The reflection effect in memory-based decisions

Regina A. Weilbächer*, Peter M. Kraemer*, Sebastian Gluth*

Department of Psychology, University of Basel, Switzerland

Author Note

Accepted Version Psychological Science, June 2020

Corresponding Author:

Regina A. Weilbächer

Address: Missionstrasse 62A, 4055 Basel

Phone: +41 61 207 39 05

E-mail: regina.weilbaecher@unibas.ch

ORCID-ID: https://orcid.org/0000-0003-1955-8664

Abstract

Previous research indicated a bias in memory-based decision making, with people preferring options that they remember better. However, the cognitive mechanisms underlying this memory bias remain elusive. Here, we propose that choosing poorly remembered options is conceptually similar to choosing options with uncertain outcomes. We predicted that the memory bias is reduced when options have negative subjective value, analogous to the reflection effect, according to which uncertainty aversion is stronger in gains than in losses. In two preregistered experiments (n = 36 each), participants made memory-based decisions between appetitive or aversive stimuli. People preferred better-remembered options in the gain domain, but this behavioral pattern reversed in the loss domain. This effect was not related to participants' ambiguity or risk attitudes, as measured in a separate task. Our results increase the understanding of memory-based decision making and connect this emerging field to well-established research on decisions under uncertainty.

Keywords: decision-making, episodic memory, uncertainty

Statement of Relevance

Many decisions in our everyday life, such as choosing whether to have the same lunch meal as yesterday, are shaped by our memories. However, we are just beginning to understand how memories and decisions interact. Based on the proposal that choosing a poorly remembered option is conceptually similar to choosing an option with uncertain outcomes, the present study draws an analogy between decisions from memory and decisions under uncertainty. In line with this rationale, we find that decisions from memory elicit a preference reversal between gains and losses that mirrors the well-known reflection effect in decisions under uncertainty: People prefer better-remembered over less-remembered options in the gain domain, but exhibit the opposite preference in the loss domain. Our findings connect two hitherto separate branches of decision-making research and have potentially broad implications for understanding the impact of aging- or disease-related changes in memory abilities on behavior.

Introduction

Many of our daily choices require us to retrieve relevant information from memory, and the role of memory in shaping such value-based decisions is receiving growing interest (Shadlen & Shohamy, 2016; Weilbächer & Gluth, 2017; Wimmer & Büchel, 2016; Weber & Johnson, 2006; Murty, Feldmanhall, Hunter, Phelps, & Davachi, 2016; Gershman & Daw, 2017). Episodic memory and decision making were studied separately for decades, but more recent studies have started to investigate how these two psychological constructs interact (Murty et al., 2016; Gershman & Daw, 2017; Wimmer & Büchel, 2016). In our previous work, we have shown that memory-based decisions can give rise to a bias in choice behavior (Gluth, Sommer, Rieskamp, & Büchel, 2015; Mechera-Ostrovsky & Gluth, 2018). More precisely, the *memory bias* in preferential choice states that people tend to prefer options they remember better to an extent that is not compatible with standard notions of optimality and utility maximization. To illustrate this, assume a decision between two hiking locations, A and B, of equal subjective value. Remembering past experiences of hiking location A more vividly than B will induce a preference for A over B. In fact, our findings suggest that even if A has somewhat lower subjective value than B, the memory bias still induces a preference for A.

An open question is why people exhibit this memory bias, or stated differently, what the cognitive mechanisms are that drive this effect. Here, we argue that uncertainty plays a critical role in decisions from memory and can explain why people show a memory bias. We assume that choosing between a vividly remembered and a poorly remembered option is conceptually similar to choosing between a certain and an uncertain option. Our argumentation follows a recent proposal that people retrieve past instances from their memory when deliberating on the likely consequences of choosing an option(Shadlen & Shohamy, 2016; Bakkour et al., 2019). Accordingly, a more vivid memory of previous encounters with an option (e.g., previous hiking trips to location A) provides higher confidence about its subjective value. In contrast, a poorly remembered option entails greater ambiguity about its potential consequences given that previous encounters of it cannot be remembered so well anymore (e.g., whether hiking location B might have included some dangerous parts). This renders the poorly remembered option an uncertain choice candidate. Importantly, research on decision making under risk (Tversky & Kahneman, 1981; Kahneman & Tversky, 1979) and ambiguity (Kahn & Sarin, 1988; Viscusi & Magat, 1992) have demonstrated a *reflection effect*, meaning that uncertainty aversion is less pronounced and sometimes even reverses in the loss as compared to the gain domain. Thus, we derive the analogous hypothesis that the memory bias is stronger in the gain domain than in the loss domain. Stated differently, when choosing between two appetitive options, we predict people to exhibit a preference for the option they remember better and whose consequences are more certain to them. But when having to choose from aversive options, this tendency should be decreased and possibly reversed, implying that people go with the less-remembered and uncertain option in hope that its consequences might not turn out to be so bad.

To investigate this hypothesis, we leveraged our *remember-and-decide* task (Gluth et al., 2015; Mechera-Ostrovsky & Gluth, 2018), in which participants first learn to associate different choice options with different locations and then make a series of value-based decisions between two highlighted locations (Fig. 1). Since only the locations but not the choice options themselves are presented, participants need to recall the options from their memory when making decisions. Following these decisions, participants are then asked to recall the option-location associations, which allows us to identify remembered and forgotten options (i.e., the label *forgotten* refers to options which could not be recalled correctly). Ultimately, the memory bias is quantified by the strength of preference for remembered over forgotten options. So far, this task has only been used with appetitive (food snack) items. As we aimed to investigate decisions in the gain as well as in the loss domain in the current study, we used appetitive and aversive images in one experiment and positive and negative monetary amounts in a second independent experiment. Beyond testing our main hypothesis, these two experiments with their different sets of stimuli also allowed us to assess to what extent the memory bias (and its putative reversal in the loss domain) generalizes to different domains of decision making.

In essence, we found that participants indeed preferred remembered over forgotten options in the gain domain, but showed the opposite pattern in the loss domain. This confirms our hypothesis of an analogy between decisions from memory and decisions under uncertainty. Thus, concerning options of positive subjective value, people stick to better remembered options and avoid the risk of choosing what they do not recall well. But when it comes to options of negative subjective value, people take the risk of choosing the unknown.

