Social framing effects in decision making

Payam Piray[1], Roshan Cools[2,3] and Ivan Toni[2]

1. Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
2. Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, The Netherlands
3. Department of Psychiatry, Radboud University Medical Center, Nijmegen, The Netherlands
Corresponding author: ppiray@princeton.edu

Abstract

Human decisions are known to be strongly influenced by the manner in which options are presented, the "framing effect". Here, we ask whether decision-makers are also influenced by how advice from other knowledgeable agents are framed, a "social framing effect". Concretely, do students learn better from a teacher who often frames advice by emphasizing appetitive outcomes, or do they learn better from another teacher who usually emphasizes avoiding options that can be harmful to their progress? We study the computational and neural mechanisms by which framing of advice affect decision-making, social learning, and trust. We found that human participants are more likely to trust and follow an adviser who often uses an appetitive frame for advice compared with another one who often uses an aversive frame. This social framing effect is implemented through a modulation of the integrative abilities of the ventromedial prefrontal cortex. At the time of choice, this region combines information learned via personal experiences of reward with social information, but the combination differs depending on the social framing of advice. Personally-acquired information is weighted more strongly when dealing with an adviser who uses an aversive frame. The findings suggest that social advice is systematically incorporated into our decisions, while being affected by biases similar to those influencing individual value-based learning.

Introduction

Human decisions are often influenced by information derived from the behavior of other agents (Baldwin and Baird, 2001; Siegal and Varley, 2002; Saxe, 2006; Behrens et al., 2009; Klucharev et al., 2009; Biele et al., 2011; Izuma and Adolphs, 2013; Stolk et al., 2013; Diaconescu et al., 2014; Chung et al., 2015; Reiter et al., 2017; Chib et al., 2018; Konovalov et al., 2018). Previous work has identified when and how decision-makers learn from knowledgeable agents (Morgan et al., 2012; Heyes, 2016). For instance, decision-makers weight socially-acquired information over personally-acquired information when the latter is considered uncertain (Morgan et al., 2012; Toelch et al., 2014), and as a function of the trustworthiness that advisers have built through their behavior over time (Behrens et al., 2008; Biele et al., 2011). Human decision-makers are also greatly biased when decision outcomes are described to them in an appetitive or aversive frame, even when those outcomes are financially equivalent (Kahneman and Tversky, 1979; Kahneman et al., 1982; De Martino et al., 2006; Rangel et al., 2008). Accordingly, framing outcomes is used by agents to influence decision-makers (McKenzie and Nelson, 2003; Sher and McKenzie, 2006; Frith and Singer, 2008; Cukier et al., 2021). Here, we ask how framing an advice, a prototypical socially-acquired information, biases the weight allocated to the advice and/or to the adviser by decision-makers, and alters their learning behavior.

Computational algorithms that support learning the trustworthiness of others over time are hypothesized to be similar to those that support learning about consequences of our own actions, even though those algorithms might be implemented by different brain circuits (Behrens et al., 2008, 2009; Nicolle et al., 2012; Ruff and Fehr, 2014). For instance, the computation of socially- and personally-acquired information is implemented by partially segregated circuits centered around the temporal-parietal junction (TPJ), often associated with social learning; and the ventral striatum, involved in reward learning (McClure et al., 2003; O'Doherty et al., 2003; Saxe, 2006). The ventromedial prefrontal cortex (vmPFC), a region implicated in value-based decision making, has been involved in combining socially- and personally-acquired information, at the time of choice (Behrens et al., 2008, 2009; Ruff and Fehr, 2014). However, it remains unclear how the combination of personally- and socially-acquired information is influenced by advice offered in appetitive or aversive frames, in particular when the reliability of information is uncertain. One possibility is that appetitively- or aversively-framed advice influence the weight allocated to the advice over personally-acquired probabilistic information, thus modulating reliance on the advice at the time of choice. The vmPFC is responsive to both appetitive and aversive values across different modalities (Pessiglione and Delgado, 2015), and it combines personally- with socially-acquired

information about the relative value of options when implementing choices (Nicolle et al., 2012; Zaki et al., 2014). Another possibility, not mutually exclusive, is that appetitively or aversively-framed advice influences the weight allocated to the adviser, i.e. her estimated knowledge or reliability in disclosing an outcome.

In this study, we address these questions through a novel paradigm, inspired by previous work (Behrens et al., 2008; Cook et al., 2014) and designed to disentangle the relative importance of socially- and personally-acquired information in a decision. Socially-acquired information was operationalized in the form of advice on binary choices between options A and B. Advice was provided by two social peers (confederates), and were closely matched in fidelity, volatility, and content, but had different framings. One adviser often gave advice with an appetitive frame, such as "Choose A". The other adviser often gave advice with an aversive frame, such as "Don't choose B". The task drives participants to track the history of reward for their choices, as well as the history of advisers' fidelity, and integrate those two sources of information when making their choice. The framing of advice can affect participants' estimation of each adviser's fidelity, participants' weighting of the two types of advice in their choice, or both. Computational modeling techniques, such as Bayesian and reinforcement learning modeling, enable us to see the world from the decision maker's perspective and calculate the estimated trustworthiness of each adviser, as well as the expected reward value for each option. Combined with functional magnetic resonance imaging (fMRI), these models allowed us to identify brain regions tracking computational parameters required for performing this task (Daw and Doya, 2006; Cohen et al., 2017), and understand how framing of advice modulates those neuro-computational mechanisms. This enabled us to disentangle whether framing of advice influences estimation of knowledge of the two advisers, or their trustworthiness, or both, in regions involved in social learning and valuation of choice, such as TPJ and vmPFC, respectively.

Results

Experimental setup

Thirty participants (16 women, average age = 23.3 years) underwent fMRI while performing a decision-making task in which they chose between blue and green options to accumulate monetary rewards. On every trial, they also received an advice from one of two advisers about which option they should or should not choose. On every trial, the adviser could help the participant by revealing the correct option or could lie. Furthermore, the adviser framed their advice in either an appetitive manner ('choose blue') or an aversive manner ('don't choose green') while keeping the content of the advice the same. While

both “choose blue” and “don’t choose green” point to the same choice for the participant, they have different frames. The advisers were social peers (confederates, one marked in yellow, the other in purple) who participants met before the imaging session began, in a practice session. The practice session was designed to expose participants to the fact that both advisers would attempt to vary fidelity of their advice from trial to trial (i.e. to lie on some trials, but not systematically), in order to earn monetary rewards. However, unbeknownst to participants, during the experiment the advice from both confederates were generated using a computer program, with matched fidelity and volatility throughout the experiment. The difference between advice of the two advisers was only in the frame of their advice. One adviser gave appetitively framed advice in the majority of trials (appetitive adviser) and the other one mostly gave aversively framed advice (aversive adviser).

