughout the task. We were specifically interested to examine the possibility that participants might have adopted an inverse heuristic. That is, they might have assumed that for every decision made in the Deliberation phase, one painting would result in a gain and the other would not. To this end, we first asked participants what was their decision strategy in the Deliberation phase and then we asked them how they decided between pairs of paintings for which they received direct feedback (chosen pairs) and for which they received no feedback (unchosen pairs) in the Final Decisions phase.

Analysis. All data and analysis codes are publicly available on GitHub⁷⁷. Data analysis was performed in R⁷⁸ using RStudio⁷⁹. All results were analyzed with generalized linear models using the "rstanarm" package⁸⁰, which performs approximate Bayesian inference over the regression coefficients. Instead of a maximum-likelihood procedure that provides a single point-estimate for each coefficient, Bayesian inference targets the full posterior distribution of each coefficient, which combines our prior assumptions and the observed data:

$$p(\beta|x) = \frac{p(\beta)p(x|\beta)}{p(x)} \quad (1)$$

In our models, the denominator of this expression is intractable to compute, preventing an exact solution for the posterior distribution. Therefore, we approximate it by taking samples from it, through a procedure called Hamiltonian Monte Carlo. We improve the fidelity of our samples by running multiple independent sampling processes ("chains"). For every model, we used default priors and ran six chains with 4000 iterations each (2000 iterations per chain were used as warm-up). To determine convergence of each chain, we made sure that for all model coefficients the effective sample size of simulation draws was >900 and the R-hat statistic was around 1.0^{81,82}. To evaluate our effects, for each regression coefficient of interest, we report the median of the posterior samples and their 95% HDI. Since a regression coefficient of zero indicates no relationship between a predictor and an outcome, we determined that a variable reliably predicts an outcome only if its 95% HDI excluded zero.

Moreover, when possible, we ran multilevel (or hierarchical) models, in which a participant's regression coefficient is drawn from a group-level coefficient distribution. Such an approach is more robust to outlier participants and observations. The parameters of this group-level distribution are of special importance, since they indicate whether an effect is present or not across all participants. In all multilevel models, all predictors varied by participants.

Analysis of choices in Final Decisions phase. We ran a multilevel Bayesian logistic regression that predicted the probability of choosing S+ over S0 as a function of (1) choice condition (chosen or unchosen pairs, centered predictor), (2) the difference in pre-task liking ratings between the items in each pair (S+ minus S0), and (3) their interaction. We first define the sigmoid function whose output is bounded between 0 and 1:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

And we model the probability of selecting S+ as

$$p(\text{select S+}) = \sigma(\beta_0 + \beta_1 * \text{choice} + \beta_2 * \Delta\text{ratings} + \beta_3 * \text{choice} : \Delta\text{ratings}) \quad (3)$$

We normalized each participant's ratings to control for differences in the overall use of the scale using a z-score normalization (i.e., some participants tend to use higher/lower values of the continuous scale), and then subtracted the normalized rating of the unrewarded painting (S0) from the rewarded painting (S+) in each decision trial. Trials where participants missed a response were excluded from analyses.

Notably, by design, chosen and unchosen pairs involve different value signals to guide choices; for chosen pairs the outcomes were presented explicitly, whereas for unchosen pairs they can only be inferred. We, therefore, expected the choice coefficient to be substantial (suggesting a difference in decision tendency in chosen versus unchosen pairs), even if participants chose randomly in trials of unchosen pairs. Consequently, we rearranged the equation, effectively decomposing our model into two separate regression models, one for chosen items and one for unchosen ones. Each such model is predicting the probability of gain choice as a function of ratings. We do so by computing a separate $\Delta\text{ratings}$ coefficient and a separate intercept term for chosen and unchosen items. The $\Delta\text{ratings}$ coefficient is a slope term indicating the influence of ratings on choice (computed as $\beta_2 +$

$\beta_3 * \text{choice}$, where $\text{choice} = 1$ for chosen pairs, and $\text{choice} = -1$ for unchosen pairs), and the intercept term quantifies the tendency to choose rewarded items when there is no difference in ratings (computed as $\beta_0 + \beta_1 * \text{choice}$). Our prime measure of interest was the intercept term.