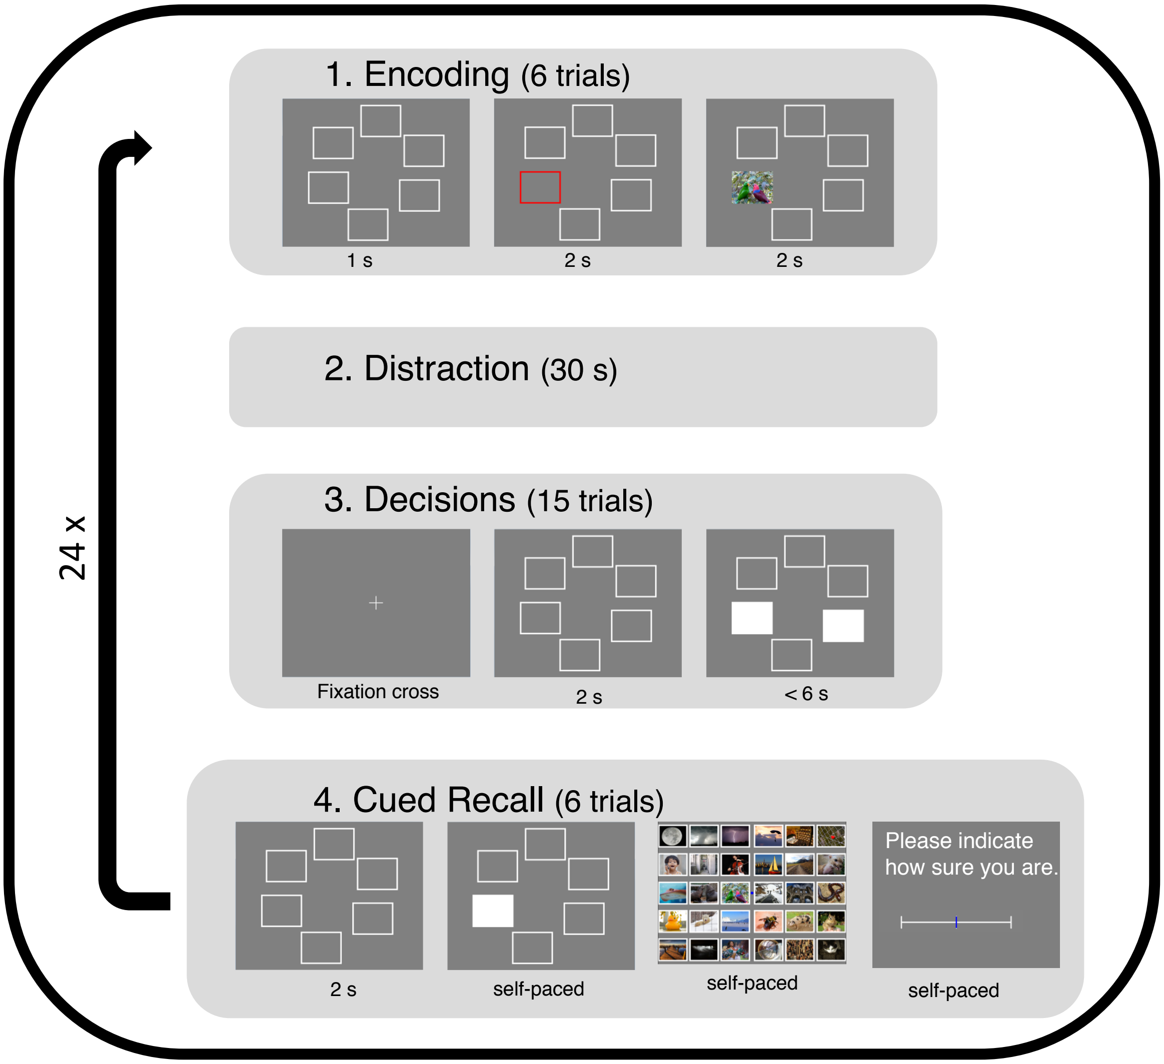

Fig. 1. The remember-and-decide task for the Images Experiment. Participants first encode the association of images with locations on the screen. After a distracting working-memory task, they make binary decisions between the images from memory. Finally, they are asked to recall each image. This procedure is repeated in 24 rounds, 12 rounds with options of positive subjective value and 12 rounds with options of negative subjective value. The figure displays one example round. The procedure was analogous for the Money Experiment but with positive and negative monetary values as choice options (the analogous figure for the Money Experiment can be found in the Supplementary Online Material [SOM]).

Method

We preregistered our two experiments (including hypotheses, experimental design, and analysis plan) on the Open Science Framework website (https://osf.io/eumj5). The complete project (including the processed data and data analysis files in the programming languages R and Python) can be found here https://osf.io/x935r/.

Sample size rationale

In the preregistration protocol, we proposed to perform a two-step analysis approach (i.e., first hierarchical Bayesian parameter estimation, second frequentist t-tests on the mean parameter estimates; see below) and performed a power analysis to estimate the required sample size.

The central hypothesis of our experiments was the difference of the memory bias in gains and losses. While the memory bias itself appears to be a strong effect (effect sizes Cohen's d in previous studies (Gluth et al., 2015; Mechera-Ostrovsky & Gluth, 2018) were between 0.7 and 1.0), the effect size of the difference between gains and losses is unknown. Therefore, we assumed a medium effect size of $d = 0.5$. We used the software program G*Power (version 3.1.9.4) to conduct a power analysis (paired-sample t-test, one-tailed, effect size $d = 0.5$, alpha error probability = .05, power = .9), which suggested a sample size of n = 36 participants. Note that we selected a power of .9 for each experiment, because we sought to achieve a power of greater than .8 across both experiments combined (i.e., $.9^2 \sim .8$).

Participants

Participants were recruited at the University of Basel (convenience sampling). In the Images Experiment a total of 53 participants started the experiment. In the Money Experiment a total of 47 participants started the experiment. Participants were between 18 and 35 years old, did not suffer from mental disorders, and were allowed to participate in only one of the two experiments.

Based on our preregistered exclusion criteria, we did not analyze the data of 17 participants from the Images Experiment for the following reasons: The participant aborted the study (n=2), did not perform all tasks (n=1), were not in the targeted age range (n=1), did not yield the minimal number of trials for the logistic regression analysis (see below; n=8), rated less than 30 images as positive or less than 30 images as negative (n=5). For the Money experiment, we did not analyze the data of 11 participants for the following reasons: Participant aborted the study (n=5), did not perform all tasks (n=2), were not in the targeted age range (n=1), did not yield the minimal number of trials for the logistic regression analysis (n=2), did not understand the n-back task (n=1). We thus included n=36 participants for the analysis of the Images Experiment (25 women, age: range 18-34, M = 23.94, SD = 4.45), and n=36 participants for the Money Experiment (26 women, age: range 18-35, M = 24.42, SD = 4.32).

Participants could only take part in the study after reading and signing the informed consent form, which had been approved by the ethics committee of north-west and central Switzerland (EKNZ). Participants were reimbursed 5 CHF for every started 15 minutes (resulting in 20 CHF per hour). Psychology students of the University of Basel had the opportunity to receive course credits instead of the monetary reimbursement. Additionally, in the Images Experiment participants received a bonus in the sense of looking at the image that they selected in a randomly selected choice trial (details provided below). In the Money Experiment participants could earn a monetary bonus between 0 and 9.50 CHF based on their decision in a randomly selected choice trial. In the additional gambles task that participants in both experiments performed on a separate day (details provided below), participants also had the opportunity to receive a monetary bonus between 0 and 60 CHF.

Apparatus and Stimuli

Participants were seated in front of a 24-in. computer screen (resolution 1680 x 1050 pixel, refresh rate 60Hz). Stimulus presentation and creation of choice sets were realized using MATLAB Version R2016a and its toolbox Cogent 2000 (version 1.33). The screen resolution was set to 1280 x 1024 pixel.