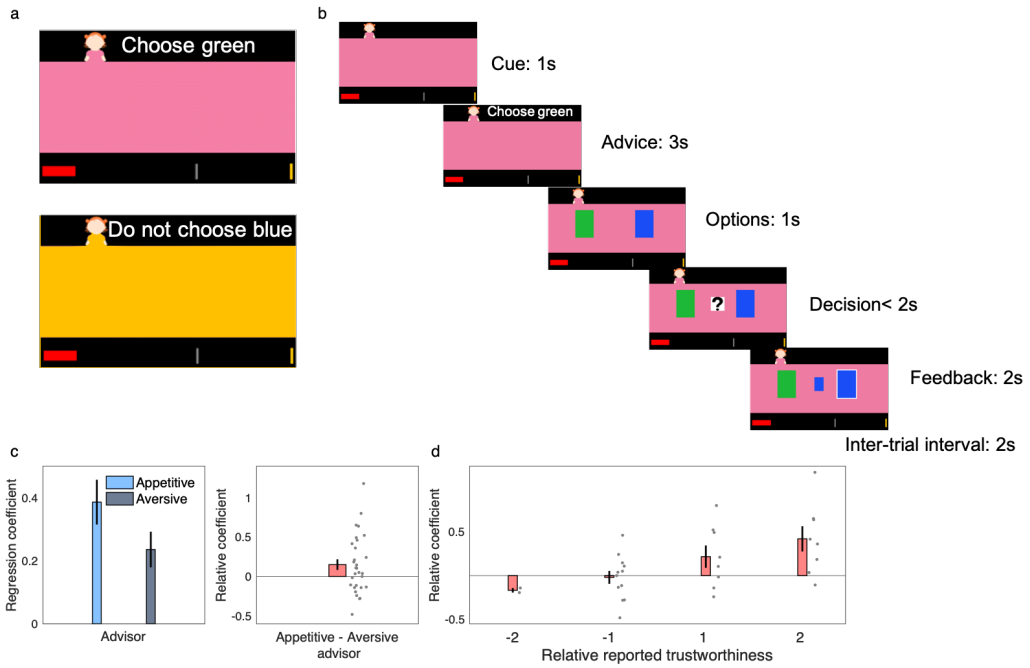

Figure 1. Experimental paradigm and behavioral finding. a) Participants received advice from one of two advisers cued with different colors. One adviser often gave appetitively framed advice, such as “Choose green” and the one often gave aversively framed advice, such as “Do not choose blue”. b) Timeline of the experimental task. On every trial, participants are presented with a color cue and then receive a written advice from the adviser. After choosing (button press) between two decision options, the correct choice (feedback) is revealed, and the total score (red bar at the bottom) increases with each correct response. c) Across all trials, participants followed the appetitive adviser more often than the aversive one, with a correspondingly positive average in the inter-participants distribution of the relative coefficient defined as the regression coefficient of appetitive- minus that of aversive- adviser. d) At the end of experiment, subjects were asked to rate trustworthiness of each adviser on a five-point scale. Across participants, relative reported trustworthiness was positively correlated with the relative

coefficients of following advisers during the experimental task. Error bars are standard errors of the mean.

Participants differentiated between the two advisers, and that difference had behavioral consequences. During the imaging session, participants' choices followed the appetitive adviser more frequently ($t(29)=2.23$, $p=0.034$; Figure 1c). After the imaging session, participants reported the trustworthiness of each adviser (explicitly, on a five-point scale), and their fairness (implicitly, through a variant of the Ultimatum Game (Sanfey et al., 2003; Tabibnia et al., 2008; Güroğlu et al., 2010)). Differential trustworthiness correlated with how closely participants' choices followed the appetitive adviser more than the aversive one (Spearman correlation: $r=0.54$, $p=0.002$, Figure 1d), and there was also a strong correlation between perceived trustworthiness and fairness of the advisers (Supplementary Figure 1).

Computational elements of choice under advice

Modeling techniques identified the computational mechanisms underlying the choice data. Optimal behavior in this task requires participants to combine social information (i.e. the advice, given the history of advisers' fidelity) with the history of correct options (i.e. which color has been more rewarding recently). Both sources of information were present on every trial outcome, but their statistics were independently manipulated throughout the experiment (Figure 2ab). Importantly, in addition to fidelity, volatility of the social information was closely matched for the two advisers. Therefore, optimal decision-making requires the decision maker to track i) the probability that the advice is correct, ii) the probability of one color being rewarded throughout the experiment, and iii) to combine the two probabilities regardless of the current adviser. We considered a Bayesian learning model, the volatile Kalman filter (VKF), an extension of the Kalman filter to volatile environments such as our task (Piray and Daw, 2020). This model assumes that the first- and second-order statistics of observations (e.g. fidelity and volatility of advice) change over time. Similar Bayesian models have been used to model choice data in probabilistic paradigms similar to the current task (Behrens et al., 2008; Diaconescu et al., 2014). The VKF predicts the probability of each option being correct given the sequence of previous outcomes, and the probability that the advice is correct (Figure 2ab). These two sources of information were then employed to generate the probability of each option by combining prior belief parameters and the predicted probabilities (likelihood) according to Bayes rule. We assumed that different belief *weight* parameters encode subject-specific beliefs about the appetitive and aversive advisers, respectively (w_{ap}, w_{av}). Higher values of the weight parameter were associated with higher probability of using social information for the trials

associated with the corresponding adviser. Probability of using reward information was always the inverse of using the social information.

The VKF model provided a better account of the choice data than a number of alternative models. We compared the VKF model with alternative simpler models (Figure 2c) using random effects Bayesian model comparison tools (Piray et al., 2019a). First, the VKF model was better than a simpler strategy, in which subjects followed the advisers blindly with different prior belief parameters, but regardless of the history of advisers' fidelity (Table 1 – Adviser's fidelity-blind model). Second, the VKF model was better than another cognitively-lean strategy, in which subjects followed the advice, possibly with different weights for appetitive and aversive frame, but regardless of the adviser (Table 1 – Adviser's frame-blind model). Finally, the VKF model was better than an alternative learning model, in which the learning strategy was simpler by tracking only the first order statistics regardless of volatility (Table 1 – Adviser's volatility-blind model). Statistics of fitted parameters of the winning model are reported in Supplementary Table 1.

	Model frequency	Exceedance probability
Volatile Kalman Filter (VKF)	0.8578	1
Advisers' fidelity-blind model	0.0334	0
Advisers' frame-blind model	0	0
Advisers' volatility-blind model	0.1088	0

Table 1. Bayesian model comparison results. The VKF model with different weights for the advisers was better than three simpler models: an advisers' fidelity-blind model in which subjects followed the advice (possibly with different weights for each adviser) regardless of their fidelity history; an advisers' frame-blind model in which subjects blindly followed the advice, possibly with different weights for each frame; and an adviser's volatility-blind model with a simpler learning strategy (constant learning rate) with different weights for the two advisers. Model frequency (ratio of participants explained by each model) and protected exceedance probability (probability that the model being the most frequent at the group level) are reported.

The belief parameters of the VKF model confirmed that participants' choices were more strongly influenced by the appetitive adviser. Individual values of the weight parameters, w_{ap} and w_{av}, indicated that participants, on average, relied less on aversive advisers: the weight parameter for using social information was significantly higher for the appetitive adviser than for the aversive one (t(29)=2.40, P=0.023; Figure 2c). Furthermore, whereas the weight probability of using reward history for making decisions for the appetitive adviser was not significantly different from 0.5 (t(29)=2.02, P=0.053), this probability was strongly higher than 0.5 for the aversive adviser (t(29)=3.83, P<0.001). This observation indicates that participants used their own experience more than the social information when the advice was given by the aversive adviser. Moreover, individual differences in the weight parameters predicted

reported trustworthiness of the two advisers. Specifically, the difference between the two belief weight parameters was correlated with the degree to which the appetitive adviser was reported to be more trustworthy than the aversive one (Spearman rank correlation: r=0.36, P=0.049; Figure 2d).