Analysis of outcome estimation. To test whether participants explicitly infer that unchosen paintings are valued in opposition to chosen paintings, we analyzed outcome estimation reports. We ran a multilevel logistic regression predicting the probability to estimate an item as S+ as a function of (1) choice (chosen or unchosen item, centered predictor), (2) actual auction outcomes (S+ or S0, for unchosen items this is the outcome of their associated chosen items, centered predictor), and (3) their interaction.

$$p(\text{estimate as S+}) = \sigma(\beta_0 + \beta_1 * \text{choice} + \beta_2 * \text{outcome} + \beta_3 * \text{choice} : \text{outcome}) \quad (4)$$

To evaluate the probability that items were estimated as rewarded in a specific condition (e.g., unchosen items previously paired with rewarded items, i.e., $S_{\text{unchosen}+}$), we plugged in the relevant predictors and summed the coefficients (e.g., for $S_{\text{unchosen}+}$, $\text{choice} = -1$, $\text{outcome} = 1$, therefore $\beta = \beta_0 + (-1) * \beta_1 + 1 * \beta_2 + 1 * (-1) * \beta_3$).

Analysis of the relationship between decision bias and outcome estimations. To test whether choices in the Final Decisions phase are predicted by outcome estimations, we computed an inverse estimation score for chosen and unchosen items separately. This is the difference in the mean probability to estimate an item as S+ between items that were rewarded (S+) and those that were not (S0, for unchosen items, the assigned outcome was that of the chosen item they were paired with). Negative scores signify an inverse inference of value, i.e., S+ items estimated as S0 and S0 items estimated as S+. Performance in the Final Decisions phase was measured as the mean probability to select S+ items (over S0), separately for chosen and unchosen pairs. To allow interpretability of the intercept terms in this model, we had to make sure that 0 value had a significant meaning in our analysis. Accordingly, we subtracted 0.5 from the mean probability to select S+ items. Negative estimates signify a tendency to choose S0 over S+, and positive estimates signify a tendency to choose S+ over S0.

We ran a linear regression predicting the probability to select S+ as a function of (1) choice (chosen or unchosen pairs, centered predictor), (2) inverse estimation score, and (3) their interaction. Because we computed inverse estimation scores separately for chosen and unchosen items, we could assign these scores to their respective pairs in the Final Decisions phase.

$$p(\text{select S+}) = \sigma(\beta_0 + \beta_1 * \text{choice} + \beta_2 * \text{estimation score} + \beta_3 * \text{choice} : \text{estimation score}) \quad (5)$$

We then rearranged the regression coefficients to effectively get two regression models, one for chosen pairs and the other for unchosen pairs. Each model is predicting the probability to select S+ as a function of inverse estimation score. For every model, we computed an intercept term ($\beta_0 + \beta_1 * \text{choice}$), quantifying the mean probability to select S+ items when there is no difference between estimation scores, and a slope term ($\beta_2 + \beta_3 * \text{choice}$), quantifying the influence of inverse estimation score on choices. Our measure of interest was the slope term.

Analysis of associative memory and inverse inference of value across and within participants. We first assessed the relationship between associative memory and inverse decision bias across participants. To this end, for each participant, we computed two measures. The first was average accuracy in pairs memory responses collected during the Surprise Memory phase. The second was a decision bias score, measuring the contrast in decision tendency between chosen and unchosen pairs. This was operationalized as the mean probability to choose a rewarded item in chosen pairs minus unchosen pairs. Values closer to 1 signify a large inverse decision bias effect. We used a Bayesian linear regression to predict inverse decision bias as a function of memory. The coefficient of interest in this model is the memory slope term.