The images for the Images Experiment were selected from the OASIS database (Kurdi, Lozano, & Banaji, 2017). The OASIS database includes a total of 900 images. To reduce the amount of images for our preference rating task, we first excluded all images with arousal and/or valence ratings <2 and >5 (ratings were on a scale from 1 to 7, rated by a large sample of $n=822$ participants; details provided in Kurdi et al. (2017)). Thereby, we excluded too arousing images (e.g., mutilations) and not arousing images, to avoid that some images were much more memorable than others. Second, we excluded all images from the category "Nudes". Third, we renamed redundant categories (e.g. "Graveyard" and "Cemetery"). Fourth, we selected one image per category (e.g., if 5 images displayed a dog, one of them which was judged to be most representative was selected). This procedure resulted in a final set of 103 images.

Experimental Procedures

We conducted two independent experiments, and each participant was allowed to take part in only one of them. Each experiment consisted of two sessions, performed with a delay of one week $\pm$ three days. The two experiments differed only in the used stimulus material: images or money. In the first session, after participants gave their informed consent, they read the instructions for the remember-and-decide task (as in Gluth et al., 2015; Mechera-Ostrovsky and Gluth, 2018). The task consists of multiple blocks of 4 phases each: 1) encoding of the association of six locations and the choice options (money or images), 2) 2-back working memory task, 3) binary choice task (in which the options need to be recalled from memory, as only the locations are presented), and 4) cued recall of the six stimuli (Fig. 1). Participants were familiarized with the task by performing two training rounds. Afterwards, they conducted 24 rounds in total – 12 times with negative stimuli and 12 times with positive stimuli. The order (positive or negative first) was counterbalanced across participants. In the Images Experiment participants made one break between the two blocks (12 rounds). In the Money Experiment participants made a break after each quarter of the rounds (6 rounds).

In the Images Experiment, participants' subjective value of the images was assessed with an incentivized rating task prior to the remember-and-decide task. Participants rated the images on a discrete rating scale ranging from -10 to 10 in steps of 1. They were asked to use the entire range of the rating scale and rated every image twice. To incentivize the rating task, participants were informed that at the end of the experiment two images were drawn randomly, and that the higher-rated image was presented to them for 3 minutes. The 103 rated images were divided into positive and negative images, based on the mean rating value. For the *remember-and-decide* task, at least 30 positively and 30 negatively rated images were needed to generate enough trials. Based on previous experience (Gluth et al., 2015; Mechera-Ostrovsky & Gluth, 2018), the two images with the most extreme negative and positive ratings were excluded, because people tend to have exceptionally good memory for these items. In case a participant had rated less than 30 images as positive on average or less than 30 images as negative on average (for example when the participant used only the negative part of the rating scale), the participant was informed that it was not possible to generate enough trials and the experiment was aborted. In this case, the data being collected so far was not used for data analysis. In the Money Experiment, the positive (appetitive) stimuli were monetary values ranging from 10 to 95 in experimental currency unit (ECU) which were translated into Swiss Francs (CHF) by being divided by 20 (e.g. 95 ECU = 4.75 CHF). Similarly, the negative (aversive) stimuli were monetary values ranging from -95 to -10 in steps of 5, resulting in 18 stimuli each. Participants could win up to 4.75 CHF from the gains trials and the loss trials, respectively, resulting in a possible bonus of 9.5 CHF. In the gains lottery, participants earned the monetary amount they chose in the randomly selected choice trial, whereas in the losses the amount of the choice was subtracted from and initial endowment of 4.75 CHF (e.g., if a trial was selected, where the participant choose -50 ECU, she received the following bonus: 4.75 - [50 ECU / 20] = 2.25).

This first session lasted approximately 75 to 90 minutes. On average the Images Experiment lasted approximately 15 to 20 minutes longer than the Money Experiment, because of the additional rating task. At the end of session 1, participants could indicate in an answer box which strategy they used to memorize the stimuli.

In the second session (approx. 30 min), participants first reported demographic information about their age, country, education, gender, handiness, income, current job and mother tongue. Afterwards, they completed two tasks. First, they performed the *estimate-your-memory* task, in which they indicated how well they remembered each possible item from the *remember-and-decide* task of session 1 (similar to Mechera-Ostrovsky and Gluth, 2018). Second, they performed a gambling task (see Fig. 3) including risky and ambiguous gambles in the gain and the loss domain. We included this task to test whether participants who exhibit a stronger reduction of the memory bias in the loss compared to the gain domain would also show a stronger reflection effect in decision under risk or ambiguity. We adapted a task from previous studies (Levy, Snell, Nelson, Rustichini, & Glimcher, 2010; Tymula, Rosenberg Belmaker, Ruderman, Glimcher, & Levy, 2013), in which participants made binary decisions between a sure gain/loss of a small amount of money (in our case ± 5 CHF) and a risky or ambiguous gamble of a larger gain/loss amount. More specifically, participants could either choose ± 5 CHF for sure or an amount between ± 6 and ± 30 CHF with a given probability. During a trial, participants first saw a fixation cross for 1s, followed by the depiction of the safe amount and the lottery. They had 10s to indicate their choice by pressing either the Q (left choice) or P (right choice) button on a keyboard. Finally, a green feedback rectangle appeared around their chosen option for 1s. The gambles included six gain/loss amounts (±6 , 12, 16, 22, 26, 30 CHF). The risky trials had five winning probability levels (0.2, 0.35, 0.5, 0.65, 0.8). In the ambiguous trials, the five levels of ambiguity (0.2, 0.35, 0.5, 0.65, 0.8) were indicated by the area of a grey bar which prevented a glimpse on the underlying probabilities. Following previous work (Levy et al., 2010; Tymula et al., 2013), the grey bar covered the red (lottery probability) and the blue (safe option probability) parts to the same extent. Therefore, if an ambiguous trial was played at the end of the experiment, a random number between the lowest winning probability and the highest (area covered by the grey bar) was drawn. Then an outcome was drawn based on this randomly selected probability. We repeated each amount twice, thus resulting in a total of 240 trials [12 unique amounts x (5 probability levels + 5 ambiguity levels) x 2 repetitions]. Seven participants in the Money Experiment did 280 trials, because an older version of the experiment was used, in which 40 catch trials with one option stochastically dominating the other option (e.g., choice between 5 CHF for sure and 5 CHF with a probability of 80%) were included. These trials were excluded for analysis.

Data exclusion

To ensure high data quality, we specified and preregistered a number of exclusion criteria. The following criteria were assessed separately for positive and negative trials: First, to reliably assess the memory bias with a hierarchical Bayesian logistic regression model, we determined a minimum number of 20 trials per participant, in which one option has been remembered, while the other has been forgotten. Moreover, we required a minimum number of 5 per observed choice (i.e., remembered option chosen; forgotten option chosen). These numbers were based on analyzes of pilot data. Additionally, we adopted a hierarchical Bayesian approach with mildly informed priors that is more robust compared to frequentist approaches (Gordóvil-Merino, Guàrdia-Olmos, & Peró-Cebollero, 2012; McNeish, 2016; Kruschke, 2010). Furthermore, participants who responded too fast (i.e., RT <200 ms) in $\geq$ 30% of trials of the decision task or in $\geq$ 30% of the gambles task were excluded (however, none of the participants had to be excluded for being too fast).