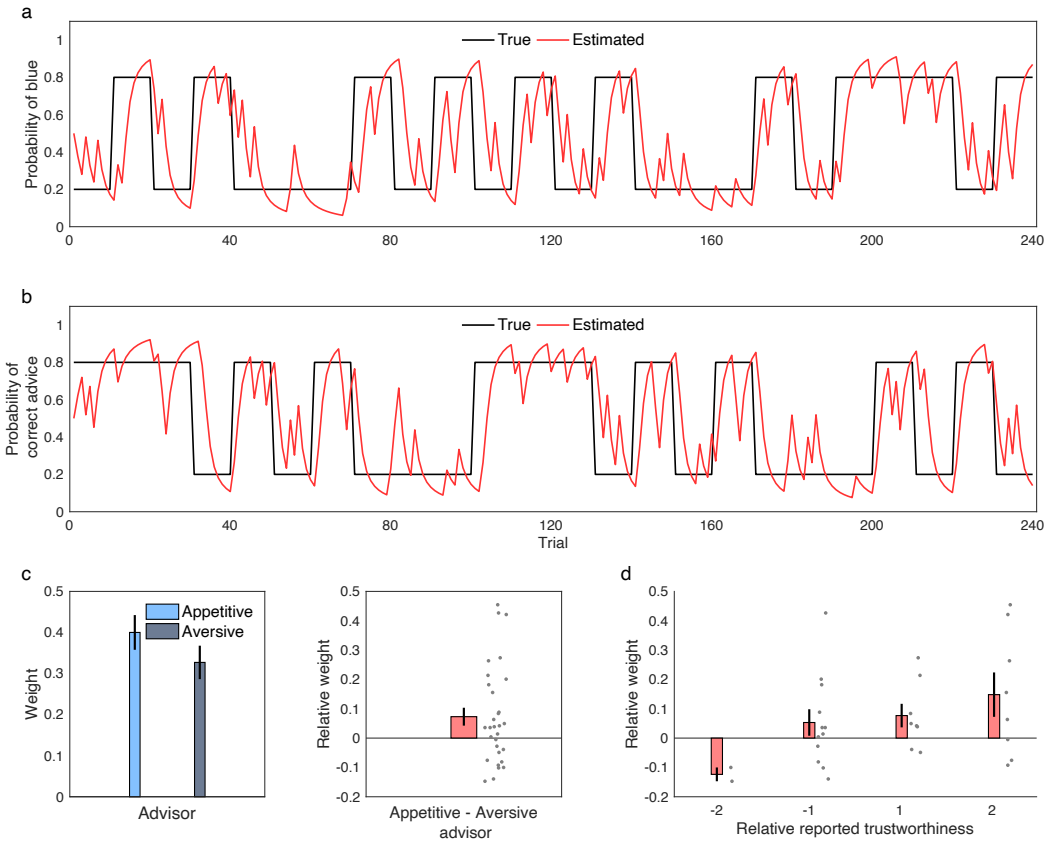

Figure 2. Computational findings. a, b) Probability schedules for the reward- and advice- based information. The black line shows the true probability of blue being correct (a) and the true probability of the advice being correct (b). The red line shows the VKF model's estimate of the probabilities. c) Weight parameters from the winning model, in which weights represent the degree that subjects employed social information on trials associated with the appetitive and aversive adviser. Participants significantly relied more on the appetitive than the aversive adviser. The distribution of the relative weight parameters (appetitive minus aversive) is also shown. d) Relative weight as a function of relative reported trustworthiness is plotted. Error bars are standard errors of the mean. See also Supplementary Figures 2-3.

A control analysis excluded that the different influence of the two advisers on participants' choices was a consequence of differences in learning the advisers' fidelity. We tested whether the association between the aversive advice ('do not choose green') and the recommended action (i.e. the blue one) might be harder to learn than the association between appetitive advice ('choose blue') and the recommended

action (i.e. the blue one). First, we performed a logistic regression analysis, similar to the analysis presented in Figure 1, with two additional regressors evaluating the impact of fidelity on the previous trial and its interaction with the identity of the adviser. This analysis revealed a highly significant effect ($t(29)=7.43$, $P<0.0001$) of fidelity on the previous trial on current choice (i.e. one-trial social learning) indicating that participants processed and used advisers' fidelity in their choice, but no significant interaction between this learning regressor and the identity of the adviser ($t(29)=-1.63$, $P=0.11$), indicating no difference in social learning between the two advisers. Second, a model-based analysis confirmed and extended these one-trial learning effects by using full models and Bayesian statistics. Specifically, two additional models were fitted to choice data to investigate whether subjects processed the history of the two advisers differently. In particular, these extended VKF models considered whether different amount of noise governed learning processes for the two advisers. Note that VKF contains two different noise parameters that encode different types of noise. When compared with the original VKF model, in which the same noise parameters govern learning for both advisers, the extended VKF models would capture the presence of any differences in the learning process between the two advisers (e.g. the possibility that participants viewed the aversive adviser as more noisy). Bayesian model comparison analysis revealed that the original VKF model with no additional parameter accounts better for choice data (model frequency = 0.94; protected exceedance probability = 1), confirming that participants did not process the history of the two advisers differently.

These computational analyses indicate that participants combined social advice and personal reward experience, relying less on aversively-framed advice to make decisions. This effect can theoretically be mediated by neural processes at the time of choice, or by neural processes at the time of receiving feedback, or a combination of both. Using fMRI, we isolated brain activity associated with computational parameters quantifying reliance on social and reward information, during choice and as a function of learning.

Choice-related effects in vmPFC

We focused on a ventromedial portion of the prefrontal cortex (vmPFC) that has been shown to combine social- and reward-related information at the time of choice in value-based decision-making (Behrens et al., 2008). We considered two decision-related regressors, time-locked to the presentation of the choice options. One regressor was based on social advice (social decision variable, SDV), and it was equal to p_t^s (estimated probability of advice being correct on trial t) and $1 - p_t^s$ on trials in which the subject chose to follow and not follow the advice, respectively (see Methods – Computational model). The other regressor

was based on personal experience of reward (reward-based decision variable, RDV), and it was equal to p_t^r (estimated probability of blue being correct on trial *t*) and $1 - p_t^r$ on trials in which the choice was the blue and the green option, respectively. Activity of the vmPFC was significantly correlated with a combination of both regressors (Behrens et al., 2008), consistent with its role in value-based decision making (Figure 3a; peak at x=–6, y=36, z=–18; voxel-level familywise small-volume corrected, k=713, t(29)=5.73, $P_{FWE\text{-}SVC}$<0.001; see also Supplementary Figure 2). Post-hoc tests revealed a stronger association with the reward-based information (RDV) when the aversive adviser was involved (Figure 3b; t(29)=2.45, P=0.021), but no significant difference on socially-acquired information (SDV) (t(29)=0.52, P=0.52). Furthermore, there was inter-participant variation in the extent to which the vmPFC signal reflected personally-acquired information when the aversive or the appetitive adviser was involved. Crucially, variation in vmPFC signal predicted variation in the weight parameter capturing reliance on reward-based probability in the computational analysis (Figure 3C; r=0.39, P=0.035). Overall, these results suggest that participants relied more on their own experience on trials involving the aversive adviser.