To assess the relationship between associative memory and inverse estimation of value within participants we used responses in the Outcome Estimation phase, where participants were asked to estimate the auction outcomes of all paintings. This allowed us to retroactively separate the deliberation pairs of each participant into two pair types: (1) direct transfer—both the chosen and unchosen paintings within a pair were later estimated with the same outcome, or (2) inverse transfer—the two paintings were later estimated with the opposite outcome. We then wanted to assess whether there was a difference in associative memory of these two pair types. To this end, we ran a multilevel logistic regression predicting memory accuracy from the Surprise Memory phase as a function of pair type (direct vs. inverse transfer). For ease of analysis, we only used responses from the "intact" trials where pairs appeared as they were presented during the Deliberation phase.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Behavioral data generated and analyzed during the current study are available in GitHub⁷⁷ with the identifier (data <https://doi.org/10.5281/zenodo.4926569>). Source data are provided with this paper.

Received: 22 February 2021; Accepted: 8 July 2021;

Published online: 30 July 2021

References

- Roese, N. J. Counterfactual thinking. *Psychol. Bull.* **121**, 133–148 (1997).
- Zeelenberg, M. et al. Emotional reactions to the outcomes of decisions: the role of counterfactual thought in the experience of regret and disappointment. *Organ. Behav. Hum. Decis. Process.* **75**, 117–141 (1998).
- Kahneman, D. in *What might have been: The Social Psychology of Counterfactual Thinking* (eds Roese, N. J. & Olson, J. M.) 375–396 (Lawrence Erlbaum Associates, Inc., 1995).
- Gilovich, T. & Medvec, V. H. The experience of regret: what, when, and why. *Psychol. Rev.* **102**, 379 (1995).
- Connolly, T. & Zeelenberg, M. Regret in decision making. *Curr. Dir. Psychol. Sci.* **11**, 212–216 (2003).
- Bell, D. E. Regret in decision making under uncertainty. *Oper. Res.* **30**, 961–981 (1982).
- Roese, N. J. & Epstude, K. *The Functional Theory of Counterfactual Thinking: New Evidence, New Challenges, New Insights. Advances in Experimental Social Psychology* (Elsevier Inc., 2017).
- Shani, Y. & Zeelenberg, M. When and why do we want to know? How experienced regret promotes post-decision information search. *J. Behav. Decis. Mak.* **20**, 207–222 (2007).
- Steiner, A. P. & Redish, A. D. The road not taken: Neural correlates of decision making in orbitofrontal cortex. *Front. Neurosci.* **6**, 1–21 (2012).
- Sutton, R. S. & Barto, A. G. *Introduction to Reinforcement Learning*. **135**, (MIT press Cambridge, 1998).
- Vlaev, I., Chater, N., Stewart, N. & Brown, G. D. A. Does the brain calculate value? *Trends Cogn. Sci.* **15**, 546–554 (2011).
- Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
- Tversky, A. & Simonson, I. Context-dependent preferences. *Manag. Sci.* **39**, 1179–1189 (1993).
- Schapiro, A. C., Kustner, L. V. & Turk-Browne, N. B. Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Curr. Biol.* **22**, 1622–1627 (2012).
- Howard, M. W. & Kahana, M. J. A distributed representation of temporal context. *J. Math. Psychol.* **46**, 269–299 (2002).