Data analysis

Memory bias estimation. The memory bias was assessed in a similar way as in our previous work (Gluth et al., 2015; Mechera-Ostrovsky & Gluth, 2018), but instead of maximum likelihood estimation we employed hierarchical Bayesian logistic regression analyses. Note that the hierarchical Bayesian framework allowed us to compare the group posterior distributions directly and provided us with an estimate of certainty (Wagenmakers et al., 2018). Moreover, it is especially recommended when the number of observations varies across participants, which is the case for our remembered-forgotten trials (McNeish, 2016). The memory bias analyses are based on trials (t) with one remembered and one forgotten option. The probability p_t to choose the remembered option over the forgotten option is given by

$$p_t = \frac{1}{1 + \exp^{-(\beta_0 + \beta_1 * x_t)}}, \tag{1}$$

where x_t refers to the standardized subjective value of the remembered option in trial t, and β_0 and β_1 refer to intercept and slope coefficients, respectively. The probability that the remembered item will be chosen is estimated by drawing from a Bernoulli distribution with success probability p_t:

$$y \sim \mathrm{Bern}(p_t), \tag{2}$$

Hierarchical priors for the two regression coefficients in the model (β_0 and β_1) and hyper priors are specified as follows:

$$\begin{aligned} \mu_\beta &\sim N(0, 1) \\ \sigma_\beta &\sim \mathrm{HalfCauchy}(5) \\ \beta &\sim N(\mu_\beta, \sigma_\beta) \end{aligned} \tag{3}$$

For each coefficient (intercept and slope) the mean μ_β was drawn from a normal distribution, and the standard deviation σ_β was drawn from a Half-Cauchy distribution.

We specified the prior distributions based on the developers' recommendations of the used estimation package.

The slope of the logistic function β_1 specifies to what extent decisions depend on the value of the remembered option, the intercept β_0 quantifies the overall tendency to prefer remembered or forgotten options, and thus the memory bias. Notably, in our previous work we introduced a corrected version of the memory bias which controls for the possibility that participants remember high-value options better than low-value options (Mechera-Ostrovsky & Gluth, 2018). This correction consists of subtracting the average value of all forgotten options from the value of the remembered option x_i. In the present study, we also implemented this correction when quantifying the memory bias.

The statistical test for an influence of memory on choice was based on the group posterior samples of the intercept parameter β_0. If the 90% HDI of the distribution did not overlap with 0, we inferred a significant memory bias (a positive memory bias if the distribution lies to the right of 0, a negative memory bias if the distribution lies to the left of 0)[1]. Moreover, to test for the difference between gains and losses, we tested for an overlap with 0 as before for the estimated difference parameter. As a sanity check that participants take the value of remembered options into account when choosing between a remembered and a forgotten option, we also checked that the posterior distribution of the mean slope parameter β_1 was larger than 0 in all conditions (gains and losses, Images and Money Experiments) by testing whether the 90% HDI (highest density interval) did not overlap with 0.

[1] At this point, we deviated from our preregistration protocol, in which we announced to fit the logistic regression model and perform frequentist tests on the means of the individual posterior distributions. Such a two-step procedure can lead to inflated results in favor of the alternative hypothesis (Boehm, Marsman, Matzke, & Wagenmakers, 2018). Therefore, here we report the fully Bayesian tests only. For completeness the SOM includes the (invalid) two-step approach as well as a (purely frequentist) random-effects regression analysis [as in Gluth et al. (2015), Mechera-Ostrovsky and Gluth (2018)].

Risk and ambiguity attitudes assessment. To assess participants' risk and ambiguity parameters we used an adapted version of a previously proposed model (Levy et al., 2010; FeldmanHall, Glimcher, Baker, & Phelps, 2016). According to this model, the subjective value of an option is given by:

$$SV = (p - \beta * \frac{A}{2}) * v^{\alpha} \tag{4}$$

where p is the probability of the gain/loss amount of the lottery, A indicates the level of ambiguity, v is the gain/loss amount, α the individual risk attitude and β the individual ambiguity attitude. Note that a loss aversion parameter is not included, because the task does not contain mixed lotteries, and risk and ambiguity attitudes are estimated separately for gains and losses. The probability of choosing the lottery is given by a logit function (as in Equation 1) with the intercept being fixed at 0.

Notably, we adopted a "bug fix" (Stewart, Scheibehenne, & Pachur, 2018) that ensures commensurability of the sensitivity parameter γ across different risk preferences. Without this bug fix the risk parameter α trades off with the sensitivity parameter, because the risk parameter determines the range of possible values (e.g., the range is much larger if α=2 compared to α=1/2). This problem is solved by transforming the subjective value SV as follows:

$$\begin{aligned} SV &= SV^{1/\alpha} \text{ for } SV \geq 0 \\ SV &= -(\mid SV \mid)^{1/\alpha} \text{ for } SV < 0 \end{aligned} \tag{5}$$

Similar to the logistic regression described above, the model prior and hyper-priors were specified as follows:

$$\begin{aligned} \mu_\beta &\sim N(0, 1) \\ \sigma_\beta &\sim \text{Inv-Gamma}(3, 0.5) \\ \beta &\sim N(\mu_\beta, \sigma_\beta) \end{aligned} \tag{6}$$

Risk and Ambiguity attitudes were estimated separately for gains and losses, and for the two experiments (Images and Money).

To test our predictions that the memory bias is related to ambiguity (more so than risk) attitudes, we estimated a Bayesian linear regression predicting the difference of the memory bias between gains and losses with the following three predictors: i) the experiment (Images and Money), ii) the difference in risk attitudes (gains - losses), and iii) the ambiguity attitudes (gains - losses). The priors of the glm module were defined as follows: intercept and regressors $\sim$ Normal(mu=0, sd=1), standard deviation $\sim$ Half-Cauchy(10). As exploratory analyses, we also correlated the mean estimates for the memory bias with the mean estimates of the risk and ambiguity attitudes (separately for gains and losses). Thereto, we used an uniform prior between -1 and 1 for the correlation coefficient r. To calculate the Bayes Factors (BF) we compared our posterior samples to samples from the prior distribution. BFs indicate the evidence provided by the data in favor of an hypothesis. We were interested in the evidence in favor of the Null hypothesis denoted as BF_{01}. A BF of 1 indicates that both hypotheses (Null and Alternative) predict the data equally well(van Doorn et al., 2019).Generally, a $BF \geq 10$ indicates strong evidence.