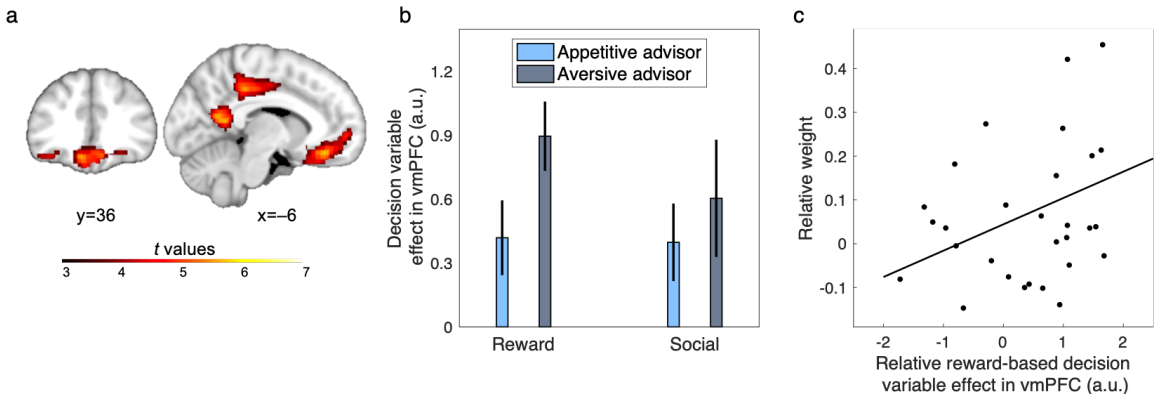

Figure 3. Decision-related activity in the vmPFC. A) Activation for the combination (mean contrast) of RDV and SDV at the time of decision making, when options were presented (thresholded at *P*<0.001 for illustration purposes). There is a significant effect in the vmPFC. B) Beta coefficients across all voxels of the VMPFC with uncorrected *P*<0.001. Decision-related activity of the vmPFC is significantly stronger for reward-based decision variable on trials associated with the aversive adviser. C) Inter-individual variability in behavior (differential use of reward-based information between appetitive and aversive advisers) was associated with neural variability in vmPFC (differential signal of personally-acquired information between appetitive and aversive advisers, as in panel B). Errorbars are standard errors of the mean. Unthresholded statistical maps are available at https://neurovault.org/collections/LINWGQZZ/. See also Supplementary Figure 4.

Outcome-related effects in TPJ

We focused on a caudal portion of the TPJ that has been shown to track social prediction error during learning (Saxe, 2006; Behrens et al., 2008; Carter et al., 2012; Diaconescu et al., 2014). We considered two learning-related regressors, time-locked to the presentation of the outcomes. One regressor was based on the social prediction error (SPE), and it was given by the difference between observed and expected fidelity of the advice, weighted by the corresponding estimated uncertainty. The other regressor was based on the reward-based prediction error (RPE), which was given by the difference between observed and expected reward, weighted by corresponding estimated uncertainty. This analysis revealed that the right TPJ was significantly correlated with the social prediction error (Figure 4a; peak at x=42, y=–58, z=28, voxel-level familywise small-volume corrected, k=27, t(29)=4.94, $P_{FWE\text{-}SVC}$=0.014). This effect was not significantly different for the appetitive and aversive advisers (across all voxels in the TPJ mask: t(29)=–1.25, p=0.22). These results concur with the behavioral and computational analyses, confirming that the TPJ, a brain region known to track social prediction error, followed the fidelity of both advisers equally well. In other words, these results suggest that observed behavioral differences in the weight given to the two advisers are not driven by differences in ability to track their fidelity. In fact, if anything, the right TPJ was more active for the aversive adviser. There was also a significant effect of reward prediction error in TPJ, but only in the left hemisphere (Figure 4b; peak at x=–64, y=–32, z=44, voxel-level familywise small-volume corrected, k=96, t(29)=5.06, $P_{FWE\text{-}SVC}$=0.006).

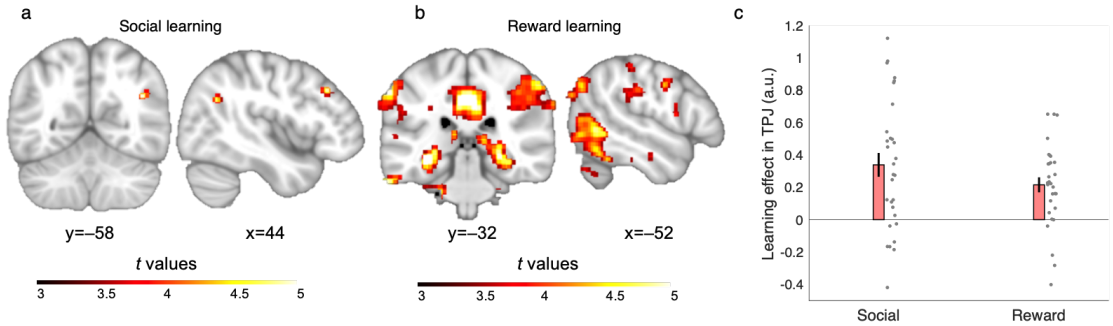

Figure 4. Learning activity in the TPJ. a) Activation for the SPE at the time of feedback, when it was revealed whether the advice was a lie or not. There is a significant effect in the TPJ. b) Activation for the RPE at the time of feedback. There is a significant effect in the TPJ, although this effect is in an anatomically distinct area compared with the social learning effect. Brain images are thresholded at *P*<0.001 uncorrected for illustration purposes. There was no significant difference between the two advisers regarding these effects. c) Beta coefficients across all voxels of the TPJ mask with uncorrected *P*<0.001 for both types of learning. Unthresholded statistical maps are available at https://neurovault.org/collections/LINWGQZZ/.

Discussion

Humans operate in a hyper-cooperative social environment, where individuals offer their knowledge to non-kins in the form of advice (Boyd and Richerson, 2009; Salali et al., 2016). Accordingly, we need computational and neural mechanisms to adaptively combine personal experience with socially mediated knowledge. Here, we identify neuro-computational characteristics of a bias in how pieces of advice are combined with personally-acquired knowledge. The bias leads participants to rely less on aversively- than appetitively- framed advice to make decisions, despite matched fidelity, volatility, and content of the advice. This social framing bias occurred at the time of choice, with a correspondingly stronger weight allocated to personal experience over aversively-framed advice when those sources of information are combined in the vmPFC. These findings suggest that social advice is systematically incorporated into the decision-making process. Fittingly, this effect was found in the vmPFC, a region known to combine volatile knowledge about social partners with personal knowledge (Stolk et al., 2015), and to guide current actions by updating the relative value of different models of the world (Nicolle et al., 2012).

Several studies in psychology, neuroscience and behavioral economics have shown that human decision making is influenced by how the decision problem is framed (Kahneman and Tversky, 1979; Kahneman et al., 1982; De Martino et al., 2006; Rangel et al., 2008). The framing effect has been often used as an example of how emotional valence might play a disruptive role in rational decision making. In that perspective, it could be argued that the current findings are a straightforward generalization of well-known Pavlovian biases in individual value-based learning (Guitart-Masip et al., 2012) to socially-mediated learning. However, that interpretation does not account for the observation that the social framing bias led participants to trust and perceive as fairer those advisers that provided predominantly appetitively-framed advice, without biasing participants' confidence in the knowledge of the advisers, and without down-weighting their influence. An alternative viewpoint suggests that socially-driven modulations of decision making might stem from appraising how an advice is communicatively embedded (McKenzie and Nelson, 2003; Sher and McKenzie, 2006; Frith and Singer, 2008). Namely, considering the communicative expectations generated by the advice might offer a more complex but comprehensive interpretation of the findings, accounting for social framing bias, matched confidence in advisers' knowledge, yet different trust in them. The two types of advice, while informationally matched, impose different cognitive demands on the participants. The aversively-framed advice imposes on the participant the need to resolve the negation and infer the suggested choice. Imposing this marginal but un-necessary cognitive demand on the recipient of the advice violates the cooperative principle of human communication (Grice, 1975),

and violates the expectation of additional pragmatic implications generated in the listener when that communicative principle is ignored.