- Ezzayat, Y. & Davachi, L. Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron* **81**, 1179–1189 (2014).
- Eichenbaum, H. Memory on time. *Trends Cogn. Sci.* **17**, 81–88 (2013).
- Kishida, K. T. et al. Subsecond dopamine fluctuations in human striatum encode superposed error signals about actual and counterfactual reward. *Proc. Natl Acad. Sci. USA* **113**, 200–5 (2016).
- Boorman, E. D., Behrens, T. E. & Rushworth, M. F. Counterfactual choice and learning in a Neural Network centered on human lateral frontopolar cortex. *PLoS Biol.* **9**, e1001093 (2011).
- Coricelli, G. et al. Regret and its avoidance: A neuroimaging study of choice behavior. *Nat. Neurosci.* **8**, 1255–1262 (2005).
- Abe, H. & Lee, D. Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron* **70**, 731–741 (2011).
- Büchel, C., Brassen, S., Yacubian, J., Kalisch, R. & Sommer, T. Ventral striatal signal changes represent missed opportunities and predict future choice. *Neuroimage* **57**, 1124–1130 (2011).
- Palminteri, S., Lefebvre, G., Kilford, E. J. & Blakemore, J. Counterfactual Confirmation Bias. 1–22 (2017).
- Hayden, B. Y., Pearson, J. M. & Platt, M. L. Fictive reward signals in the anterior cingulate cortex. *Science* **324**, 948–950 (2009).
- Shohamy, D. & Wagner, A. D. Integrating memories in the human brain: hippocampal-midbrain encoding of overlapping events. *Neuron* **60**, 378–389 (2008).
- Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).
- Dusek, J. A. & Eichenbaum, H. The hippocampus and memory for orderly stimulus relations (ratsentorhinal and perirhinal corticesrelational representationdeclarative memory). *Psychology* **94**, 7109–7114 (1997).
- Wimmer, G. E. & Shohamy, D. Preference by Association. *Science* **338**, 270–273 (2012).
- Wong, F. S., Westbrook, R. F. & Holmes, N. M. 'Online' integration of sensory and fear memories in the rat medial temporal lobe. *Elife* **8**, e47085 (2019).
- Liashenko, A., Dizaji, A. S., Melloni, L. & Schwiedrzik, C. M. Memory guidance of value-based decision making at an abstract level of representation. *Sci. Rep.* **10**, 1–12 (2020).
- Mather, M., Shafir, E. & Johnson, M. K. Misremembrance of options past: source monitoring and choice. *Psychol. Sci.* **11**, 132–138 (2000).
- Lind, M., Visentini, M., Mäntylä, T. & Del Missier, F. Choice-supportive misremembering: a new taxonomy and review. *Front. Psychol.* **8**, (2017).
- Palminteri, S., Lefebvre, G., Kilford, E. J. & Blakemore, S. J. Confirmation bias in human reinforcement learning: evidence from counterfactual feedback processing. *PLoS Comput. Biol.* **13**, e1005684 (2017).
- Henkel, L. A. & Mather, M. Memory attributions for choices: how beliefs shape our memories. *J. Mem. Lang.* **57**, 163–176 (2007).
- Voigt, K., Murawski, C. & Bode, S. Endogenous formation of preferences: choices systematically change willingness-to-pay for goods. *J. Exp. Psychol. Learn. Mem. Cogn.* **43**, 1872–1882 (2017).
- Sharot, T., Velasquez, C. M. & Dolan, R. J. Do decisions shape preference? Evidence from blind choice. *Psychol. Sci.* **21**, 1231–1235 (2010).
- Festinger, L. *A Theory of Cognitive Dissonance*. (Stanford university press, 1957).