Bayesian parameter estimation details. Bayesian models for estimating the memory bias were implemented using the pymc3 Python library. We sampled four chains, with 10000 samples each (5000 tuning samples), using the no-U-turn sampler (NUTS). Bayesian models estimating the risk/ambiguity attitudes were implemented using the rstan R library. We sampled two chains, with 5000 samples each (2000 tuning samples), using NUTS. Convergence was diagnosed using the Gelman–Rubin criterion ($|\mathrm{R}\hat{} - 1| < 0.05$) for all analyses. Effects were declared as statistically meaningful either when the 90% HDI excluded zero or when 90% of the posterior density was above (below) zero. In the latter case, we also reported the proportion of the posterior mass above (below) zero, directly indicating the posterior probability of the effect being larger (smaller) than zero. (Kruschke, 2014).Bayesian model estimation for the assessment of the memory bias, the Bayesian linear regressions and Bayesian correlations for the relationship of the memory bias and the risk/ambiguity attitudes were performed in Python v3.6.9, using the NumPy v1.17.2, Pandas v0.25.1, Theano v1.0.4 and PyMC3 v3.7 libraries. All other analyses (frequenstist tests in the SOM, descriptives, figures and data-preprocessing, Bayesian risk/ambiguity attitude estimation) were performed in R v3.6.1, using additionally the libraries psych v1.8.12, ggplot2 v3.2.1, rstan v2.19.2 and bayestestR v0.4.0.

Results

The memory bias in preferential choice in gains and losses. Our central hypothesis was that the memory bias, that is, the tendency to prefer remembered over forgotten options, is more positive in the gain as compared to the loss domain. To test this hypothesis we performed hierarchical Bayesian logistic regression analyses for trials with one remembered and one forgotten option, and predicted the choice of the remembered option based on its value. Before testing for the memory bias, however, we checked whether participants were more likely to choose remembered options of higher subjective value. In line with this, we found that the HDI of the group-level posterior distributions of the logistic slope coefficient was positive and did not overlap with 0 in all conditions (Images Experiment, gains: M = 0.47, SD = 0.10, 90% HDI = [0.30,0.64], losses: M = 0.33, SD = 0.09, 90% HDI:[0.17,0.47, difference gains - losses: M = 0.14, SD = 0.14, 90% HDI = [-0.07,0.38]; Money Experiment, gains: M = 0.88, SD = 0.13, 90% HDI = [0.62,1.10], losses: M = 0.76, SD = 0.15, 90% HDI = [0.53,1.01], difference gains - losses: M = 0.12, SD = 0.20, 90% HDI = [-0.16,0.44]).

More importantly, to test for a more positive memory bias in gains compared to losses we contrasted the group-level posterior distributions of the logistic intercept coefficient between gains and losses. In both experiments, we found that the memory bias was more positive in the gain than in the loss domain, and that the overlap of the two posterior distributions was less than 5% (i.e., 0.47% in the Images Experiment and 4.96% in the Money Experiment; Images Experiment difference gains - losses: M = 0.34, SD = 0.13, 90% HDI = [0.13,0.56]), Money Experiment difference gains - losses: M = 0.15, SD =

0.09, 90% HDI = [0.01,0.31], which confirmed our hypothesis (Fig. 2). In addition, we tested whether the memory bias was positive in the gain domain and negative in the loss domain (in absolute terms). Descriptively, this was the case in both experiments, but only in the gain condition of the Images Experiment the 90% HDI did not overlap with 0 (Images Experiment, gains: M = 0.24, SD = 0.09, 90% HDI = [0.10,0.39], losses: M =-0.10, SD = 0.09, 90% HDI:[-0.25,0.06], ; Money Experiment, gains: M = 0.09, SD = 0.07, 90% HDI = [-0.03,0.20], losses: M = -0.06, SD = 0.06, 90% HDI = [-0.16,0.03]).

Taken together, participants in both experiments indeed preferred remembered over forgotten options in the gain domain but forgotten over remembered options in the loss domain, with the difference between gains and losses being credible.

Testing an association of the memory bias with risk and ambiguity aversion. In addition to our main hypothesis, we tested whether the difference of the memory bias in gains vs. losses is correlated with the difference in risk or ambiguity aversion in gains vs. losses. We predicted to find an association with ambiguity but not risk, because choosing a less-remembered option whose consequences are uncertain should be conceptually similar to choosing a lottery option whose probabilities are not even known. To test this hypothesis, participants in both experiments performed an additional task, in which they made binary decisions between a sure gain or loss and either a risky or ambiguous lottery (Fig. 2a and b). We modeled their decisions to derive individual risk and ambiguity attitudes separately for gains and losses in a hierarchical Bayesian framework. Then, we linked the individual risk and ambiguity attitude parameters (individuals' mean estimates) with the memory bias parameter using a combined Bayesian multiple linear regression analysis for both experiments. We found that neither the ambiguity nor the risk attitudes as measured by the gambles task were related to the memory bias, as the 90% HDI included 0. However, we observed an effect of experiment, as the size of the memory bias differed if monetary rewards are used or images (intercept: M = 0.34, SD = 0.07, 90% HDI = [0.23,0.47], Experiment (money as reference): M = -0.19,