An important aspect of our experimental paradigm was the fact that participant made binary choices. This was important because such a binary choice guaranteed that the content of information was matched for both frames (i.e. that “Choose green” is equivalent to “Do not choose blue”). However, real-life decisions often require choosing between more than two options, and in such environments, an aversively framed advice is less informative for decision making than an appetitive frame. In naturalistic settings, it could be argued that an appetitively-framed advice, on average, leads to faster and better decisions, possibly resulting in a systematic bias even when choosing between equivalent options. On the other hand, aversive frames are more likely to be consequential for survival, i.e. to signal threats, in particular when combined with nonverbal bodily signals. Our data are clearly consistent with the former possibility. It remains to be seen whether threat-related social signals might alter the social framing effect reported here, e.g. by influencing computations at the time of learning rather than choice, as observed in reinforcement learning studies (Piray et al., 2019b).

Methods

Experimental paradigm

The experimental paradigm and the cover story were taken from (Behrens et al., 2008). In this study, verbal advice was given to subjects by two confederates: one confederate often used an appetitive frame to give advice (in 80% of trials) and the other confederate often used an aversive frame to give advice.

Participants performed a probabilistic decision-making task in the scanner, in which they chose between blue and green tokens to accumulate reward (i.e. 1 or 0 point). Participants were instructed that on every trial, one of the colors leads to reward, while the other one leads to no reward. Therefore, it was clear for the participants that the probability of win given blue and green was inverse of each other, i.e. $p(\text{win}|\text{blue}) = 1 - p(\text{win}|\text{green})$. Participants' accumulated rewards were shown on the screen using a red score-bar, the length of which was proportional to the accumulated reward. A gold and a silver target were also shown on the screen and subjects were instructed that if they land their red score-bar in either of those targets, they receive 8 and 4 euros bonus, respectively. On each trial, the subject was first presented with a color cue (yellow or purple) and a cartoon cue indicating the adviser of this trial. Participants were instructed that each color is associated with one adviser throughout the experiment. One second after seeing the cue, the subject received a written advice (in Dutch), which could take 1 of 4 forms: "Choose blue", "Choose green", "Don't choose green" and "Don't choose blue". The advice remained on the screen for 3 seconds to ensure that subjects have sufficient time to process the advice. Next, two options (blue and green tokens) appeared on the screen for 1 second, during which subjects were not allowed to make their choice. This was followed by a question mark on the screen, indicating that subjects were allowed to make their choice (in less than 2 seconds). The correct option on that trial was then revealed on the screen .25 second after subjects made their choice and remained onscreen for 1.75 second. If the subject's choice was correct, a coin, and otherwise, a red cross appeared on the screen 1 second after the correct option was revealed. Trials were terminated by an inter-trial interval of 1.5-2.5 seconds (except three trials throughout the task with longer intervals of about 20 seconds).

Cover story

The cover story was similar to that of (Behrens et al., 2008). Participants were introduced to two actresses before the experiment began and all three of them were brought to a behavioral lab in which they received instructions and practiced the task together. In the behavioral lab, subjects and confederates were instructed that there were two different roles in the experiment: adviser and receiver. Participants

were always assigned as the receiver and confederates as advisers (using a counterfeit test). The advisers' task was to give trial-by-trial advice and their goal was to land the receiver's score bar in either a golden or a silver range to win 8 or 4 euro of bonus, respectively. Importantly, instructions about the advisers' task were given in front of participants. Critically, advisers were told that they need to mix helpful and harmful advice to maximize their bonus reward.

With examples, it was made clear that advisers' motive to give a helpful or harmful advice might change during the task given the receiver's score-bar location and the length and location of the golden and silver ranges. For example, if the receiver's score is far behind the golden range, the advisers are motivated to give her helpful advice to help her to reach to the golden range. However, as soon as the receiver is already in the golden range, advisers are motivated to give harmful advice to keep her in that range. The advisers were told that each trial will be randomly assigned to one of them, in which they should choose between four options as their advice: "Choose [correct]", "Not choose [incorrect], "Choose [incorrect]" and "Not choose [correct]". The advisers were also told that they don't know which color is correct or incorrect, but their advice will be translated according to the correct color on that specific trial. For example, if they select "Choose [correct]" and the blue is the correct color on that trial, their advice which appears on the receiver's screen would be "Choose blue". Therefore, it was made clear that advisers can choose not only the content of their advice (correct or incorrect), but also the frame of their advice (choose or not choose). They were told that they will perform the experiment in different labs and they don't see each other afterwards.

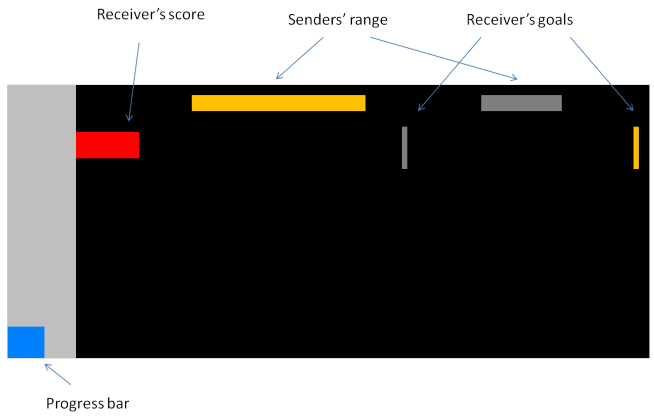

Advisers' task. The two advisers were instructed and practiced their task (in front of the subject). They were told that their task was to use their advice to facilitate that the receiver's red score-bar would land within their own golden range or a silver range, associated with 8 or 4 euros bonus. A progress bar signaled the number of trials passed. The advisers were told that their silver and golden ranges were the same and each trial would be randomly assigned to one of them.

The goal of the cover story was to make clear for the subject that 1) advisers have rational motives to give helpful or harmful advice in different phases of experiments (depending on the location and length of the golden and silver ranges); and 2) those motivations are the same for both advisers; 3) however, advisers' advice is independent of each other (as they don't see or communicate with each other). 4) Advisers' advice is also independent of the subject's choices, and, therefore, the subject can go against the advice or follow it as long as they want, 5) the advice is an extra and independent piece of information, because it is independent of the actual color being correct. Several examples were given during the instructions to make these points explicitly clear.

At the end of the experiment, subjects were asked to report which adviser they found to be more trustworthy. Next, they were asked to rate each adviser trustworthiness on a scale from 1 to 5. These data were used for correlation analysis in Figures 1-2.