- Bem, D. J. Self-perception: an alternative interpretation of cognitive dissonance phenomena. *Psychol. Rev.* **74**, 183 (1967).
- Chen, M. & Bargh, J. A. Consequences of automatic evaluation: immediate behavioral predispositions to approach or avoid the stimulus. *Personal. Soc. Psychol. Bull.* **25**, 215–224 (1999).
- Duckworth, K. L., Bargh, J. A., Garcia, M. & Chaiken, S. The automatic evaluation of novel stimuli. *Psychol. Sci.* **13**, 513–519 (2002).
- Palminteri, S., Khamassi, M., Joffily, M. & Coricelli, G. Contextual modulation of value signals in reward and punishment learning. *Nat. Commun.* **6**, 8096 (2015).
- Boorman, E. D., Behrens, T. E. & Rushworth, M. F. Counterfactual choice and learning in a Neural Network centered on human lateral frontopolar cortex. *PLoS Biol.* **9**, e1001093 (2011).
- Sweis, B. M., Thomas, M. J. & Redish, A. D. Mice learn to avoid regret. *PLoS Biol.* **16**, 1–21 (2018).
- Jones, J. & Schoenbaum, G. Inferred but not cached values. *Science* **338**, 953–956 (2012).
- Sharot, T., Fleming, S. M., Yu, X., Koster, R. & Dolan, R. J. Is choice-induced preference change long lasting? *Psychol. Sci.* **23**, 1123–1129 (2012).
- Festinger, L. *Conflict, Decision, and Dissonance*. (Stanford University Press, Stanford CA, 1964).
- Chen, M. K. & Risen, J. L. How choice affects and reflects preferences: revisiting the free-choice paradigm. *J. Pers. Soc. Psychol.* **99**, 573–594 (2010).
- Chammat, M. et al. Cognitive dissonance resolution depends on episodic memory. *Sci. Rep.* **7**, 1–10 (2017).
- Salti, M., El Karoui, I., Maillet, M. & Naccache, L. Cognitive dissonance resolution is related to episodic memory. *PLoS ONE* **9**, 1–8 (2014).
- DuBrow, S., Eberts, E. A. & Murty, V. P. A common mechanism underlying choice's influence on preference and memory. *Psychon. Bull. Rev.* **26**, 1958–1966 (2019).
- Hagège, J., Chammat, M., Tandetnik, C. & Naccache, L. Suggestion of self-(in) coherence modulates cognitive dissonance. *PLoS ONE* **13**, 1–10 (2018).
- Biderman, N., Bakkour, A. & Shohamy, D. What are memories for? the hippocampus bridges past experience with future decisions. *Trends Cogn. Sci.* **24**, 542–556 (2020).
- Hertwig, R. & Erev, I. The description-experience gap in risky choice. *Trends Cogn. Sci.* **13**, 517–523 (2009).
- Gershman, S. J. & Daw, N. D. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annu. Rev. Psychol.* **68**, 101–128 (2017).
- Bornstein, A. M., Khaw, M. W., Shohamy, D. & Daw, N. D. Reminders of past choices bias decisions for reward in humans. *Nat. Commun.* **8**, 1–9 (2017).
- Murty, V. P., FeldmanHall, O., Hunter, L. E., Phelps, E. A. & Davachi, L. Episodic memories predict adaptive value-based decision-making. *J. Exp. Psychol. Gen.* **145**, 1–11 (2016).
- Duncan, K., Semmler, A. & Shohamy, D. Modulating the use of multiple memory systems in value-based decisions with contextual novelty. *J. Cogn. Neurosci.* **31**, 1455–1467 (2019).
- Murty, V. P., DuBrow, S. & Davachi, L. The simple act of choosing influences declarative memory. *J. Neurosci.* **35**, 6255–6264 (2015).
- Murty, V. P., DuBrow, S. & Davachi, L. Decision-making increases episodic memory via postencoding consolidation. *J. Cogn. Neurosci.* **31**, 1308–1317 (2019).

60. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
61. Steiner, A. P. & Redish, A. D. Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nat. Neurosci.* **17**, 995–1002 (2014).
62. Abram, S. V., Hanke, M., Redish, A. D. & MacDonald, A. W. Neural signatures underlying deliberation in human foraging decisions. *Cogn. Affect. Behav. Neurosci.* **19**, 1492–1508 (2019).
63. Kolling, N., Behrens, T. E. J., Mars, R. B. & Rushworth, M. F. S. Neural mechanisms of foraging. *Science* **336**, 95–98 (2012).
64. Phillips, J., Morris, A. & Cushman, F. How we know what not to think. *Trends Cogn. Sci.* **23**, 1026–1040 (2019).
65. Wikenheiser, A. M. & Schoenbaum, G. Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nat. Rev. Neurosci.* **17**, 513–523 (2016).
66. Schacter, D. L., Benoit, R. G., De Brigard, F. & Szpunar, K. K. Episodic future thinking and episodic counterfactual thinking: intersections between memory and decisions. *Neurobiol. Learn. Mem.* **117**, 14–21 (2015).
67. De Brigard, F. & Parikh, N. Episodic counterfactual thinking. *Curr. Dir. Psychol. Sci.* **28**, 59–66 (2019).
68. Gerlach, K. D., Dornblaser, D. W. & Schacter, D. L. Adaptive constructive processes and memory accuracy: consequences of counterfactual simulations in young and older adults. *Memory* **22**, 145–162 (2014).
69. Petrocelli, J. V. & Harris, A. K. Learning inhibition in the monty hall problem: the role of dysfunctional counterfactual prescriptions. *Personal. Soc. Psychol. Bull.* **37**, 1297–1311 (2011).
70. Hafner, R. J., White, M. P. & Handley, S. J. Spoilt for choice: the role of counterfactual thinking in the excess choice and reversibility paradoxes. *J. Exp. Soc. Psychol.* **48**, 28–36 (2012).
71. Simonsohn, U. Small telescopes: detectability and the evaluation of replication results. *Psychol. Sci.* **26**, 559–569 (2015).
72. Bates, D., Sarkar, D., Bates, M. D. & Matrix, L. The lme4 package. *R. Packag. version 2*, 74 (2007).
73. Champely, S. et al. Package 'pwr'. *R Packag. version 1–2* (2018).
74. Navarro, D. *Learning statistics with R: A tutorial for psychology students and other beginners: Version 0.5*. (University of Adelaide Adelaide, Australia, 2013).
75. de Leeuw, J. R. jsPsych: a JavaScript library for creating behavioral experiments in a Web browser. *Behav. Res. Methods* **47**, 1–12 (2015).
76. Durkin, C., Hartnett, E., Shohamy, D. & Kandel, E. R. An objective evaluation of the beholder's response to abstract and figurative art based on construal level theory. *Proc. Natl Acad. Sci. USA* **117**, 19809–19815 (2020).
77. Biderman, N. & Shohamy, D. Memory and decision making interact to shape the value of unchosen options. *GitHub* <https://doi.org/10.5281/zenodo.4926568> (2021).
78. R Core Team (2013). R: A language and environment for statistical computing. *R Found. Stat. Comput. Vienna, Austria.* (2017). S0103-64402004000300015.
79. Team, R. RStudio: integrated development for R. RStudio, PBC, Boston. (2020).
80. Gabry, J. & Goodrich, B. rstanarm: Bayesian applied regression modeling via stan. *R. package version 2.1* (2016).
81. Gelman, A. et al. *Bayesian Data Analysis: Third Edition. Book* (Academic Press, 2013).
82. Gelman, A. & Rubin, D. B. Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**, 457–472 (1992).

Acknowledgements

We thank Dan Biderman for insightful discussions and for providing feedback on an earlier draft. We also thank Celia Durkin for comments on an earlier draft and the Shohamy lab for helpful conversations on this project. We are grateful for funding support from the NSF (award # 1822619), NIMH/NIH (award # MH121093), and the Templeton Foundation (grant #60844).

Author contributions

N.B. and D.S. designed the experiments. N.B. programmed the experiments, collected the data and performed data analysis. N.B. and D.S. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-021-24907-x>.

Correspondence and requests for materials should be addressed to N.B. or D.S.

Peer review information *Nature Communications* thanks Brice Kuhl, Felipe De Brigard, Vishnu Murty and the other, anonymous, reviewer for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permission information is available at <http://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) 2021