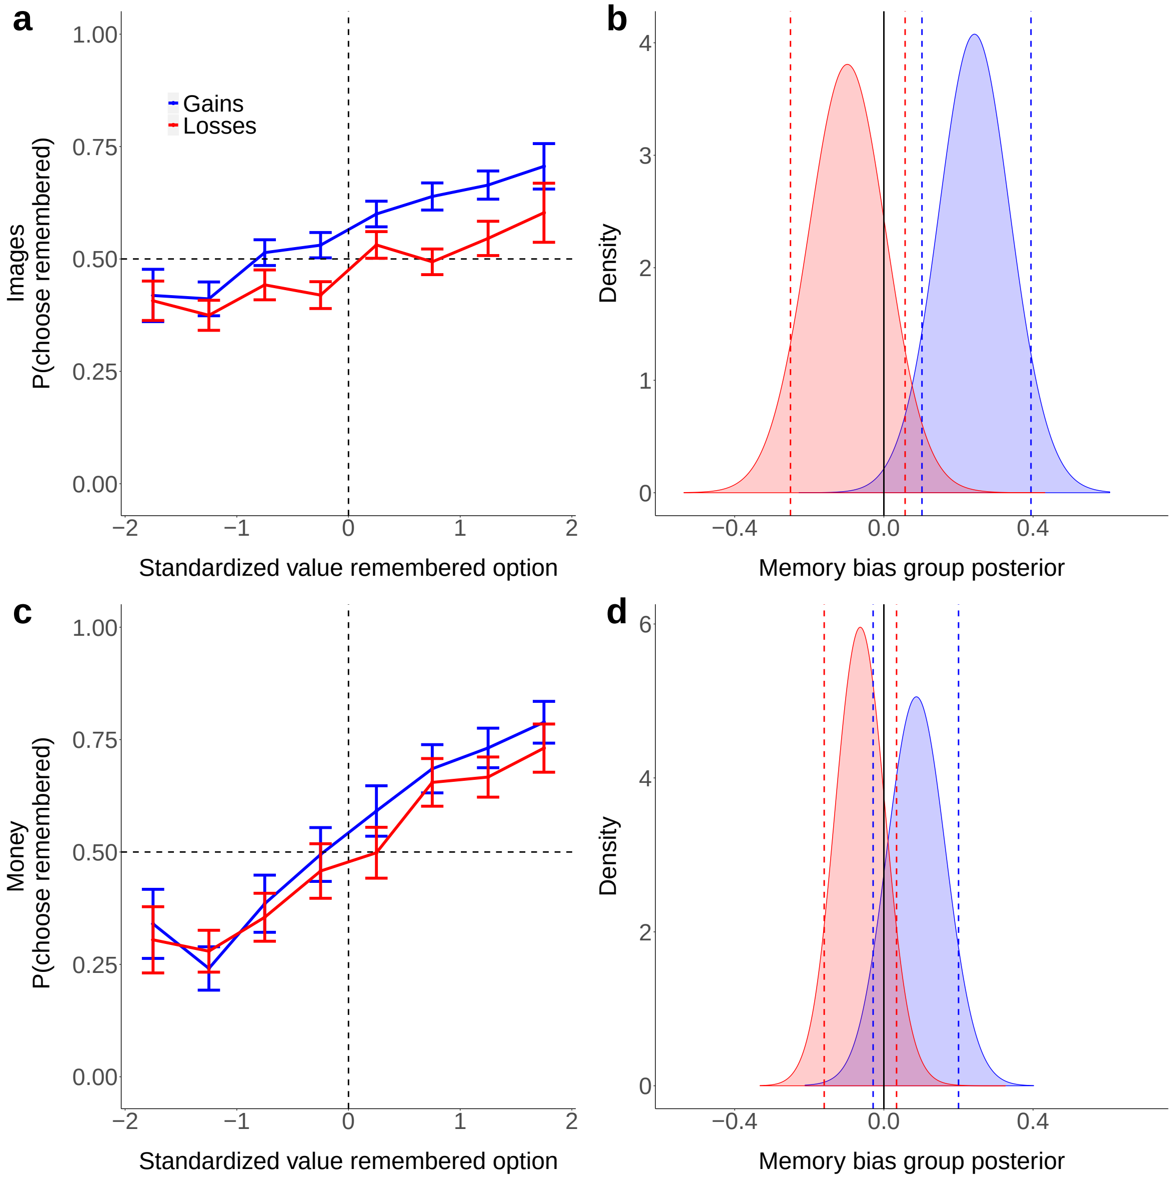

Fig. 2. Memory bias for gains and losses in both experiments. The upper panels refer to the Images Experiment, the lower panels refer to the Money Experiment. The left panels depict the probability to choose remembered over forgotten options as a function of the remembered option's subjective value. The right panels depict the posterior samples of the group-level intercept coefficient of the logistic regression, that is, the memory bias parameter. Error bars in the left panels indicate 95% CI. In the right panels, the dashed lines indicate the 90% HDI of the posterior distribution.

SD = 0.10, 90% HDI = [-0.36,-0.04], effect of risk: M = 0.00, SD = 0.09, 90% HDI = [-0.16,1.15], effect of ambiguity: M = 0.01, SD = 0.08, 90% HDI = [-0.11,0.14]). To quantify the evidence in favor of the Null, we also computed Bayes Factors, which suggest that there is strong evidence in favor of the Null for an effect of ambiguity (BF_{01} = 13.01) and strong evidence in favor of the Null for an effect of risk (BF_{01} = 10.86).

As additional exploratory analyses, we correlated the gain-loss difference in the memory bias with the gain-loss difference in risk and ambiguity attitudes. Results indicate that neither the risk attitude nor the ambiguity attitude as measured by the gambles task were related to the memory bias (Fig. 3c and d). More specifically, we calculated the correlations separate per experiment (money or images), finding no credible correlation ($r_{\text{images,risk}}$: M = -0.05, SD = 0.17, 90% HDI = [-0.32,0.23], BF_{01} = 4.18; $r_{\text{money,risk}}$: M = 0.14, SD = 0.16, 90% HDI = [-0.13,0.41], BF_{01} = 4.21; $r_{\text{images,ambiguity}}$: M = -0.04, SD = 0.17, 90% HDI = [-0.31,0.24], BF_{01} = 3.00; $r_{\text{money,ambiguity}}$: M = 0.22, SD = 0.16, 90% HDI = [-0.04,0.48], BF_{01} = 1.87).

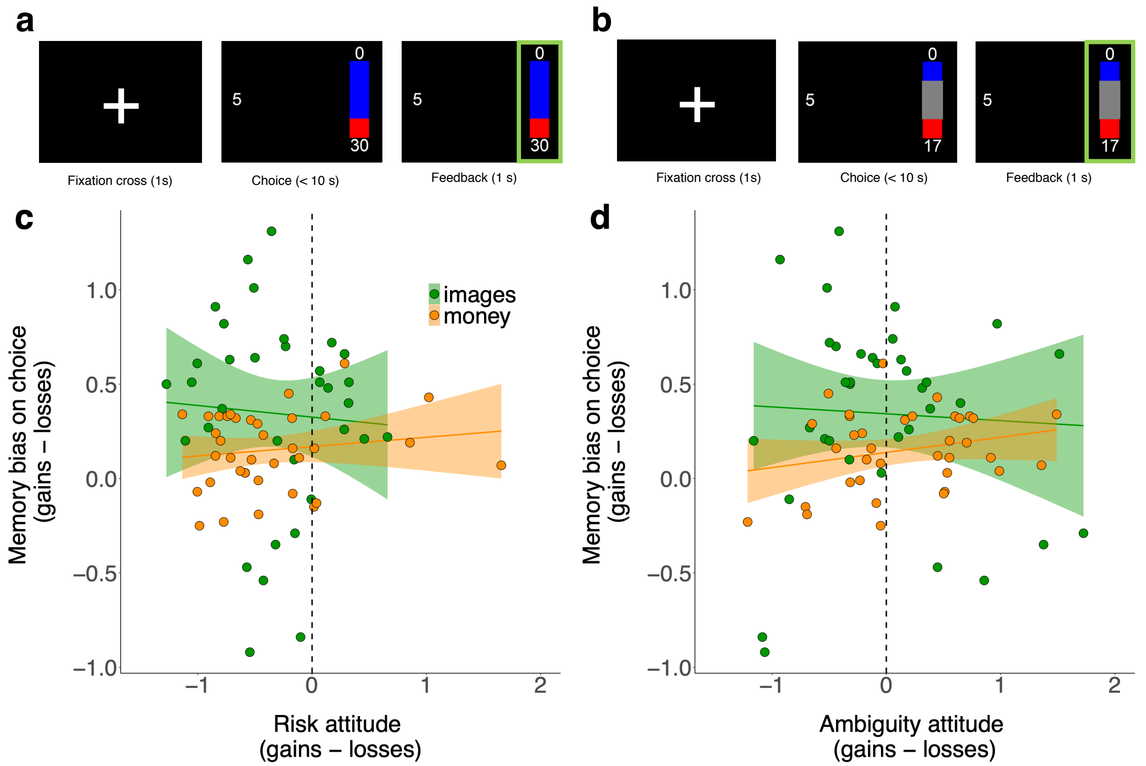

Fig. 3. Ambiguity and risk attitudes. In an additional lottery task, participants made binary decisions between a sure gain or loss and a risky (a) or ambiguous (b) lottery. The colored areas indicate the probability of the upper and lower amounts of the lottery. In case of ambiguous options, parts of the probability information are occluded. The gain-loss difference in the memory bias was not related to the gain-loss difference in risk (c) or ambiguity (d). Regression lines are added separately per experiment with their 95% CI.