Computational modeling

The main model of interest in this study is a hierarchical Bayesian learning model, called volatile Kalman filter (Piray and Daw, 2020), extends the celebrated Kalman filter model to volatile environments in which both the first- and second- order statistics of the environment evolve over time. In addition to the error-correcting rule of Kalman filter and other classical models (e.g. Rescorla-Wagner model) for learning observations, the VKF also learns volatility according to a second error-correcting rule. We used the binary version of this algorithm here as observations are binary. Specifically, if o_t is the observation on trial t, the VKF tracks the mean, m_t, of the hidden state of the environment using an uncertainty-weighted prediction error signal:

$$m_t = m_{t-1} + \sqrt{\sigma_{t-1} + v_{t-1}}(o_t - s(m_{t-1})) \qquad \text{Equation 1}$$

where σ_{t-1} is the estimated uncertainty on the previous trial, v_{t-1} is the estimated volatility on the previous trial, o_t is the current observation (binary: 0 or 1) and $s(.)$ is the sigmoid function:

$$(x) = \frac{1}{1 + e^{-x}} \qquad \text{Equation 2}$$

On every trial, VKF updates its uncertainty and volatility estimates:

$$k_t = (\sigma_{t-1} + v_{t-1})/(\sigma_{t-1} + v_{t-1} + \omega), \qquad \text{Equation 3}$$

$$\sigma_t = (1 - k_t)(\sigma_{t-1} + v_{t-1}), \qquad \text{Equation 4}$$

$$\sigma_{t-1,t} = (1 - k_t)\sigma_{t-1}, \qquad \text{Equation 5}$$

$$v_t = v_{t-1} + \lambda\left((m_t - m_{t-1})^2 + \sigma_{t-1} + \sigma_t - 2\sigma_{t-1,t} - v_{t-1}\right) \qquad \text{Equation 6}$$

The VKF model requires three constant parameters: $0 < \lambda < 1$ indicating volatility update rate, $v_0 > 0$ determining the initial value of volatility and $\omega > 0$ indicating observation noise. Note that volatility estimates are updated using an (second-order) error correcting rule. Also, note that if $\lambda = 0$, then the VKF is equivalent to a (binary) Kalman filter in which the process variance parameter is v_0. Both mean and uncertainty of the VKF were initiated at 0.

For every subject, the VKF algorithm was used twice. First, we used VKF to quantify value of each option based on reward history, m_t^r. In this case, observations were the color of the correct option (1=blue, 0=green). Thus, larger values of m_t^r indicate subject's estimated probability that the correct option is be the blue one. This is learnt using Equations 1-6 and based on three free parameters, λ_r, ω_r and v_0.

Second, we used the VKF to quantify value of each option based on the social experience and the advice, m_t^s. In this case, observations were the correctness of the advice (1=correct, 0=incorrect). Thus, larger values indicate subject's belief that the advice should be followed. This is also learnt using Equations 1-6 and based on λ_s, ω_s and v_0.

Next, a choice model was used to generate probability of choosing each option based on m_t^r and m_t^s. First, probability of the advice based on the social information was quantified using a sigmoid (softmax) function:

$$p_t^s = \sigma(\beta m_t^s + b)$$

where $\beta > 0$ is the decision noise parameter and b is the value-independent bias parameter for following the advice. Second, we computed the probability of the advice given the personal information:

$$p_t^r = \sigma(\beta \widehat{m}_t^r)$$

where $\widehat{m}_t^r = m_t^r$ if the advice is the blue option (i.e. "Choose blue" or "Don't choose green") and $\widehat{m}_t^r = -m_t^r$ if the advice is the green option.

Finally, these two probabilities were combined to construct probability of following the advice using a mixture model with two prior belief parameters (weights):

$$p_t = w_{ap} p_t^r + (1 - w_{ap}) p_t^s \text{ for the trials associated with the appetitive adviser}$$

$$p_t = w_{av} p_t^r + (1 - w_{av}) p_t^s \text{ for the trials associated with the aversive adviser}$$

where w_{ap} and w_{av} indicate weight parameters corresponding to the trials associated with the appetitive and aversive adviser, respectively. Both w_{ap} and w_{av} are constrained in the unit range and indicate the weight of using own experience. For example, if $w_{ap} = 1$, the subject relies entirely on her own experience and ignores the social information. Conversely, if $w_{ap} = 0$, the subject relies entirely on social information and ignores her own experience from past trials. Therefore, in total, this model has 9 free parameters: w_{ap}, w_{av}, β, b, v_0, ω^s, λ^s, λ^r and ω^r. This is model M1.

We considered three simpler models. First, we considered the possibility that subjects do not follow the intention of the confederate and follow advice (or go against advice) blindly, although with different weights for the advisers (M2, adviser-dependent blind). For learning from reward feedback, this model is identical to model M1. Therefore, this model has six parameters: w_{ap}, w_{av}, β, v_0, λ_r and ω_r. We also considered a similar model in which different weights were given to appetitively- and aversively- framed advice (M3, frame-dependent blind). Note that although this model has some overlap with the previous one, but is not identical to that one as advisers did not use their dominant frame in all trials. This account, therefore, considers the possibility that subjects blindly follow the advice given the frame, but not the advisers' intention (e.g. because unpacking the aversive frame is too difficult due to negation).

We also considered a simpler reinforcement learning strategy, which is based on only following the first-order statistics and fewer free parameters (M4, simpler learning strategy). Instead of equations 1-6, the reinforcement learning model learns value of two actions, a_1 and a_2 (e.g. blue and green) by updating the value of chosen action using the prediction error:

$$Q_t(a) = Q_{t-1}(a) + \kappa(r_t - Q_{t-1}(a))$$

where a is the chosen action on trial t, r_t is the reward on this trial, Q_t is the action-value on trial t and κ is the learning rate (a constant parameter bounded in the unit range). We then define $m_t = Q_t(a = 1) - Q_t(a = 2)$. We used this algorithm twice: first with reward-related actions ($a = 1$ for choosing blue and $a = 2$ for choosing green) and κ_r to obtain m_t^r; and then with advice-related actions ($a = 1$ for following the advice and $a = 2$ for not following) and κ_s to obtain m_t^s. The same choice model as M1 (containing a decision noise, a bias and two weight parameters) was then used to obtain probability of choices. Therefore, this model contains 6 free parameters: w_{ap}, w_{av}, β, b, κ_r, κ_s.

Finally, we performed a control analysis to confirm that subjects were able to follow both advisers equally well. For this analysis, we considered two models with additional parameters. These models were variants of M1. The first model (M1v1), assumed that different volatility learning parameters (that also give the degree of noise at the volatility level) govern learning for the two advisers (λ_s was different for the two advisers). The second model (M1v2) assumed that different noise parameters, ω_s, govern learning for the two advisers. As reported in the Results, model comparison analysis revealed that the original model (M1) with no additional parameter fitted choice data better than these models.

Model comparison and parameter estimation

We then used hierarchical Bayesian inference, HBI (Piray et al., 2019a), to compare these models. The HBI constructs subject-level priors based on the group-level statistics (empirical priors) and takes a random effects approach to model identity (different models might be the true model in different subjects). This method has been shown to be more robust than other methods for model selection (Piray et al., 2019a). This method fits parameters in the infinite real-space and transforms them to obtain *actual* parameters fed to the models. Therefore, appropriate transform functions were used for this purpose: the sigmoid function to transform parameters bounded in the unit range or with an upper bound and the exponential function to transform parameters bounded to positive value. In model M1, we assumed an upper bound of 10 for the observation noise parameters (ω_r and ω_s) to have robust estimates. The HBI indicated that the VKF model (M1) is the most likely model across population given the data (protected exceedance probability of M1 was not distinguishable from 1). The HBI algorithm is available online at https://github.com/payampiray/cbm.

For subject-level analyses of the parameters, we used individually fitted parameters of model M1. Unlike parameters estimated by the HBI that are regularized according to all subjects' data, the individually fitted parameters are independently estimated and therefore can be used in regular statistical tests. These parameters were obtained using a maximum-a-posteriori procedure, in which the prior mean and variance for all parameters (before transformation) were assumed to be 0 and 3.5, respectively. This initial variance is chosen to ensure that the parameters could vary in a wide range with no substantial effect of prior. Specifically, for any parameter bounded in the unit range, the log-effect of this prior is less than one chance-level choice (i.e log0.5) for any value of that parameter between 0.1 and 0.9.

For statistical analyses involving the weight parameters, we used the individually fitted parameters with appropriate parametric tests (i.e. t-test). Note that these parameters were fitted independently and

therefore it is valid to use classical tests for their analysis. For analyses involved the trustworthiness rating data, we used non-parametric Spearman correlation as these data were not normally distributed.

fMRI acquisition and preprocessing

Functional images were acquired using a 3T Skyra MRI system (Siemens), using a multiband sequence(Moeller et al., 2010) with acceleration factor 3 (Flip angle = 60 degrees, FOV = 64x64x33 voxels, voxel resolution of 3.5 mm isotropic, TR/TE = 1000/30 ms). Structural images were acquired using a T1-weighted MP-Rage sequence (Flip angle = 8 degrees, FOV = 192x256x256, voxel size 1x1x1, TR/TE = 2300/3.03ms). The task was presented and time-locked to fMRI data using Presentation (Neurobehavioural Systems, USA).

Data was preprocessed using FSL tools. Motion correction was applied using rigid body registration to the central volume and high pass temporal filtering was applied using a Gaussian-weighted running lines filter, with a 3dB cutoff of 100s. The T1-weighted images were spatially coregistered to functional images using linear (FLIRT) and non-linear tools (FNIRT). Noise effects in data were removed using FMRIB's ICA-based Xnoiseifier tool(Salimi-Khorshidi et al., 2014), which uses independent component analysis (ICA) and classification techniques to identify noise components in data (using a trained version based on HCP_hp2000 dataset and R=10). We also measured physiological signals for the purpose of denoising fMRI data using RETROICOR (Glover et al., 2000). Pulse was recorded using a pulse oximeter transducer affixed to the middle or ring finger of the non-dominant hand and respiration was measured using a respiration belt. Recorded signals were then used to generate covariate nuisance regressors (25 in total) using a fifth-order Fourier model of the cardiac and respiratory phase-related modulation of the BOLD signal (Birn et al., 2006; Shmueli et al., 2007; van Buuren et al., 2009).

fMRI analysis

We performed FMRI general linear model (GLM) analyses using SPM12. Similar to Behrens et al. (Behrens et al., 2008), we considered two GLM models, where the first looked for decision-related activity and the second for learning-related activity. In both GLMs, four sets of two regressors, each containing one regressor per cue type, were considered. These regressors were time-locked to the presentation of cues, advice, options and outcomes onscreen. In both GLMs, we also included 25 noise regressors obtained using the RETROICOR analysis of the pulse and respiration signals. Twelve motion regressors representing six motion parameters obtained from the brain-realignment procedure and their first derivative were also included.

In the decision-related GLM, two decision-variable parametric regressors, one based on reward information and the other one based on social information, were also included separately for each cue, both time-locked to presentation of options. The reward-based decision variable (RDV) was defined based on the interaction of p_t^r and the reward-related choice. Thus, RDV was equal to p_t^r and $1-p_t^r$ on trials in which the choice was the blue and the green option, respectively. The social decision variable (SDV) was defined based on the interaction of p_t^s and the advice-related choice. Thus, SDV was equal to p_t^s and $1-p_t^s$ on trials in which the subject chose to follow and not follow the advice, respectively.

In the learning-related GLM, four parametric regressors were included, which were all time-locked to the presentation of outcomes. Two error regressors, an uncertainty-weighted reward-based prediction error (RPE) regressor (in the choice frame) and an uncertainty-weighted social prediction error (SPE) regressor were included (the update term in Equation 1). These regressors were defined based on equation 1. Two other learning regressors encoding volatility estimate were also included separately for each cue.

We used a common set of (group-level) parameters obtained by fitting a variant of M1 to data using the HBI to generate parametric regressors for the above GLM FMRI analyses. The only difference between this model and M1 was the assumption that the weights are the same for both advisers. This is critical because it ensures that any difference in fMRI data is not simply due to having regressors generated based on different weights. Furthermore, any between-subject variability in coefficients of GLM FMRI cannot be attributed to parameters (correlation analysis of those coefficients and model parameters is valid and not circular). We used anatomical masks for region of interest (ROI) analyses according to recent connectivity-based parcellation studies(Mars et al., 2012; Neubert et al., 2015). Clusters 1 and 2 found by Mars et al. (Mars et al., 2012) was used as the TPJ mask (after smoothing with a 6mm kernel to avoid sharp edges). Cluster 4 of (Neubert et al., 2015) was used as the vmPFC mask.

	v_0	λ_r	ω_r	λ_s	ω_s	β	b	w_{ap}	w_{av}
Hierarchical fit	1.82	0.05	5.28	0.62	0.3	1.39	0.95	0.64	0.64

Supplementary Table 2. Parameters used for generating regressors for fMRI analysis based on a variant of M1, which has equal weight parameters for both advisers (i.e. $w_{ap} = w_{av}$).

References

Baldwin DA, Baird JA (2001) Discerning intentions in dynamic human action. Trends Cogn Sci 5:171–178.

Behrens TEJ, Hunt LT, Rushworth MFS (2009) The computation of social behavior. Science 324:1160–1164.

Behrens TEJ, Hunt LT, Woolrich MW, Rushworth MFS (2008) Associative learning of social value. Nature 456:245–249.

Biele G, Rieskamp J, Krugel LK, Heekeren HR (2011) The Neural Basis of Following Advice. PLoS Biol 9 Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3119653/ [Accessed November 30, 2019].

Birn RM, Diamond JB, Smith MA, Bandettini PA (2006) Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. NeuroImage 31:1536–1548.

Boyd R, Richerson PJ (2009) Culture and the evolution of human cooperation. Philos Trans R Soc B Biol Sci 364:3281–3288.

Carter RM, Bowling DL, Reeck C, Huettel SA (2012) A distinct role of the temporal-parietal junction in predicting socially guided decisions. Science 337:109–111.

Chib VS, Adachi R, O'Doherty JP (2018) Neural substrates of social facilitation effects on incentive-based performance. Soc Cogn Affect Neurosci 13:391–403.

Chung D, Christopoulos GI, King-Casas B, Ball SB, Chiu PH (2015) Social signals of safety and risk confer utility and have asymmetric effects on observers' choices. Nat Neurosci 18:912–916.

Cohen JD, Daw N, Engelhardt B, Hasson U, Li K, Niv Y, Norman KA, Pillow J, Ramadge PJ, Turk-Browne NB, Willke TL (2017) Computational approaches to fMRI analysis. Nat Neurosci 20:304–313.

Cook JL, den Ouden HEM, Heyes CM, Cools R (2014) The social dominance paradox. Curr Biol CB 24:2812–2816.