Discussion

In the current preregistered study, we investigated an analogy between decisions from memory and decisions under uncertainty. More specifically, we tested whether the memory bias on preferential choice underlies characteristics of the well-known reflection effect (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981; Kahn & Sarin, 1988; Viscusi & Magat, 1992). If so, it should be reduced and possibly even reversed in the loss domain, meaning that people should prefer less-remembered over better-remembered options of negative subjective value. We conducted two experiments in which participants made preferential choices from memory between images and money amounts. Both experiments were carried out within the gain and loss domain. In both experiments, we observed that participants preferred better-remembered options in the gain domain but less-remembered options in the loss domain, with the gain-loss difference being credible. These results confirm our hypothesis that the memory bias shares characteristics with decisions under uncertainty.

By drawing a link between memory and uncertainty, our work connects two hitherto separate branches of decision-making research. It suggests that the uncertainty entailed in weak memories influence our choice behavior. Importantly, this connection offers several new avenues for future research. First, it will be important to further specify the nature of memory-induced uncertainty in more detail. Along this line, we speculate that the strength of memory for an option could be conceptualized as the probability weight assigned to it. Thus, a parametric effect of memory strength could exhibit a similar profile as the probability weighting function of prospect theory (Tversky & Kahneman, 1992) and lead to similar effects on behavior. Among such effects are the certainty and the possibility effect, according to which the subjective weighting of sure (i.e., 100%) and impossible (i.e., 0%) events are exceptionally larger/smaller than those of almost sure (e.g., 99%) and almost impossible (e.g., 1%) events. If memory strength exhibits a similar weighting profile, then remembering an option "for sure" (i.e., in all its episodic details) and not remembering an option at all should have exceptionally strong influences on our decisions. Second, the link between memory and uncertainty could stimulate research on the impact of inter-individual differences in memory abilities on decision making. For example, the fact that episodic memory shows a considerable decline over the lifespan (Nyberg, Lövdén, Riklund, Lindenberger, & Bäckman, 2012) should have important implications for older adults' attitudes toward uncertainty, at least with respect to those decisions that rely heavily on memory retrieval. Third, it will be critical to test whether our notion of a memory-uncertainty link can be supported by neuroscientific data. More specifically, neuroimaging research on decision under risk and uncertainty suggest a brain circuitry comprising the amygdala, the orbitofrontal cortex, and the dorsomedial prefrontal cortex (dmPFC) being involved in ambiguous choices (Hsu, Bhatt, Adolphs, Tranel, & Camerer, 2005; Huettel, Stowe, Gordon, Warner, & Platt, 2006), as well as the dmPFC and the anterior insula (aIns) being critical to risky choices (Morriss, Gell, & van Reekum, 2019; Mohr, Biele, & Heekeren, 2010). The aIns is also central to the processing of aversive stimuli (Nitschke, Sarinopoulos, Mackiewicz, Schaefer, & Davidson, 2006). Therefore, we assert that these areas should also be involved in memory-based decisions, possibly as a (negative) function of the vividness with which the chosen option is remembered.

Contrary to our prediction, the gain-loss difference in the memory bias was not related to the corresponding difference in participants' ambiguity (or risk) attitudes. We discuss three possible explanations for this null finding. First, it could be due to a lack of statistical power. This notion is partially supported by the comparatively low Bayes Factors in favor of the Null hypothesis (which were all below 10 when computing the correlations, thus never suggesting strong evidence). Second, the null result may relate to the finding that behavioral risk measures appear to have a low test-retest reliability (Frey, Pedroni, Mata, Rieskamp, & Hertwig, 2017), rendering them less suitable for studying inter-individual differences. Hence, it might be that an actual relationship between the reflection effect in memory-based decisions and the reflection effect in lottery decisions was concealed by the poor reliability of the later (and possibly of the former as well, since we have not assessed the test-retest reliability of the memory bias, yet). In this light, future studies may consider adding self-report measures of risk and ambiguity, as these measures appear to have higher reliability. Third, it is conceivable that uncertainty induced by poor memories of choice options and uncertainty induced by risk and ambiguity (i.e., known and unknown probabilities of outcomes) are only weakly related to each other. Notably, previous research has shown that risk attitudes are indeed highly domain-specific (Weber, Blais, & Betz, 2002; Blais & Weber, 2006), and behavioral measures do not only suffer from low reliability but also appear to exhibit low convergent validity (Frey et al., 2017). Therefore, even though the finding of a reflected memory bias supports our notion that weak memories induce a feeling of uncertainty, this form of uncertainty may be distinct from the uncertainty induced by not knowing whether a potential monetary amount will be paid out. Along all these lines, it is interesting to note that we found positive (albeit not significant) correlations between the memory bias and participants' risk and ambiguity attitudes in the Money Experiment but not in the Images Experiment. We speculate that this may reflect the similarity of choosing between monetary amounts retrieved from memory and of choosing between (uncertain) monetary rewards in the gambles task - a similarity not given in the Images Experiment. Certainly, a comprehensive understanding of the exact nature of memory-induced uncertainty in decision making requires more research efforts in the future, and possibly testing a larger sample.

Importantly, we do not consider uncertainty to be the sole driver of the memory bias on preferential choice. Our previous work showed that, in the gain domain, people believe to remember high-value options better than low-value options, and that the strength of this subjective belief was associated with the strength of the memory bias (Mechera-Ostrovsky & Gluth, 2018). Remarkably, in the current study, we found that not only participants' preferences but also their beliefs were inverted in the loss domain. That is, participants believed to remember strongly negative items better than weakly negative items (see SOM). Yet, after taking these value- and domain-dependent beliefs into account, the gain-loss difference of the memory bias remained significant (see SOM). Thus, the influence of memory on decisions appears to be multifaceted and to depend on both, what we infer about poorly remembered choice options (belief) and how we feel about choosing such options (uncertainty).

In sum, our two experiments showed that the influence of memory on preferential decisions generalizes to different types of choice options and exhibits a striking parallel to decisions from uncertainty: In the gain domain, people prefer better-remembered items, but in the loss domain they tend to prefer less-remembered options. We take this finding as evidence for a conceptual similarity between choosing poorly remembered options and choosing options with uncertain outcomes, thus connecting two different branches of decision-making research. Further research that should include neuroimaging and computational modeling approaches will be required to develop a comprehensive theory of the interplay between memory, uncertainty and preferential choice.

Supplemental Material

Additional supporting information can be found in the Supporting Online Material that accompany this manuscript.

Open Practices

All data and data analysis files have been made publicly available via the Open Science Framework (OSF) and can be accessed at https://osf.io/x935r/. The experiments were preregistered at the OSF (https://osf.io/eumj5/).

References

Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H., Reid, A., Verfaellie, M., … Shohamy, D. (2019). The hippocampus supports deliberation during value-based decisions. *eLife*, *8*, 1–28. doi:10.7554/elife.46080

Blais, A.-r. & Weber, E. U. (2006). A Domain-Specific Risk-Taking (DOSPERT) scale for adult populations. *Judgment and Decision making*, *1*(1), 33–47.

Boehm, U., Marsman, M., Matzke, D., & Wagenmakers, E.-J. (2018). On the importance of avoiding shortcuts in applying cognitive models to hierarchical data. *Behavior Research Methods*, *50*(4), 1614–1631. doi:10.3758/s13428-018-1054-3

FeldmanHall, O., Glimcher, P., Baker, A. L., & Phelps, E. A. (2016). Emotion and decision-making under uncertainty: Physiological arousal predicts increased gambling during ambiguity but not risk. *Journal of Experimental Psychology: General*, *145*(10), 1255–1262. doi:10.1037/xge0000205. arXiv: 15334406

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, *3*(10), e1701381. doi:10.1126/sciadv.1701381

Gershman, S. J. & Daw, N. D. (2017). Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework. *Annual Review of Psychology*, *68*(1), 101–128. doi:10.1146/annurev-psych-122414-033625

Gluth, S., Sommer, T., Rieskamp, J., & Büchel, C. (2015). Effective Connectivity between Hippocampus and Ventromedial Prefrontal Cortex Controls Preferential Choices from Memory. *Neuron*, *86*(4), 1078–1090. doi:10.1016/j.neuron.2015.04.023

Gordóvil-Merino, A., Guàrdia-Olmos, J., & Peró-Cebollero, M. (2012). Estimation of logistic regression models in small samples. A simulation study using a weakly informative default prior distribution. *Psicológica*, *33*, 345–361.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural Systems Responding to Degree of Uncertainty in human Decision-Making. *Science*, *310*(5754), 1680. doi:10.1126/science.1115327

Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., & Platt, M. L. (2006). Neural Signatures of Economic Preferences for Risk and Ambiguity. *Neuron*, *49*(5), 765–775. doi:10.1016/j.neuron.2006.01.024

Kahn, B. E. & Sarin, R. K. (1988). Modeling Ambiguity in Decisions Under Uncertainty. *Journal of Consumer Research*, *15*(2), 265. doi:10.1086/209163

Kahneman, D. & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263. doi:10.2307/1914185

Kruschke, J. K. (2010). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences*, *14* (7), 293–300. doi:10.1016/j.tics.2010.05.001. arXiv: arXiv:1011.1669v3

Kruschke, J. K. (2014). *Doing bayesian data analysis: A tutorial with r, jags, and stan.* Academic Press.

Kurdi, B., Lozano, S., & Banaji, M. R. (2017). Introducing the Open Affective Standardized Image Set (OASIS). *Behavior Research Methods*, *49* (2), 457–470. doi:10.3758/s13428-016-0715-3

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural Representation of Subjective Value Under Risk and Ambiguity. *Journal of Neurophysiology*, *103* (2), 1036–1047. doi:10.1152/jn.00853.2009

McNeish, D. (2016). On Using Bayesian Methods to Address Small Sample Problems. *Structural Equation Modeling*, *23* (5), 750–773. doi:10.1080/10705511.2016.1186549

Mechera-Ostrovsky, T. & Gluth, S. (2018). Memory Beliefs Drive the Memory Bias on Value-based Decisions. *Scientific Reports*, *8* (1), 10592. doi:10.1038/s41598-018-28728-9

Mohr, P. N. C., Biele, G., & Heekeren, H. R. (2010). Neural Processing of Risk. *Journal of Neuroscience*, *30* (19), 6613–6619. doi:10.1523/JNEUROSCI.0003-10.2010

Morriss, J., Gell, M., & van Reekum, C. M. (2019). The uncertain brain: A co-ordinate based meta-analysis of the neural signatures supporting uncertainty during different contexts. *Neuroscience and Biobehavioral Reviews*, *96* (July 2018), 241–249. doi:10.1016/j.neubiorev.2018.12.013

Murty, V., Feldmanhall, O., Hunter, L. E., Phelps, E. A., & Davachi, L. (2016). Episodic memories predict adaptive value-based decision-making. *Journal of Experimental Psychology: General*, *145* (5), 548–558. doi:10.1037/xge0000158

Nitschke, J. B., Sarinopoulos, I., Mackiewicz, K. L., Schaefer, H. S., & Davidson, R. J. (2006). Functional neuroanatomy of aversion and its anticipation. *NeuroImage*, *29*(1), 106–116. doi:10.1016/j.neuroimage.2005.06.068

Nyberg, L., Lövdén, M., Riklund, K., Lindenberger, U., & Bäckman, L. (2012). Memory aging and brain maintenance. *Trends in Cognitive Sciences*, *16*(5), 292–305. doi:10.1016/j.tics.2012.04.005

Shadlen, M. N. & Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. *Neuron*, *90*(5), 927–939. doi:10.1016/j.neuron.2016.04.036

Stewart, N., Scheibehenne, B., & Pachur, T. (2018). Psychological parameters have units: A bug fix for stochastic prospect theory and other decision models (PREPRINT), 1–11.

Tversky, A. & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458. doi:10.1126/science.7455683

Tversky, A. & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. doi:10.1007/BF00122574

Tymula, A., Rosenberg Belmaker, L. A., Ruderman, L., Glimcher, P. W., & Levy, I. (2013). Like cognitive function, decision making across the life span shows profound age-related changes. *Proceedings of the National Academy of Sciences*, *110*(42), 17143–17148. doi:10.1073/pnas.1309909110

van Doorn, J., van den Bergh, D., Bohm, U., Dablander, F., Derks, K., Draws, T., … Wagenmakers, E.-J. (2019). The JASP Guidelines for Conducting and Reporting a Bayesian Analysis. *PsyArxiv Preprint*, (February), 0–31. doi:10.31234/osf.io/yqxfr. arXiv: osf.io/yqxfr

Viscusi, W. K. & Magat, W. A. (1992). Bayesian decisions with ambiguous belief aversion. *Journal of Risk and Uncertainty*, *5*(4), 371–387. doi:10.1007/BF00122576

Wagenmakers, E. J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., … Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin and Review*, *25*(1), 35–57. doi:10.3758/s13423-017-1343-3

Weber, E. U., Blais, A. R., & Betz, N. E. (2002). A Domain-specific Risk-attitude Scale: Measuring Risk Perceptions and Risk Behaviors. *Journal of Behavioral Decision Making*, *15*(4), 263–290. doi:10.1002/bdm.414

Weber, E. U. & Johnson, E. J. (2006). Constructing Preferences from Memory. *SSRN Electronic Journal*, 1–26. doi:10.2139/ssrn.1301075

Weilbächer, R. & Gluth, S. (2017). The Interplay of Hippocampus and Ventromedial Prefrontal Cortex in Memory-Based Decision Making. *Brain Sciences*, *7*(12), 4. doi:10.3390/brainsci7010004

Wimmer, G. E. & Büchel, C. (2016). Reactivation of reward-related patterns from single past episodes supports memory-based decision making. *The Journal of Neuroscience*, *36*(10), 2868–2880. doi:10.1523/JNEUROSCI.3433-15.2016