Cukier K, Mayer-Schönberger V, Véricourt F de (2021) Framers: Human Advantage in an Age of Technology and Turmoil. New York, New York: Dutton.

Daw ND, Doya K (2006) The computational neurobiology of learning and reward. Curr Opin Neurobiol 16:199–204.

De Martino B, Kumaran D, Seymour B, Dolan RJ (2006) Frames, Biases, and Rational Decision-Making in the Human Brain. Science 313:684–687.

Diaconescu AO, Mathys C, Weber LAE, Daunizeau J, Kasper L, Lomakina EI, Fehr E, Stephan KE (2014) Inferring on the Intentions of Others by Hierarchical Bayesian Learning. PLoS Comput Biol 10 Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4154656/ [Accessed May 6, 2019].

Frith CD, Singer T (2008) The role of social cognition in decision making. Philos Trans R Soc B Biol Sci 363:3875–3886.

Glover GH, Li TQ, Ress D (2000) Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magn Reson Med 44:162–167.

Grice P (1975) Logic and conversation. In: Syntax and semantics, pp 41–58. New York: Academic Press.

Guitart-Masip M, Huys QJM, Fuentemilla L, Dayan P, Duzel E, Dolan RJ (2012) Go and no-go learning in reward and punishment: interactions between affect and effect. NeuroImage 62:154–166.

Güroğlu B, van den Bos W, Rombouts SARB, Crone EA (2010) Unfair? It depends: neural correlates of fairness in social context. Soc Cogn Affect Neurosci 5:414–423.

Heyes C (2016) Who Knows? Metacognitive Social Learning Strategies. Trends Cogn Sci 20:204–213.

Izuma K, Adolphs R (2013) Social manipulation of preference in the human brain. Neuron 78:563–573.

Kahneman D, Slovic P, Tversky A (1982) Judgment under Uncertainty: Heuristics and Biases, 1st ed. Cambridge University Press.

Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. Econometrica 47:263–291.

Klucharev V, Hytönen K, Rijpkema M, Smidts A, Fernández G (2009) Reinforcement learning signal predicts social conformity. Neuron 61:140–151.

Konovalov A, Hu J, Ruff CC (2018) Neurocomputational approaches to social behavior. Curr Opin Psychol 24:41–47.

Mars RB, Sallet J, Schüffelgen U, Jbabdi S, Toni I, Rushworth MFS (2012) Connectivity-based subdivisions of the human right "temporoparietal junction area": evidence for different areas participating in different cortical networks. Cereb Cortex N Y N 1991 22:1894–1903.

McClure SM, Berns GS, Montague PR (2003) Temporal prediction errors in a passive learning task activate human striatum. Neuron 38:339–346.

McKenzie CRM, Nelson JD (2003) What a speaker's choice of frame reveals: reference points, frame selection, and framing effects. Psychon Bull Rev 10:596–602.

Moeller S, Yacoub E, Olman CA, Auerbach E, Strupp J, Harel N, Uğurbil K (2010) Multiband Multislice GE-EPI at 7 Tesla, With 16-Fold Acceleration Using Partial Parallel Imaging With Application to High Spatial and Temporal Whole-Brain FMRI. Magn Reson Med Off J Soc Magn Reson Med Soc Magn Reson Med 63:1144–1153.

Morgan TJH, Rendell LE, Ehn M, Hoppitt W, Laland KN (2012) The evolutionary basis of human social learning. Proc R Soc B Biol Sci 279:653–662.

Neubert F-X, Mars RB, Sallet J, Rushworth MFS (2015) Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. Proc Natl Acad Sci U S A 112:E2695-2704.

Nicolle A, Klein-Flügge MC, Hunt LT, Vlaev I, Dolan RJ, Behrens TEJ (2012) An agent independent axis for executed and modeled choice in medial prefrontal cortex. Neuron 75:1114–1121.

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. Neuron 38:329–337.

Pessiglione M, Delgado MR (2015) The good, the bad and the brain: Neural correlates of appetitive and aversive values underlying decision making. Curr Opin Behav Sci 5:78–84.

Piray P, Daw ND (2020) A simple model for learning in volatile environments. PLoS Comput Biol 16:e1007963.

Piray P, Dezfouli A, Heskes T, Frank MJ, Daw ND (2019a) Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies. PLOS Comput Biol 15:e1007043.

Piray P, Ly V, Roelofs K, Cools R, Toni I (2019b) Emotionally Aversive Cues Suppress Neural Systems Underlying Optimal Learning in Socially Anxious Individuals. J Neurosci 39:1445–1456.

Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. Nat Rev Neurosci 9:545–556.

Reiter A, Diaconescu A, Li S-C, Eppinger B (2017) 872. Mechanisms of Peer Influence on Decision-Making in Adolescence. Biol Psychiatry 81:S352–S353.

Ruff CC, Fehr E (2014) The neurobiology of rewards and values in social decision making. Nat Rev Neurosci 15:549–562.

Salali GD, Chaudhary N, Thompson J, Grace OM, van der Burgt XM, Dyble M, Page AE, Smith D, Lewis J, Mace R, Vinicius L, Migliano AB (2016) Knowledge-Sharing Networks in Hunter-Gatherers and the Evolution of Cumulative Culture. Curr Biol CB 26:2516–2521.

Salimi-Khorshidi G, Douaud G, Beckmann CF, Glasser MF, Griffanti L, Smith SM (2014) Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. NeuroImage 90:449–468.

Sanfey AG, Rilling JK, Aronson JA, Nystrom LE, Cohen JD (2003) The Neural Basis of Economic Decision-Making in the Ultimatum Game. Science 300:1755–1758.

Saxe R (2006) Uniquely human social cognition. Curr Opin Neurobiol 16:235–239.

Sher S, McKenzie CRM (2006) Information leakage from logically equivalent frames. Cognition 101:467–494.

Shmueli K, van Gelderen P, de Zwart JA, Horovitz SG, Fukunaga M, Jansma JM, Duyn JH (2007) Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. NeuroImage 38:306–320.

Siegal M, Varley R (2002) Neural systems involved in "theory of mind." Nat Rev Neurosci 3:463–471.

Stolk A, D'Imperio D, di Pellegrino G, Toni I (2015) Altered communicative decisions following ventromedial prefrontal lesions. Curr Biol CB 25:1469–1474.

Stolk A, Verhagen L, Schoffelen J-M, Oostenveld R, Blokpoel M, Hagoort P, van Rooij I, Toni I (2013) Neural mechanisms of communicative innovation. Proc Natl Acad Sci U S A 110:14574–14579.

Tabibnia G, Satpute AB, Lieberman MD (2008) The sunny side of fairness: preference for fairness activates reward circuitry (and disregarding unfairness activates self-control circuitry). Psychol Sci 19:339–347.

Toelch U, Bruce MJ, Newson L, Richerson PJ, Reader SM (2014) Individual consistency and flexibility in human social information use. Proc R Soc B Biol Sci 281:20132864.

van Buuren M, Gladwin TE, Zandbelt BB, van den Heuvel M, Ramsey NF, Kahn RS, Vink M (2009) Cardiorespiratory effects on default-mode network activity as measured with fMRI. Hum Brain Mapp 30:3031–3042.

Zaki J, López G, Mitchell JP (2014) Activity in ventromedial prefrontal cortex co-varies with revealed social preferences: evidence for person-invariant value. Soc Cogn Affect Neurosci 9:464–469.